
Beyond Computer Ethics

A reader for ECS 188
University of California, Davis
compiled by Phillip Rogaway
March 29, 2009

Contents

0. A Brief Note to the Student (P. Rogaway)	5
1. Views of Technology (I. Barbour)	7
2. Do Machines Make History? (R. Heilbroner)	21
3. Do Artifacts Have Politics? (L. Winner)	32
4. Five Things We Need to Know About Technological Change (N. Postman)	44
5. McLuhan Interview (M. McLuhan)	49
7. Technology and Social Justice (F. Dyson)	65
8. Why the Future Doesn't Need Us (B. Joy)	75
9. The Technological Subversion of Environmental Ethics (D. Strong)	94
10. Philosophical Ethics (D. Johnson)	109
11. The Altered Nature of Human Action (H. Jonas)	126
12. Industrial Society and Technological Systems (R. Schwartz)	140
13. The Lexus and the Olive Tree (T. Friedman)	154
14. A Road Map for Natural Capitalism (A. Lovins, L. Lovins, and G. Hawken)	167
15. The Tragedy of the Commons (G. Hardin)	183
16. The World as a Polder (J. Diamond)	191
17. War (B. Orend)	211
18. Computers, Ethics, and Collective Violence (C. Summers and E. Markusen)	225
19. Farewell Address to the Nation (D. Eisenhower)	239
20. Fencing Off Ideas: Enclosure & The Disappearance of the Public Domain (J. Boyle)	243
21. The GNU Manifesto (R. Stallman)	253
22. The Darknet and the Future of Content Distribution (P. Biddle <i>et al.</i>)	261
23. Microsoft Research DRM Talk (C. Doctorow)	277
24. Bigger Monster, Weaker Chains (the ACLU Technology and Liberty Program)	289
27. Bhopal Lives (S. Mehta)	311
28. Medical Devices: The Therac-25 (N. Leveson)	321
29. The Fifty-Nine Story Crisis (J. Morgenstern)	346
30.1 ACM Code of Ethics (the ACM)	357
30.2 (IEEE Code of Ethics)	362
30.3 Software Engineering Code of Ethics (the IEEE Computer Society and the ACM)	364
31. Unlocking the Clubhouse: Women and Computing (J. Margolis and A. Fisher)	369
32. The Future of Our Profession (B. Dahlbom and L. Mathiassen)	383
33. Disciplined Minds (J. Schmidt)	393
34. Some Pledges (various authors)	410

A Brief Note to the Student

Throughout your university studies, and indeed throughout your lives in the USA, you've been quietly inculcated with a rather specific world view concerning technology, ethics, and society. Some of its tenants are that technology is about the gadgets that we build and use; that it's an outgrowth of the sciences; that, overwhelmingly, technology makes things better; that it liberates us, empowers us, and helps everyone to prosper; and that our technology is fundamentally apolitical, areligious, and amoral. The same world view holds that the individual is the primary agent that drives technological change, as well as the locus of responsibility for that change. Correspondingly, the individual scientist or engineer behaves ethically and appropriately when he abides by the law, by professional standards, and by cultural norms.

It is all, I am afraid, far more more false than true. It is a kind of fiction we have spun to make it easier to do the things we do. At the same time, the viewpoint I have sketched is so much a part of our culture, our institutions, and our selves that we can hardly even see it's there. And, now, this blindness imperils our very existence.

Many of the readings I've assembled here are intended to push you, at least a bit, to question assumptions like those above. Hopefully one or two of them will do their job.

At the end of the term, when you do your course evaluations, the statement of course goals will say, in part, that I wanted you to think about, and act upon, the ethical implications of your personal and professional choices, and our collective work as technologists. At one level, this may sound kind of easy, perhaps like something you've always done. But, in fact, I suspect it's rarely done, and a terribly hard thing to do. Overturning this rock reveals a world both difficult to understand and uncomfortable to see.

A colleague once commented that he had never met anyone who regarded his own behavior as anything but proper and good. And yet, collectively, it seems to me that we are routinely committing a massive amount of wrong. Is it really possible that we could each behave well and yet, somehow, our collective behavior should end up so rank? I will leave you to ponder your own answer to this riddle, and close by wishing you wisdom—certainly more than I have ever found—in your own struggles with the issues of this note and of this course.

Kind regards,



Davis, California, USA
April 2008

Redacted (PR) from Chapter 1 of *Ethics in an Age of Technology: The Gifford Lectures 1989-1991, Volume 2*, Harper San Francisco. December 25, 1992.

Views of Technology

by

Ian G. Barbour

Technology, the source of the problem, will once again prove to contain within itself the germs of a solution compatible with the betterment of man's lot and dignity.

CHARLES SUSSKIND¹

Our enslavement to the machine has never been more complete.

JOHN ZERMAN AND ALICE CARNES²

What we call Man's power over Nature turns out to be a power exercised by some men over other men with Nature as its instrument.

C. S. LEWIS³

Appraisals of modern technology diverge widely. Some see it as the beneficent source of higher living standards, improved health, and better communications. They claim that any problems created by technology are themselves amenable to technological solutions. Others are critical of technology, holding that it leads to alienation from nature, environmental destruction, the mechanization of human life, and the loss of human freedom. A third group asserts that technology is ambiguous, its impacts varying according to the social context in which it is designed and used, because it is both a product and a source of economic and political power.⁴

In this chapter, views of technology are grouped under three headings: Technology as Liberator, Technology as Threat, and Technology as Instrument of Power. In each case the underlying assumptions and value judgments are examined. I will indicate why I agree with the third of these positions, which emphasizes the social construction and use of particular technologies. The issues cut across disciplines; I draw from the writings of engineers, historians, sociologists, political scientists, philosophers, and theologians. The human and environmental values relevant to the appraisal of technology are further analyzed in chapters 2 and 3. These three chapters provide the ethical categories and principles for examining policy decisions about particular technologies in later chapters.

Technology may be defined as *the application of organized knowledge to practical tasks by ordered systems of people and machines*.⁵ There are several advantages to such a broad definition. "Organized knowledge" allows us to include technologies based on practical experience and invention as well as those based on scientific theories. The "practical tasks" can include both the production of mater (in industry and agriculture, for instance) and the provision of services (by computers, communications media, and biotechnologies, among others). Reference to "ordered systems of people and machines" directs attention to social institutions as well as to the hardware of technology. The breadth of the definition also reminds us that there are major differences among technologies.

1. TECHNOLOGY AS LIBERATOR

Throughout modern history, technological developments have been enthusiastically welcomed because of their potential for liberating us from hunger, disease, and poverty. Technology has been celebrated as the source of material progress and human fulfillment.

1. THE BENEFITS OF TECHNOLOGY

Defenders of technology point out that four kinds of benefits can be distinguished if one looks at its recent history

and considers its future:

1. *Higher Living Standards.* New drugs, better medical attention, and improved sanitation and nutrition have more than doubled the average life span in industrial nations within the past century. Machines have released us from much of the backbreaking labor that in previous ages absorbed most of people's time and energy. Material progress represents liberation from the tyranny of nature. The ancient dream of a life free from famine and disease is beginning to be realized through technology. The standard of living of low-income families in industrial societies has doubled in a generation, even though relative incomes have changed little. Many people in developing nations now look on technology as their principal source of hope. Productivity and economic growth, it is said, benefit everyone in the long run.

2. *Opportunity for Choice.* Individual choice has a wider scope today than ever before because technology has produced new options riot previously available arid a greater range of products and services. Social and geographical mobility allow a greater choice of jobs and locations. In air urban industrial society, a person's options are not as limited by parental or community expectations as they were in a small-town agrarian society. The dynamism of technology can liberate people from static and confining traditions to assume responsibility for their own lives. Birth control techniques, for example, allow a couple to choose the size and tinning of their family. Power over nature gives greater opportunities for the exercise of human freedom.⁶

3. *More Leisure.* Increases in productivity have led to shorter working hours. Computers and automation hold the promise of eliminating much of the monotonous work typical of earlier industrialism. Through most of history, leisure and cultural pursuits have been the privilege of the few, while the mass of humanity was preoccupied with survival. In an affluent society there is time for continuing education, the arts, social service, sports, and participation in community life. Technology can contribute to the enrichment of human life and the flowering of creativity. Laborsaving devices free as to do what machines cannot do. Proponents of this viewpoint say that people can move be

4. *Improved Communications.* With new forms of transportation, one can in a few hours travel to distant cities that once took months to reach. With electronic technologies (radio, television, computer networks, and so on), the speed, range, and scope of communication have vastly increased. The combination of visual image and auditory message have an immediacy not found in the linear sequence of the printed word. These new media offer the possibility of instant worldwide communication, greater interaction, understanding, and mutual appreciation in the "global village." It has been suggested that by dialing coded numbers on telephones hooked into computer networks, citizens could participate in an instant referendum on political issues. According to its defenders, technology brings psychological and social benefits as well as material progress.

In part 2 we will encounter optimistic forecasts of each of the particular technologies examined. In agriculture, some experts anticipate that the continuing Green Revolution and the genetic engineering of new crops will provide adequate food for a growing world population. In the case of energy, it is claimed that breeder reactors and fusion will provide environmentally benign power to replace fossil fuels. Computer enthusiasts anticipate the Information Age in which industry is automated and communications networks enhance commercial, professional, and personal life. Biotechnology promises the eradication of genetic diseases, the improvement of health, and the deliberate design of new species—even the modification of humanity itself. In subsequent chapters we will examine each of these specific claims as well as the general attitudes they reveal.

2. OPTIMISTIC VIEWS OF TECHNOLOGY

Let us look at some authors who have expressed optimism regarding technology. Melvin Kranzberg, a prominent historian of technology, has presented a very positive picture of the technological past and future. He argues that urban industrial societies offer *more freedom* than rural ones and provide greater choice of occupations, friends, activities, and life-styles. The work week has been cut in half, and human wants have been dramatically fulfilled.⁷ Emanuel Mesthene, former director of the Harvard Program in Technology and Society, grants that every technology brings risks as well as benefits, but he says that our task is the rational management of risk. Some technologies poison the environment, but others reduce pollution. A new technology may displace some workers but it also creates new jobs. Nineteenth-century factories and twentieth-century assembly lines did involve dirty and monotonous work, but the newer technologies allow greater creativity and individuality.⁸

A *postindustrial society*, it is said, is already beginning to emerge. In this new society, according to the sociologist Daniel Bell, power will be based on knowledge rather than property. The dominant class will be scientists, engineers, and technical experts; the dominant institutions will be intellectual ones (universities, industrial laboratories, and research institutes). The economy will be devoted mainly to services rather than material goods.

Decisions will be made on rational-technical grounds, marking “the end of ideology.” There will be a general consensus on social values; experts will coordinate social planning, using rational techniques such as decision theory and systems analysis. This will be a fixture-oriented society, the age of the professional managers, the technocrats.⁹ A bright picture of the coming technological society has been given by many “futurists,” including Buckminster Fuller, Herman Kahn, and Alvin Toffler.¹⁰

Samuel Florman is an articulate engineer and author who has written extensively *defending* technology against its detractors. He insists that the critics have romanticized the life of earlier centuries and rural societies. Living standards were actually very low, work was brutal, and roles were rigidly defined. People have much greater freedom in technological societies. The automobile, for example, enables people to do what they want and enhances geographical and class mobility. People move to cities because they prefer life there to “the tedium and squalor of the countryside.” Florman says that worker alienation in industry is rare, and many people prefer the comfortable monotony of routine tasks to the pressures of decision and accountability. Technology is not an independent force out of control: it is the product of human choice, a response to public demand expressed through the marketplace.¹¹

Florman grants that technology often has undesirable side effects, but he says that these are amenable to *technological solutions*. One of his heroes is Benjamin Franklin, who “proposed technological ways of coping with the unpleasant consequences of technology.”¹² Florman holds that environmental and health risks are inherent in every technical advance. Any product or process can be made safer, but always at an economic cost. Economic growth and lower prices for consumers are often more important than additional safety, and absolute safety is an illusory goal. Large-scale systems are usually more efficient than small-scale ones. It is often easier to find a “technical fix” for a social problem than to try to change human behavior or get agreement on political policies.¹³

Florman urges us to rely on *the judgment of experts* in decisions about technology. He says that no citizen can be adequately informed about complex technical questions such as acid rain or radioactive waste disposal. Public discussion of these issues only leads to anxiety and erratic political actions. We should rely on the recommendations of experts on such matters.¹⁴ Florman extols the “unquenchable spirit” and “irrepressible human will” evident in technology:

For all our apprehensions, we have no choice but to press ahead. We must do so, first, in the name of compassion. By turning our backs on technological change, we would be expressing our satisfaction with current world levels of hunger, disease, and privation. Further, we must press ahead in the name of the human adventure. Without experimentation and change our existence would be a dull business. We simply cannot stop while there are masses to feed and diseases to conquer, seas to explore and heavens to survey.¹⁵

Some theologians have also given very positive appraisals of technology. They see it as a source not only of higher living standards but also of *greater freedom and creative expression*. In his earlier writings, Harvey Cox held that freedom to master and shape the world through technology liberates us from the confines of tradition. Christianity brought about the desecralization of nature and allowed it to be controlled and used for human welfare.¹⁶ Norris Clarke sees technology as an instrument of human fulfillment and self-expression in the use of our God-given intelligence to transform the world. Liberation from bondage to nature, he says, is the victory of spirit over matter. As cocreators with God we can celebrate the contribution of reason to the enrichment of human life.¹⁷ Other theologians have affirmed technology as an instrument of love and compassion in relieving human suffering—a modern response to the biblical command to feed the hungry and help the neighbor in need.

The Jesuit paleontologist Pierre Teilhard de Chardin, writing in the early years of nuclear power, computers, and molecular biology, expressed a *hopeful vision of the technological future*. He envisioned computers and electronic communication in a network of interconnected consciousness, a global layer of thought that he called “the noosphere.” He defended eugenics, “artificial neo-life,” and the remodeling of the human organism by manipulation of the genes. With this new power over heredity, he said, we can replace the crude forces of natural selection and “seize the tiller” to control the direction of future evolution. We will have total power over matter, “reconstructing the very stuff of the universe.” He looked to a day of interplanetary travel and the unification of our own planet, based on intellectual and cultural interaction.¹⁷

Teilhard’s writings present us with a magnificent sweep of time from past to future. But they do not consider the institutional structures of economic power and self-interest that now control the directions of technological development. Teilhard seldom acknowledged the tragic hold of social injustice on human life. He was writing before the destructive environmental impacts of technology were evident. When Teilhard looked to the past, he portrayed humanity as an integral part of the natural world, interdependent with other creatures. But when he looked to the future, he expected that because of our technology and our spirituality we will be increasingly separated from other creatures. Humanity will move beyond dependence on the organic world. Though he was

ultimately theocentric (centered on God), and he talked about the redemption of the whole cosmos, many of his images are anthropocentric (centered on humanity) and imply that other forms of life are left behind in the spiritualization of humankind that technology will help to bring about.

3. A REPLAY TO THE OPTIMISTS

First, *the environmental costs and human risks* of technology are dismissed too rapidly. The optimists are confident that technical solutions can be found for environmental problems. Of course, pollution abatement technologies can treat many of the effluents of industry, but often unexpected, indirect, or delayed consequences occur. The effects of carcinogens may not show up for twenty-five years or more. The increased death rates among shipyard workers exposed to asbestos in the early 1940s were not evident until the late 1960s. Toxic wastes may contaminate groundwater decades after they have been buried. The hole in the ozone layer caused by the release of chlorofluorocarbons had not been anticipated by any scientists. Above all, soil erosion and massive deforestation threaten the biological resources essential for human life, and global warming from our use of fossil fuels threatens devastating changes in world climates.

Second, environmental destruction is symptomatic of a deeper problem: *alienation from nature*. The idea of human domination of nature has many roots. Western religious traditions have often drawn a sharp line between humanity and other creatures (see chapter 3). Economic institutions treat nature as a resource for human exploitation. But technological enthusiasts contribute to this devaluation of the natural world if they view it as an object to be controlled and manipulated. Many engineers are trained in the physical sciences and interpret living things in mechanistic rather than ecological terms. Others spend their entire professional lives in the technosphere of artifacts, machines, electronics, and computers, cut off from the world of nature. To be sure, sensitivity to nature is sometimes found among technological optimists, but it is more frequently found among the critics of technology.

Third, technology has contributed to the *concentration of economic and political power*. Only relatively affluent groups or nations can afford the latest technology; the gaps between rich and poor have been perpetuated and in many cases increased by technological developments. In its world of limited resources, it also appears impossible for all nations to sustain the standards of living of industrial nations today, much less the higher standards that industrial nations expect in the future. Affluent nations use a grossly disproportionate share of the world's energy and resources. Commitment to justice within nations also requires a more serious analysis of the distribution of the costs and benefits of technology. We will find many technologies in which one group enjoys the benefits while another group is exposed to the risks and social costs.

Fourth, *large-scale technologies* typical of industrial nations today are particularly problematic. They are capital-intensive rather than labor-intensive, and they add to unemployment in many parts of the world. Large-scale systems tend to be vulnerable to error, accident, or sabotage. The near catastrophe at the Three Mile Island nuclear plant in 1979 and the Chernobyl disaster in 1986 were the products of human errors, faulty equipment, poor design, and unreliable safety procedures. Nuclear energy is a prime example of a vulnerable, centralized, capital-intensive technology. Systems in which human or mechanical failures can be disastrous are risky even in a stable society, quite apart from additional risks under conditions of social unrest. The large scale of many current systems is as much the product of government subsidies, tax and credit policies, and particular corporate interests as of any inherent economies of scale.

Fifth, greater *dependence on experts* for policy decisions would not be desirable. The technocrats claim that their judgments are value free; the technical elite is supposedly nonpolitical. But those with power seldom use it rationally and objectively when their own interests are at stake. When social planners think they are deciding for the good of all—whether in the French or Russian revolution or in the proposed technocracy of the future—the assumed innocence of moral intentions is likely to be corrupted in practice. Social controls over the controllers are always essential. I will suggest that the most important form of freedom is participation in the decisions affecting our lives.

Lastly, we must question the linear view of the *science-technology-society relationship*, which is assumed by many proponents of optimistic views. Technology is taken to be applied science, and it is thought to have an essentially one-way impact on society. The official slogan of the Century of Progress exposition in Chicago in 1933 was: “Science Finds—Industry Applies—Man Conforms.” This has been called “the assembly-line view” because it pictures science at the start of the line and its stream of technological products pouring off the end of the line.¹⁹ If technology is fundamentally benign, there is no need for government interference except to regulate the most serious risks. Whatever guidance is needed for technological development is supplied by the expression of consumer preferences through the marketplace. In this view, technologies develop from the “push” of science and the “pull” of economic profits.

II. TECHNOLOGY AS THREAT

At the opposite extreme are the critics of modern technology who see it as a threat to authentic human life. We will confine ourselves here to criticisms of the human rather than environmental consequences of technology.

1. THE HUMAN COSTS OF TECHNOLOGY

Five characteristics of industrial technology seem to its critics particularly inimical to human fulfillment.²⁰

1. *Uniformity in a Mass Society.* Mass production yields standardized products, and mass media tend to produce a uniform national culture. Individuality is lost and local or regional differences are obliterated in the homogeneity of industrialization. Nonconformity hinders efficiency, so cooperative and docile workers are rewarded. Even the interactions among people are mechanized and objectified. Human identity is defined by roles in organizations. Conformity to a mass society jeopardizes spontaneity and freedom. According to the critics, there is little evidence that an electronic, computerized, automated society will produce more diversity than earlier industrialism did.

2. *Narrow Criteria of Efficiency.* Technology leads to rational and efficient organization, which requires fragmentation, specialization, speed, the maximization of output. The criterion is efficiency in achieving a single goal or a narrow range of objectives; side effects and human costs are ignored. Quantitative criteria tend to crowd out qualitative ones. The worker becomes the servant of the machine, adjusting to its schedule and tempo, adapting to its requirements. Meaningful work roles exist for only a small number of people in industrial societies today. Advertising creates demand for new products, whether or not they fill real needs, in order to stimulate a larger volume of production and a consumer society.

3. *Impersonality and Manipulation.* Relationships in a technological society are specialized and functional. Genuine community and interpersonal interaction are threatened when people feel like cogs in a well-oiled machine. In a bureaucracy, the goals of the organization are paramount and responsibility is diffused, so that no one feels personally responsible. Moreover, technology has created subtle ways of manipulating people and new techniques of electronic surveillance and psychological conditioning. When the technological mentality is dominant, people are viewed and treated like objects.

4. *Uncontrollability.* Separate technologies form an interlocking system, a total, mutually reinforcing network that seems to lead a life of its own. "Runaway, technology" is said to be like a vehicle out of control, with a momentum that cannot be stopped. Some critics assert that technology is not just a set of adaptable tools for human use but an all-encompassing form of life, a pervasive structure with its own logic and dynamic. Its consequences are unintended and unforeseeable. Like the sorcerer's apprentice who found the magic formula to make his broom carry water but did not know how to make it stop, we have set in motion forces that we cannot control. The individual feels powerless facing a monolithic system.

5. *Alienation of the Worker.* The worker's alienation was a central theme in the writing of Karl Marx. Under capitalism, he said, workers do not own their own tools or machines, and they are powerless in their work life. They can sell their labor as a commodity, but their work is not a meaningful form of self-expression. Marx held that such alienation is a product of capitalist ownership and would disappear under state ownership. He was optimistic about the use of technology in a communist economic order, and thus he belongs with the third group below, the contextualists, but his idea of alienation has influenced the pessimists.

More recent writers point out that alienation has been common in state-managed industrial economies too and seems to be a product of the division of labor, rationalization of production, and hierarchical management in large organizations, regardless of the economic system. Studs Terkel and others have found in interviews that resentment, frustration, and a sense of powerlessness are widespread among American industrial workers. This contrasts strongly with the greater work autonomy, job satisfaction, and commitment to work found in the professions, skilled trades, and family-owned farms.²¹ Other features of technological development since World War II have evoked widespread concern. The allocation of more than two-thirds of the U.S. federal research and development budget to military purposes has diverted expertise from environmental problems and urgent human needs. Technology also seems to have contributed to the impoverishment of human relationships and a loss of community. The youth counterculture of the 1970s was critical of technology and sought harmony with nature, intensity of personal experience, supportive communities, and alternative life-styles apart from the prevailing industrial order. While many of its expressions were short-lived, many of its characteristic attitudes, including

disillusionment with technology, have persisted among some of the younger generation.²²

2. RECENT CRITICS OF TECHNOLOGY

To the French philosopher and social critic Jacques Ellul, technology is *an autonomous and uncontrollable force* that dehumanizes all that it touches. The enemy is “technique”—a broad term Ellul uses to refer to the technological mentality and structure that he sees pervading not only industrial processes, but also all social, political, and economic life affected by them. Efficiency and organization, he says, are sought in all activities. The machine enslaves people when they adapt to its demands. Technology has its own inherent logic and inner necessity. Rational order is everywhere imposed at the expense of spontaneity and freedom.

Ellul ends with a *technological determinism*, since technique is self-perpetuating, all-pervasive, and inescapable. Any opposition is simply absorbed as we become addicted to the products of technology. Public opinion and the state become the servants of technique rather than its masters. Technique is global, monolithic, and unvarying among diverse regions and nations. Ellul offers us no way out, since all our institutions, the media, and our personal lives are totally in its grip. He holds that biblical ethics can provide a viewpoint transcending society from which to judge the sinfulness of the technological order and can give us the motivation to revolt against it, but he holds out little hope of controlling it.²³ Some interpreters see in Ellul’s recent writings a very guarded hope that a radical Christian freedom that rejects cultural illusions of technological progress might in the long run lead to the transformation rather than the rejection of technology. But Ellul does not spell out such a transformation because he holds that the outcome is in God’s hands, not ours, and most of his writings are extremely pessimistic about social change.²⁴

The political scientist Langdon Winner has given a sophisticated version of the argument that technology is *an autonomous system* that shapes all human activities to its own requirements. It makes little difference who is nominally in control—elected politicians, technical experts, capitalist executives, or socialist managers—if decisions are determined by the demands of the technical system. Human ends are then adapted to suit the techniques available rather than the reverse. Winner says that large-scale systems are self-perpetuating, extending their control over resources and markets and molding human life to fit their own smooth functioning. Technology is not a neutral means to human ends but an all-encompassing system that imposes its patterns on every aspect of life and thought.²⁵

The philosopher Hans Jonas is impressed by *the new scale of technological power* and its influence on events distant in time and place. Traditional Western ethics have been anthropocentric and have considered only short-range consequences. Technological change has its own momentum, and its pace is too rapid for trial-and-error readjustments. Now genetics gives us power over humanity itself. Jonas calls for a new ethic of responsibility for the human future and for nonhuman nature. We should err on the side of caution, adopting policies designed to avert catastrophe rather than to maximize short-run benefits. “The magnitude of these stakes, taken with the insufficiency of our predictive knowledge, leads to the pragmatic rule to give the prophecy of doom priority over the prophecy of bliss.”²⁶ We should seek “the least harm,” not “the greatest good.” We have no right to tamper genetically with human nature or to accept policies that entail even the remote possibility of the extinction of humanity in a nuclear holocaust.

Another philosopher, Albert Borgmann, does not want to return to a pretechnological past, but he urges the selection of technologies that encourage *genuine human fulfillment*. Building on the ideas of Heidegger, he holds that authentic human existence requires the engagement and depth that occur when simple things and practices focus our attention and center our lives. We have let technology define the good life in terms of production and consumption, and we have ended with mindless labor and mindless leisure. A fast-food restaurant replaces the family meal, which was an occasion of communication and celebration. The simple pleasures of making music, hiking and running, gathering with friends around the hearth, or engaging in creative and self-reliant work should be our goals. Borgmann thinks that some large-scale capital-intensive industry is needed (especially in transportation and communication), but he urges the development of small-scale labor-intensive, locally owned enterprises (in arts and crafts, health care, and education, for example). We should challenge the rule of technology and restrict it to the limited role of supporting the humanly meaningful activities associated with a simpler life.²⁷

In *Technology and Power*, the psychologist David Kipnis maintains that those who control a technology have power over other people and this affects personal attitudes as well as social structures. Power holders interpret technological superiority as moral superiority and tend to look down on weaker parties. Kipnis shows that military and transportation technologies fed the conviction of colonists that they were superior to colonized peoples. Similarly, medical knowledge and specialization have led doctors to treat patients as impersonal cases and to keep patients at arms length with a minimum of personal communication. Automation gave engineers and managers

increased power over workers, who no longer needed special skills. In general, “power corrupts” and leads people to rationalize their use of power for their own ends. Kipnis claims that the person with technological knowledge often has not only a potent instrument of control but also a self-image that assumes superiority over people who lack that knowledge and the concomitant opportunities to make decisions affecting their lives.²⁸

Some Christian groups are critical of the *impact of technology on, human life*. The Amish, for example, have resolutely turned their backs on radios, television, and even automobiles. By hard work, community cooperation, and frugal ways, they have prospered in agriculture and have continued their distinctive life styles and educational patterns. Many theologians who do not totally reject technology criticize its tendency to generate a Promethean pride and its quest for unlimited power. The search for omnipotence is a denial of creaturehood. Unqualified devotion to technology as a total way of life, they say, is a form of idolatry. Technology is finally thought of as the source of salvation, the agent of secularized redemption.²⁹ In an affluent society, a legitimate concern for material progress readily becomes a frantic pursuit of comfort, a total dedication to self-gratification. Such an obsession with things distorts our basic values as well as our relationships with other persons. Exclusive dependence on technological rationality also leads to a truncation of experience, a loss of imaginative and emotional life, and an impoverishment of personal existence.

Technology is *imperialistic and addictive*, according to these critics. The optimists may think that, by fulfilling our material needs, technology liberates us from materialism and allows us to turn to intellectual, artistic, and spiritual pursuits. But it does not seem to be working out that way. Our material wants have escalated and appear insatiable. Yesterday’s luxuries are today’s necessities. The rich are usually more anxious about their future than the poor. Once we allow technology to define the good life, we have excluded many important human values from consideration.

3. A REPLY TO THE PESSIMISTS

In replying to these authors, we may note first that there are *great variations among technologies*, which are ignored when they are lumped together and condemned wholesale. Computerized offices differ greatly from steel mills and auto assembly lines, even if they share some features in common. One survey of journal articles finds that philosophers and those historians who trace broad trends (in economic and urban history, for example) often claim that technology determines history, whereas the historians or sociologists who make detailed studies of particular technologies are usually aware of the diversity of social, political, and economic interests that affect the design of a machine and its uses.³⁴ I will maintain that the uses of any technology vary greatly depending on its social contexts. To be sure, technological systems are interlocked, but they do not form a monolithic system impervious to political influence or totally dominating all other social forces. In particular, technology assessment and legislation offer opportunities for controlling technology, as we shall see.

Second, technological pessimists neglect possible avenues for *the redirection of technology*. The “inevitability” or “inherent logic” of technological developments is not supported by historical studies. We will note below some cases in which there were competing technical designs and the choice among them was affected by various political and social factors. Technological determinism underestimates the diversity of forces that contribute to technological change. Unrelieved pessimism undercuts human action and becomes a self fulfilling prophecy. If we are convinced that nothing can be done to improve the system, we will indeed do nothing to try to improve it. This would give to the commercial sponsors of technology the choices that are ours as responsible citizens.

Third, technology can be the *servant of human values*. Life is indeed impoverished if the technological attitudes of mastery and power dominate one’s outlook. Calculation and control do exclude mutuality and receptivity in human relationships and prevent the humility and reverence that religious awareness requires. But I would submit that the threat to these areas of human existence comes not from technology itself but from preoccupation with material progress and unqualified reliance on technology. We can make decisions about technology within a wider context of human and environmental values.

III. TECHNOLOGY AS INSTRUMENT OF POWER

A third basic position holds that technology is neither inherently good nor inherently evil but is an ambiguous instrument of power whose consequences depend on its social context. Some technologies seem to be neutral if they can be used for good or evil according to the goals of the users. A knife can be used for surgery or for murder. An isotope separator can enrich uranium for peaceful nuclear reactors or for aggression with nuclear weapons. But

historical analysis suggests that most technologies are already molded by particular interests and institutional goals. Technologies are social constructions, and they are seldom neutral because particular purposes are already built into their design. Alternative purposes would lead to alternative designs. Yet most designs still allow some choice as to how they are deployed.

1. TECHNOLOGY AND POLITICAL POWER

Like the authors in the previous group, those in this group are critical of many features of current technology. But they offer hope that technology can be used for more humane ends, either by political measures for more effective guidance within existing institutions or by changes in the economic and political systems themselves.

The people who make most of the decisions about technology today are not a technical elite or technocrats trying to run society rationally or disinterested experts whose activity was supposed to mark “the end of ideology.” The decisions are made by managers dedicated to the *interests of institutions*, especially industrial corporations and government bureaucracies. The goals of research are determined largely by the goals of institutions: corporate profits, institutional growth, bureaucratic power, and so forth. Expertise serves the interests of organizations and only secondarily the welfare of people or the environment.

The interlocking structure of *technologically based government agencies and corporations*, sometimes called the “technocomplex,” is wider than the “military-industrial complex.” Many companies are virtually dependent on government contracts. The staff members of regulatory agencies, in turn, are mainly recruited from the industries they are supposed to regulate. Networks of industries with common interests form lobbies of immense political power. For example, U.S. legislation supporting railroads and public mass transit systems was blocked by a coalition of auto manufacturers, insurance companies, oil companies, labor unions, and the highway construction industry. But citizens can also influence the direction of technological development. Public opposition to nuclear power plants was as important as rising costs in stopping plans to construct new plants in almost all Western nations.

The historian Arnold Pacey gives many examples of *the management of technology for power and profit*. This is most clearly evident in the defense industries with their close ties to government agencies. But often the institutional biases associated with expertise are more subtle. Pacey gives as one example the Western experts in India and Bangladesh who in the 1960s advised the use of large drilling rigs and diesel pumps for wells, imported from the West. By 1975, two thirds of the pumps had broken down because the users lacked the skills and maintenance networks to operate them. Pacey calls for greater public participation and a more democratic distribution of power in the decisions affecting technology. He also urges the upgrading of indigenous technologies, the exploration of intermediate-scale processes, and greater dialogue between experts and users. Need-oriented values and local human benefits would then play a larger part in technological change.³⁵

2. THE REDIRECTION OF TECHNOLOGY

The political scientist Victor Ferkiss expresses hope about the redirection of technology. He thinks that both the optimists and the pessimists have neglected the diversity among different technologies and *the potential role of political structures* in reformulating policies. In the past, technology has been an instrument of profit, and decisions have been motivated by short-run private interests. Freedom understood individualistically became license for the economically powerful. Individual rights were given precedence over the common good, despite our increasing interdependence. Choices that could only be made and enforced collectively—such as laws concerning air and water pollution—were resisted as infringements on free enterprise. But Ferkiss thinks that economic criteria can be subordinated to such social criteria as ecological balance and human need. He believes it is possible to combine centralized, systemwide planning in basic decisions with decentralized implementation, cultural diversity, and citizen participation.³⁶

There is a considerable range of views among *contemporary Marxists*. Most share Marx’s conviction that technology is necessary for solving social problems but that under capitalism it has been an instrument of exploitation, repression, and dehumanization. In modern capitalism, according to Marxists, corporations dominate the government and political processes serve the interests of the ruling class. The technical elite likewise serves the profits of the owners. Marxists grant that absolute standards of living have risen for everyone under capitalist technology. But relative inequalities have increased, so that class distinctions and poverty amidst luxury remain. Marxists assign justice a higher priority than freedom. Clearly they blame capitalism rather than technology for these evils of modern industrialism. They believe that alienation and inequality will disappear and technology will

be wholly benign when the working class owns the means of production. The workers, not the technologists, are the agents of liberation. Marxists are thus as critical as the pessimists concerning the consequences of technology within capitalism but as enthusiastic as the optimists concerning its potentialities—within a proletarian economic order.

How, then, do Western Marxists view the human effects of *technology in Soviet history*? Reactions vary, but many would agree with Bernard Gendron that in the Soviet Union workers were as alienated, factories as hierarchically organized, experts as bureaucratic, and pollution and militarism as rampant as in the United States. But Gendron insists that the Soviet Union did not follow Marx’s vision. The means of production were controlled by a small group within the Communist party, not by the workers. Gendron maintains that in a truly democratic socialism, technology would be humane and work would not be alienating.³⁷ Most commentators hold that the demise of communism in Eastern Europe and the Soviet Union was a product of both its economic inefficiency and its political repression. It remains to be seen whether any distinctive legacy from Marxism will remain there after the economic and political turmoil of the early nineties.

3. THE SOCIAL CONSTRUCTION OF TECHNOLOGY

How are science, technology, and society related? Three views have been proposed (see Fig. 1).

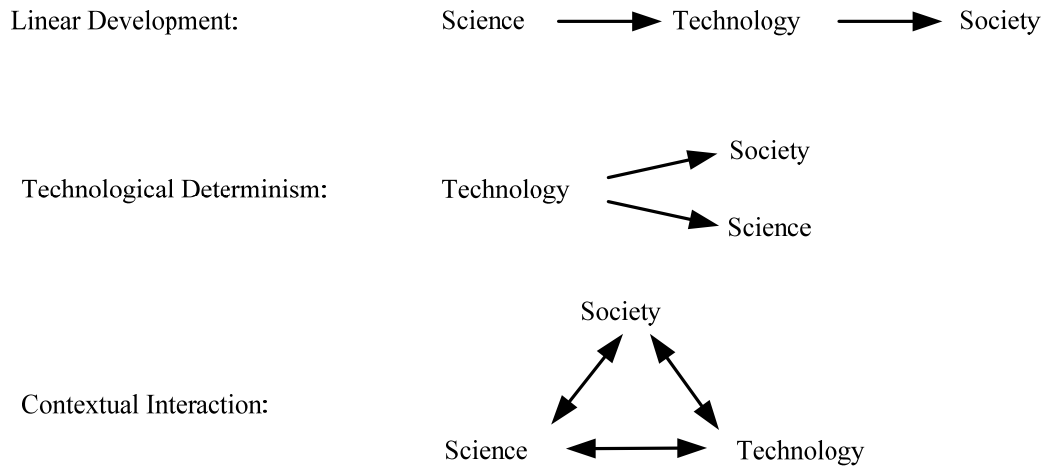


Fig. 1. Views of the Interaction of Science, Technology, and Society

1. *Linear development.* In linear development it is assumed that science leads to technology, which in turn has an essentially one-way impact on society. The deployment of technology is primarily a function of the marketplace. This view is common among the optimists. They consider technology to be predominantly beneficial, and therefore little government regulation or public policy choice is needed; consumers can influence technological development by expressing their preferences through the marketplace.

2. *Technological Determinism.* Several degrees and types of determinism can be distinguished. Strict determinism asserts that only one outcome is possible. A more qualified claim is that there are very strong tendencies present in technological systems, but these could be at least partly counteracted if enough people were committed to resisting them. Again, technology may be considered an autonomous interlocking system, which develops by its own inherent logic, extended to the control of social institutions. Or the more limited claim is made that the development and deployment of technology in capitalist societies follows only one path, but the outcomes might be different in other economic systems. In all these versions, science is itself driven primarily by technological needs. Technology is either the “independent variable” on which other variables are dependent, or it is the overwhelmingly predominant force in historical change.

Technological determinists will be pessimists if they hold that the consequences of technology are on balance socially and environmentally harmful. Moreover, any form of determinism implies a limitation of human

freedom and technological choice. However, some determinists retain great optimism about the consequences of technology. On the other hand, pessimists do not necessarily accept determinism, even in its weaker form. They may acknowledge the presence of technological choices but expect such choices to be misused because they are pessimistic about human nature and institutionalized greed. They may be pessimistic about our ability to respond to a world of global inequities and scarce resources. Nevertheless, determinism and pessimism are often found together among the critics of technology.

3. *Contextual Interaction.* Here there are six arrows instead of two, representing the complex interactions between science, technology, and society. Social and political forces affect the design as well as the uses of particular technologies. Technologies are not neutral because social goals and institutional interests are built into the technical designs that are chosen. Because there are choices, public policy decisions about technology play a larger role here than in either of the other views. Contextualism is most common among our third group, those who see technology as an ambiguous instrument of social power.

Contextualists also point to *the diversity of science-technology interactions*. Sometimes a technology was indeed based on recent scientific discoveries. Biotechnology, for example, depends directly on recent research in molecular biology. In other cases, such as the steam engine or the electric power system, innovations occurred with very little input from new scientific discoveries. A machine or process may have been the result of creative practical innovation or the modification of an existing technology. As Frederick Ferré puts it, science and technology in the modern world are both products of the combination of theoretical and practical intelligence, and “neither gave birth to the other.”⁴⁴ Technology has its own distinctive problems and builds up its own knowledge base and professional community, though it often uses science as a resource to draw on. The reverse contribution of technology to science is also often evident. The work of astronomer’s, for instance, has been dependent on its succession of new technologies, from optical telescopes to microwave antennae and rockets. George Wise writes, “Historical studies have shown that the relations between science and technology need not be those of domination and subordination. Each has maintained its distinctive knowledge base and methods while contributing to the other and to its patrons as well.”⁴⁵

In the previous volume, I discussed the “*social construction of science*” thesis, in which it is argued that not only the direction of scientific development but also the concepts and theories of science are determined by cultural assumptions and interests. I concluded that the “strong program” among sociologists and philosophers of science carries this historical and cultural relativism too far, and I defended a reformulated understanding of objectivity, which gives a major role to empirical data while acknowledging the influence of society on interpretive paradigms.

The case for “*the social construction of technology*” seems to me much stronger. Values are built into particular technological designs. There is no one “best way” to design a technology. Different individuals and groups may define a problem differently and may have diverse criteria of success. Bijker and Pinch show that in the late nineteenth century inventors constructed many different types of bicycles. Controversies developed about the relative size of front and rear wheels, seat location, air tires, brakes, and so forth. Diverse users were envisioned (workers, vacationers, racers, men and women) and diverse criteria (safety, comfort, speed, and so forth). In addition, the bicycle carried cultural meanings, affecting a person’s self-image and social status. There was nothing logically or technically necessary about the model that finally won out and is now found around the world.⁴⁶

The historian John Staudenmaier writes that

contextualism is rooted in the proposition that technical designs cannot be meaningfully interpreted in abstraction from their human context. The human fabric is not an envelope around a culturally neutral artifact. The values and world views, the intelligence and stupidity, the biases and vested interests of those who design, accept and maintain a technology are embedded in the technology itself.⁴⁷

Both the linear and the determinist view imply that technology determines *work organization*. It is said that the technologies of the Industrial Revolution imposed their own requirements and made repetitive tasks inevitable. The contextualists reply that the design of a technology is itself affected by social relations. The replacement of workers by machines was intended not only to reduce labor costs but also to assert greater control by management over labor. For instance, the spinning mule helped to break the power of labor unions among skilled textile workers in nineteenth-century England.

Other contextualists have pointed to the role of technology in *the subordination of women*. Engineering was once considered heavy and dirty work unsuitable for women, but long after it became a clean and intellectual profession, there are still few women in it. Technology has been an almost exclusively male preserve, reflected in toys for boys, the expectations of parents and teachers, and the vocational choices and job opportunities open to men and women. Most technologies are designed by men and add to the power of men.

Strong *gender divisions* are present among employees of technology-related companies. When telephones were introduced, women were the switchboard operators and record keepers, while men designed and repaired the equipment and managed the whole system. Typesetting in large printing frames once required physical strength and mechanical skills and was a male occupation. But men continued to exclude women from compositors' unions when linotype, and more recently computer formatting, required only typing and formatting skills.⁴⁸ Today most computer designers and programmers are men, while in offices most of the data are entered at computer keyboards by women. With many middle-level jobs eliminated, these lower-level jobs often become dead ends for women.⁴⁹ A study of three computerized industries in Britain found that women were the low-paid operators, while only men understood and controlled the equipment, and men almost never worked under the supervision of women.⁵⁰

Note that contextualism allows for a *two-way interaction* between technology and society. When technology is treated as merely one form of cultural expression among others, its distinctive characteristics may be ignored. In some renditions, the way in which technology shapes culture are forgotten while the cultural forces on technology are scrutinized. The impact of technology on society is particularly important in the transfer of a technology to a new cultural setting in a developing country. Some Third World authors have been keenly aware of technology as an instrument of power, and they portray a two-way interaction between technology and society across national boundaries.

IV. CONCLUSIONS

The optimists stress the contribution of technology to *economic development*. They hold that greater productivity improves standards of living and makes *food and health* more widely available. For most of them, the most important form of *participatory freedom* is the economic freedom of the marketplace, though in general they are also committed to political democracy. These authors say that social justice and environmental protection should not be ignored, but they must not be allowed to jeopardize economic goals. The optimists usually evaluate technology in a utilitarian framework, seeking to maximize the balance of costs over benefits.

The pessimists typically make *personal fulfillment* their highest priority, and they interpret fulfillment in terms of human relationships and community life rather than material possessions. They are concerned about individual rights and the dignity of persons. They hold that *meaningful work* is as important as economic productivity in policies for technology. The pessimists are dedicated to resource sustainability and criticize the high levels of consumption in industrial societies today. They often advocate *respect for all creatures* and question the current technological goal of mastery of nature.

The contextualists are more likely to give prominence to *social justice* because they interpret technology as both a product and an instrument of social power. For them the most important forms of *participatory freedom* are opportunities for participation in political processes and in work-related decisions. They are less concerned about economic growth than about how that growth is distributed and who receives the costs and the benefits. Contextualists often seek environmental protection because they are aware of the natural as well as the social contexts in which technologies operate.

I am most sympathetic with the contextualists, though I am indebted to many of the insights of the pessimists. Four issues seem to me particularly important in analyzing the differences among the positions outlined above.

1. *Defense of the Personal.* The pessimists have defended human values in a materialistic and impersonal society. The place to begin, they say, is one's own life. Each of us can adopt individual life-styles more consistent with human and environmental values. Moreover, strong protest and vivid examples are needed to challenge the historical dominance of technological optimism and the disproportionate resource consumption of affluent societies. I admire these critics for defending individuality and choice in the face of standardization and bureaucracy. I join them in upholding the significance of personal relationships and a vision of personal fulfillment that goes beyond material affluence. I affirm the importance of the spiritual life, but I do not believe that it requires a rejection of technology. The answer to the destructive features of technology is not less technology but technology of the right kind.

2. *The Role of Politics.* Differing models of social change are implied in the three positions. The first group usually assumes a free market model. Technology is predominantly beneficial, and the reduction of any undesirable side effects is itself a technical problem for the experts. Government intervention is needed only to regulate the most harmful impacts. Writers mentioned in the second section, by contrast, typically adopt some variant of technological determinism. Technology is dehumanizing and uncontrollable. They see runaway technology as all autonomous and all-embracing system that molds all of life, including the political sphere, to its requirements. The individual is

helpless within the system. The views expressed in the third section presuppose a “social conflict” model. Technology influences human life but is itself part of a cultural system; it is an instrument of social power serving the purposes of those who control it. It does systematically impose distinctive forms on all areas of life, but these can be modified through political processes. Whereas the first two groups give little emphasis to politics, the third, with which I agree, holds that conflicts concerning technology must be resolved primarily in the political arena.

3. *The Redirection of Technology.* I believe that we should neither accept uncritically the past directions of technological development nor reject technology *in toto* but redirect it toward the realization of human and environmental values. In the past, technological decisions have usually been governed by narrowly economic criteria, to the neglect of environmental and human costs. In a later chapter we will look at technology assessment, a procedure designed to use a broad range of criteria to evaluate the diverse consequences of an emerging technology—*before* it has been deployed and has developed the vested interests and institutional momentum that make it seem uncontrollable. I will argue that new policy priorities concerning agriculture, energy, resource allocation, and the redirection of technology toward basic human needs can be achieved within democratic political institutions. The key question will be: What decision-making processes and what technological policies can contribute to human and environmental values?

4. *The Scale of Technology.* Appropriate technology can be thought of as an attempt to achieve some of the material benefits of technology outlined in the first section without the destructive human costs discussed in the second section, most of which result from large-scale centralized technologies. Intermediate-scale technology allows decentralization and greater local participation in decisions. The decentralization of production also allows greater use of local materials and often a reduction of impact on the environment. Appropriate technology does not imply a return to primitive and prescientific methods; rather, it seeks to use the best science available toward goals different from those that have governed industrial production in the past.

Industrial technology was developed when capital and resources were abundant, and we continue to assume these conditions. Automation, for example, is capital-intensive and labor saving. Yet in *developing nations* capital is scarce and labor is abundant. The technologies needed there must be relatively inexpensive and labor-intensive. They must be of intermediate scale so that jobs can be created in rural areas and small towns, to slow down mass migration to the cities. They must fulfill basic human needs, especially for food, housing, and health. Alternative patterns of modernization are less environmentally and socially destructive than the path that we have followed. It is increasingly evident that many of these goals are desirable also in industrial nations. I will suggest that we should develop a mixture of large- and intermediate-scale technologies, which will require deliberate encouragement of the latter.

The redirection of technology will be no easy task. Contemporary technology is so tightly tied to industry, government, and the structures of economic power that changes in direction will be difficult to achieve. As the critics of technology recognize, the person who tries to work for change within the existing order may be absorbed by the establishment. But the welfare of humankind requires a creative technology that is economically productive, ecologically sound, socially just, and personally fulfilling.

References

1. Charles Susskind, *Understanding Technology* (Baltimore: Johns Hopkins University Press, 1973), p. 132.
2. John Zerman and Alice Carnes, eds., *Questioning Technology* (Santa Cruz, CA: New Society Publishers, 1991), p. 217.
3. C. S. Lewis, *The Abolition of Man* (New York: Macmillan, 1965), p. 69.
4. Among the volumes dealing with broad attitudes toward technology are Albert H. Teich, ed., *Technology and the Future*, 5th ed. (New York: St Martin's Press, 1989), and Carl Mitcham and Robert Mackey, eds., *Philosophy and Technology* (New York: Free Press, 1972).
5. This is close to the definition given by Arnold Pacey in *The Culture of Technology* (Cambridge: MIT Press, 1983), p. 6. Pacey adds “living things” among the “ordered systems” (in order to include agriculture, medicine, and biotechnology), but I suggest that these are already included under the rubric of “practical tasks.” Frederick Ferré, *Philosophy of Technology* (Englewood Cliffs, NJ: Prentice-Hall, 1988), defines technology as “the practical implementation of intelligence” and argues that intelligence itself has both practical and theoretical forms.
6. Emanuel Mesthene, *Technological Change: Its Impact on Man and Society* (New York: New American Library, 1970).
7. Melvin Kranzberg, “Technology the Liberator,” in *Technology at the Turning Point*, ed. William Pickett (San Francisco: San Francisco Press, 1977). See also Charles Susskind, *Understanding Technology*.
8. Emanuel Mesthene, “Technology as Evil: Fear or Lamentation?” in *Research in Philosophy and Technology*, vol. 7, ed. Paul Durbin (Greenwich, CT: JAI Press, 1984).
9. Daniel Bell, *The Coming Postindustrial Society* (New York: Basic Books, 1973).
10. Buckminster Fuller, *The Critical Path* (New York: St. Martin's Press, 1981); Herman Kahn et al., *The Next 200 Years*

- (New York: William Morrow, 1976); Alvin Toffler, *Future Shock* (New York: Bantam, 1971) and *The Third Wave* (New York: William Morrow, 1980).
11. Samuel Florman, *The Existential Pleasure of Engineering* (New York: St. Martin's Press, 1977) and *Blaming Technology: The Irrational Search for Scapegoats*, (New York: St. Martin's Press, 1981).
 12. Florman, *Blaming Technology*, p. 188.
 13. Cf. Alvin Weinberg, "Can Technology Replace Social Engineering," in *Technology and the Future*, ed. Tecih.
 14. Samuel Florman, "Science for Public Consumption: More Than We Can Chew?" *Technology Review* 86 (April 1983) 12-13.
 15. Florman, *Blaming Technology*, p. 193.
 16. Harvey Cox, *The Secular City* (New York: Macmillan, 1965), and "The Responsibility of the Christian in a World of Technology," in *Science and Religion*, ed. Ian G. Barbour (New York: Harper & Row, 1968).
 17. W. Norris Clarke, S.J., "Technology and Man: A Christian Vision," in *Science and Religion*, ed. Barbour.
 18. Pierre Teilhard de Chardin, *The Future of Man*, trans. Norman Denny (New York: Harper & Row, 1961), chaps. 8, 9, and 10. See also "The Place of Technology in a General Biology of Mankind," and "On Looking at a Cyclotron," in *The Activation of Energy* (New York: Harcourt Brace Jovanovich, 1971).
 19. George Wise, "Science and Technology," *Osiris*, 2d ser., 1 (1985): 229-46.
 20. See for example Lewis Mumford, *The Myth of the Machine*, vol. 1, *Technics and Human Development*, and vol. 2, *The Pentagon of Power* (New York: Harcourt Brace Jovanovich, 1967 and 1969).
 21. Studs Terkel, *Working* (New York: Pantheon, 1972); Robert Schrag, *Ten Thousand Working Days* (Cambridge: MIT Press, 1978); William A. Faunce, *Problems of an Industrial Society*, 2d ed. (New York: McGraw-Hill, 1981).
 22. Theodore Roszak, *The Making of a Counter Culture* (New York: Doubleday, 1969), and *Where the Wasteland Ends* (New York: Doubleday, 1972); see Ian G. Barbour. "Science, Religion, and the Counterculture," *Zygon* 10 (1975): 380-97.
 23. Jacques Ellul, *The Technological Society*, trans. J. Wilkinson (New York: Knopf, 1964); also *The Technological System*, trans. J. Neugroschel (New York: Continuum, 1980), and *The Technological Bluff*, trans. G. Bromiley (Grand Rapids: Eedmans, 1999).
 24. Darrell Fasching, "The Dialectic of Apocalypse and Utopia in the Theological Ethics of Jacques Ellul," in *Research in Philosophy and Technology*, vol. 10, ed. Frederick Ferré (Greenwich, CT: JAI Press, 1990).
 25. Langdon Winner, *Autonomous Technology* (Cambridge: MIT Press, 1977) and *The Reactor and the Whale* (Chicago: University of Chicago Press, 1986).
 26. Hans Jonas, *The Imperative of Responsibility: In Search of an Ethics for the Technological Age* (Chicago: University of Chicago Press, 1984), p. x.
 27. Albert Borgmann, *Technology and the Character of Contemporary Life* (Chicago: University of Chicago Press, 1984); Martin Heidegger, *The Question Concerning Technology*, trans. William Lovitt (New York: Harper & Row, 1977).
 28. David Kipnis, *Technology and Power* (Berlin: Springer-Verlag, 1990).
 29. Langdon Gilkey, *Religion and the Scientific Future* (New York: Harper & Row, 1970).
 30. Paul Tillich, "The Person in a Technological Society," in *Social Ethics*, ed. Gibson Winter (New York: Harper & Row, 1968).
 31. Gabriel Marcel, "The Sacred in the Technological Age," *Theology Today* 19 (1962): 27-38.
 32. Martin Buber. *I and Thou*, trans. R. G. Smith (New York: Charles Scribner's Sons, 1937).
 33. P. Hans Sun, "Notes on How to Begin to Think about Technology in a Theological Way," in *Theology and Technology*, ed. Carl Mitcham and Jim Grote (New York: University Press of America, 1984).
 34. Thomas Misa, "How Machines Make History, and How Historians (and Others) Help Them Do So," *Science, Technology & Human Values* 13 (1988): 308-31.
 35. Arnold Pacey, *Culture of Technology*.
 36. Victor Ferkiss, *Technological Man and The Future of Technological Civilizations* (New York: George Braziller, 1969 and 1974).
 37. Bernard Gendron, *Technology and the Human Condition* (New York: St. Martin's Press, 1977).
 38. Norman Faramelli, *Technethics* (New York: Friendship Press, 1971).
 39. J. Edward Carothers, Margaret Mead, Daniel McCracken, and Roger Shinn, eds., *To Love or to Perish: The Technological Crisis and the Churches* (New York: Friendship Press, 1972); Paul Abrecht and Roger Shinn, eds., *Faith and Science in an Unjust World* (Geneva: World Council of Churches, 1980).
 40. Thomas Derr, "Conversations about Ultimate Matters: Theological Motifs in WCC Studies on the Technological Future," *International Review of Missions* 66 (1977): 123-34.
 41. Egbert Schuurman, *Technology and the Future* (Toronto: Wedge Publishing, 1980), also "The Modern Babylon Culture," in *Technology and Responsibility*, ed. Paul Durbin (Dordrecht, Holland: D. Reidel,

- 1987), and “A Christian Philosophical Perspective on Technology,” in *Theology and Technology*, ed. Mitcham and Grote. Schuurman was also a contributor to Stephen Monsma, ed. *Responsible Technology: A Christian Perspective* (Grand Rapids: Eerdmans, 1986).
42. Roger Shinn, *Forced Options: Social Decisions for the 21st Century*, 3d ed. (Cleveland: Pilgrim Press, 1991).
 43. Richard Niebuhr, *Christ and Culture* (New York: Harper & Brothers, 1951). See also Carl Mitcham, “Technology as a Theological Problem in the Christian Tradition,” in *Theology Technology*, ed. Mitcham and Grote.
 44. Ferré, *Philosophy of Technology*, p. 44.
 45. Wise, “Science and Technology.”
 46. Trevor Pinch and Wiebe Bijker, “The Social Construction of Facts and Artifacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit from Each Other,” in *The Social Construction of Technological Systems*, ed. Wiebe Bijker, Thomas Hughes, and Trevor Pinch (Cambridge: MIT Press, 1987).
 47. John W. Staudenmaier, *Technology’s Storyteller* (Cambridge: MIT Press, 1985), p. 165.
 48. Cynthia Cockburn, “The Material of Male Power,” in *The Social Shaping of Technology*, ed. Donald McKenzie and Judy Wajcman (Milton Keynes, England: Open University Press, 1985).
 49. Roslyn Feldberg and Evelyn Nakano Glenn, “Technology and Work Degradation: Effects of Office Automation on Women Clerical Workers,” in *Machina Ex Dea: Feminist Perspectives on Technology*, ed. Joan Rothschild (New York: Pergamon Press, 1983); see also articles by Cheris Kramarae, Anne Machung, and others in *Technology and Women’s Voices*, ed. Cheris Kramarae (New York and London: Routledge & Kegan Paul, 1988).
 50. Cynthia Cockburn, *Machinery of Dominance: Women, Men, and Technological Know-How* (London: Pluto Press, 1985).

Do Machines Make History?

ROBERT L. HEILBRONER

The hand-mill gives you society with the feudal lord; the steam-mill, society with the industrial capitalist.

MARX, *The Poverty of Philosophy*

That machines make history in some sense—that the level of technology has a direct bearing on the human drama—is of course obvious. That they do not make all of history, however that word be defined, is equally clear. The challenge, then, is to see if one can say something systematic about the matter, to see whether one can order the problem so that it becomes intellectually manageable.

To do so calls at the very beginning for a careful specification of our task. There are a number of important ways in which machines make history that will not concern us here. For example, one can study the impact of technology on the *political* course of history, evidenced most strikingly by the central role played by the technology of war. Or one can study the effect of machines on the *social* attitudes that underlie historical evolution: one thinks of the effect of radio or television on political behavior. Or one can study technology as one of the factors shaping the changeful content of life from one epoch to another: when we speak of “life” in the Middle Ages or today we define an existence much of whose texture and substance is intimately connected with the prevailing technological order.

None of these problems will form the focus of this essay. Instead, I propose to examine the impact of technology on history in another area—an area defined by the famous quotation from Marx that stands beneath our title. The question we are interested in, then, concerns the effect of technology in determining the nature of the *socioeconomic order*. In its simplest terms the question is: did medieval technology bring about feudalism? Is industrial technology the necessary and sufficient condition for capitalism? Or, by extension, will the technology of

PROF. HEILBRONER, of the New School for Social Research, is the author of *The Worldly Philosophers*, *The Limits of American Capitalism*, and other books dealing with economic theory and development.

336 *Robert L. Heilbroner*

the computer and the atom constitute the ineluctable cause of a new social order?

Even in this restricted sense, our inquiry promises to be broad and sprawling. Hence, I shall not try to attack it head-on, but to examine it in two stages:

1. If we make the assumption that the hand-mill does “give” us feudalism and the steam-mill capitalism, this places technological change in the position of a prime mover of social history. Can we then explain the “laws of motion” of technology itself? Or to put the question less grandly, can we explain why technology evolves in the sequence it does?

2. Again, taking the Marxian paradigm at face value, exactly what do we mean when we assert that the hand-mill “gives us” society with the feudal lord? Precisely how does the mode of production affect the superstructure of social relationships?

These questions will enable us to test the empirical content—or at least to see if there *is* an empirical content—in the idea of technological determinism. I do not think it will come as a surprise if I announce now that we will find *some* content, and a great deal of missing evidence, in our investigation. What will remain then will be to see if we can place the salvageable elements of the theory in historical perspective—to see, in a word, if we can explain technological determinism historically as well as explain history by technological determinism.

I

We begin with a very difficult question hardly rendered easier by the fact that there exist, to the best of my knowledge, no empirical studies on which to base our speculations. It is the question of whether there is a fixed sequence to technological development and therefore a necessary path over which technologically developing societies must travel.

I believe there is such a sequence—that the steam-mill follows the hand-mill not by chance but because it is the next “stage” in a technical conquest of nature that follows one and only one grand avenue of advance. To put it differently, I believe that it is impossible to proceed to the age of the steam-mill until one has passed through the age of the hand-mill, and that in turn one cannot move to the age of the hydroelectric plant before one has mastered the steam-mill, nor to the nuclear power age until one has lived through that of electricity.

Before I attempt to justify so sweeping an assertion, let me make a few reservations. To begin with, I am fully conscious that not all societies are interested in developing a technology of production or in

Do Machines Make History? 337

channeling to it the same quota of social energy. I am very much aware of the different pressures that different societies exert on the direction in which technology unfolds. Lastly, I am not unmindful of the difference between the discovery of a given machine and its application as a technology—for example, the invention of a steam engine (the aeolipile) by Hero of Alexandria long before its incorporation into a steam-mill. All these problems, to which we will return in our last section, refer however to the way in which technology makes its peace with the social, political, and economic institutions of the society in which it appears. They do not directly affect the contention that there exists a determinate sequence of productive technology for those societies that are interested in originating and applying such a technology.

What evidence do we have for such a view? I would put forward three suggestive pieces of evidence:

1. *The Simultaneity of Invention*

The phenomenon of simultaneous discovery is well known.¹ From our view, it argues that the process of discovery takes place along a well-defined frontier of knowledge rather than in grab-bag fashion. Admittedly, the concept of “simultaneity” is impressionistic,² but the related phenomenon of technological “clustering” again suggests that technical evolution follows a sequential and determinate rather than random course.³

2. *The Absence of Technological Leaps*

All inventions and innovations, by definition, represent an advance of the art beyond existing base lines. Yet, most advances, particularly in retrospect, appear essentially incremental, evolutionary. If nature makes no sudden leaps, neither, it would appear, does technology. To make

¹ See Robert K. Merton, “Singletons and Multiples in Scientific Discovery: A Chapter in the Sociology of Science,” *Proceedings of the American Philosophical Society*, CV (October 1961), 470–86.

² See John Jewkes, David Sawers, and Richard Stillerman, *The Sources of Invention* (New York, 1960 [paperback edition]), p. 227, for a skeptical view.

³ “One can count 21 basically different means of flying, at least eight basic methods of geophysical prospecting; four ways to make uranium explosive; . . . 20 or 30 ways to control birth. . . . If each of these separate inventions were autonomous, i.e., without cause, how could one account for their arriving in these functional groups?” S. C. Gilfillan, “Social Implications of Technological Advance,” *Current Sociology*, I (1952), 197. See also Jacob Schmookler, “Economic Sources of Inventive Activity,” *Journal of Economic History* (March 1962), pp. 1–20; and Richard Nelson, “The Economics of Invention: A Survey of the Literature,” *Journal of Business*, XXXII (April 1959), 101–19.

338 *Robert L. Heilbroner*

my point by exaggeration, we do not find experiments in electricity in the year 1500, or attempts to extract power from the atom in the year 1700. On the whole, the development of the technology of production presents a fairly smooth and continuous profile rather than one of jagged peaks and discontinuities.

3. *The Predictability of Technology*

There is a long history of technological prediction, some of it ludicrous and some not.⁴ What is interesting is that the development of technical progress has always seemed *intrinsically* predictable. This does not mean that we can lay down future timetables of technical discovery, nor does it rule out the possibility of surprises. Yet I venture to state that many scientists would be willing to make *general* predictions as to the nature of technological capability twenty-five or even fifty years ahead. This too suggests that technology follows a developmental sequence rather than arriving in a more chancy fashion.

I am aware, needless to say, that these bits of evidence do not constitute anything like a “proof” of my hypothesis. At best they establish the grounds on which a *prima facie* case of plausibility may be rested. But I should like now to strengthen these grounds by suggesting two deeper-seated reasons why technology *should* display a “structured” history.

The first of these is that a major constraint always operates on the technological capacity of an age, the constraint of its accumulated stock of available knowledge. The application of this knowledge may lag behind its reach; the technology of the hand-mill, for example, was by no means at the frontier of medieval technical knowledge, but technical realization can hardly precede what men generally know (although experiment may incrementally advance both technology and knowledge concurrently). Particularly from the mid-nineteenth century to the present do we sense the loosening constraints on technology stemming from successively yielding barriers of scientific knowledge—loosening constraints that result in the successive arrival of the electrical, chemical, aeronautical, electronic, nuclear, and space stages of technology.⁵

⁴ Jewkes *et al.* (see n. 2) present a catalogue of chastening mistakes (p. 230 f.). On the other hand, for a sober predictive effort, see Francis Bello, “The 1960s: A Forecast of Technology,” *Fortune*, LIX (January 1959), 74–78; and Daniel Bell, “The Study of the Future,” *Public Interest*, I (Fall 1965), 119–30. Modern attempts at prediction project likely avenues of scientific advance or technological function rather than the feasibility of specific machines.

⁵ To be sure, the inquiry now regresses one step and forces us to ask whether there are inherent stages for the expansion of knowledge, at least insofar as it ap-

Do Machines Make History? 339

The gradual expansion of knowledge is not, however, the only order-bestowing constraint on the development of technology. A second controlling factor is the material competence of the age, its level of technical expertise. To make a steam engine, for example, requires not only some knowledge of the elastic properties of steam but the ability to cast iron cylinders of considerable dimensions with tolerable accuracy. It is one thing to produce a single steam-machine as an expensive toy, such as the machine depicted by Hero, and another to produce a machine that will produce power economically and effectively. The difficulties experienced by Watt and Boulton in achieving a fit of piston to cylinder illustrate the problems of creating a technology, in contrast with a single machine.

Yet until a metal-working technology was established—indeed, until an embryonic machine-tool industry had taken root—an industrial technology was impossible to create. Furthermore, the competence required to create such a technology does not reside alone in the ability or inability to make a particular machine (one thinks of Babbage's ill-fated calculator as an example of a machine born too soon), but in the ability of many industries to change their products or processes to "fit" a change in one key product or process.

This necessary requirement of technological congruence⁶ gives us an additional cause of sequencing. For the ability of many industries to co-operate in producing the equipment needed for a "higher" stage of technology depends not alone on knowledge or sheer skill but on the division of labor and the specialization of industry. And this in turn hinges to a considerable degree on the sheer size of the stock of capital itself. Thus the slow and painful accumulation of capital, from which springs the gradual diversification of industrial function, becomes an independent regulator of the reach of technical capability.

In making this general case for a determinate pattern of technological evolution—at least insofar as that technology is concerned with production—I do not want to claim too much. I am well aware that reasoning about technical sequences is easily faulted as *post hoc ergo propter hoc*. Hence, let me leave this phase of my inquiry by suggesting no more

plies to nature. This is a very uncertain question. But having already risked so much, I will hazard the suggestion that the roughly parallel sequential development of scientific understanding in those few cultures that have cultivated it (mainly classical Greece, China, the high Arabian culture, and the West since the Renaissance) makes such a hypothesis possible, provided that one looks to broad outlines and not to inner detail.

⁶ The phrase is Richard LaPiere's in *Social Change* (New York, 1965), p. 263 f.

340 *Robert L. Heilbroner*

than that the idea of a roughly ordered progression of productive technology seems logical enough to warrant further empirical investigation. To put it as concretely as possible, I do not think it is just by happenstance that the steam-mill follows, and does not precede, the hand-mill, nor is it mere fantasy in our own day when we speak of the coming of the automatic factory. In the future as in the past, the development of the technology of production seems bounded by the constraints of knowledge and capability and thus, in principle at least, open to prediction as a determinable force of the historic process.

II

The second proposition to be investigated is no less difficult than the first. It relates, we will recall, to the explicit statement that a given technology imposes certain social and political characteristics upon the society in which it is found. Is it true that, as Marx wrote in *The German Ideology*, “A certain mode of production, or industrial stage, is always combined with a certain mode of cooperation, or social stage,”⁷ or as he put it in the sentence immediately preceding our hand-mill, steam-mill paradigm, “In acquiring new productive forces men change their mode of production, and in changing their mode of production they change their way of living—they change all their social relations”?

As before, we must set aside for the moment certain “cultural” aspects of the question. But if we restrict ourselves to the functional relationships directly connected with the process of production itself, I think we can indeed state that the technology of a society imposes a determinate pattern of social relations on that society.

We can, as a matter of fact, distinguish at least two such modes of influence:

1. *The Composition of the Labor Force*

In order to function, a given technology must be attended by a labor force of a particular kind. Thus, the hand-mill (if we may take this as referring to late medieval technology in general) required a work force composed of skilled or semiskilled craftsmen, who were free to practice their occupations at home or in a small atelier, at times and seasons that varied considerably. By way of contrast, the steam-mill—that is, the technology of the nineteenth century—required a work force composed of semiskilled or unskilled operatives who could work only at the factory site and only at the strict time schedule enforced by turning the machinery on or off. Again, the technology of the electronic age has

⁷ Karl Marx and Friedrich Engels, *The German Ideology* (London, 1942), p. 18.

Do Machines Make History? 341

steadily required a higher proportion of skilled attendants; and the coming technology of automation will still further change the needed mix of skills and the locale of work, and may as well drastically lessen the requirements of labor time itself.

2. The Hierarchical Organization of Work

Different technological apparatuses not only require different labor forces but different orders of supervision and co-ordination. The internal organization of the eighteenth-century handicraft unit, with its typical man-master relationship, presents a social configuration of a wholly different kind from that of the nineteenth-century factory with its men-manager confrontation, and this in turn differs from the internal social structure of the continuous-flow, semi-automated plant of the present. As the intricacy of the production process increases, a much more complex system of internal controls is required to maintain the system in working order.

Does this add up to the proposition that the steam-mill gives us society with the industrial capitalist? Certainly the class characteristics of a particular society are strongly implied in its functional organization. Yet it would seem wise to be very cautious before relating political effects exclusively to functional economic causes. The Soviet Union, for example, proclaims itself to be a socialist society although its technical base resembles that of old-fashioned capitalism. Had Marx written that the steam-mill gives you society with the industrial *manager*, he would have been closer to the truth.

What is less easy to decide is the degree to which the technological infrastructure is responsible for some of the sociological features of society. Is anomie, for instance, a disease of capitalism or of all industrial societies? Is the organization man a creature of monopoly capital or of all bureaucratic industry wherever found? These questions tempt us to look into the problem of the impact of technology on the existential quality of life, an area we have ruled out of bounds for this paper. Suffice it to say that superficial evidence seems to imply that the similar technologies of Russia and America are indeed giving rise to similar social phenomena of this sort.

As with the first portion of our inquiry, it seems advisable to end this section on a note of caution. There is a danger, in discussing the structure of the labor force or the nature of intrafirm organization, of assigning the sole causal efficacy to the visible presence of machinery and of overlooking the invisible influence of other factors at work. Gilfillan, for instance, writes, "engineers have committed such blunders as saying

342 *Robert L. Heilbroner*

the typewriter brought women to work in offices, and with the typesetting machine made possible the great modern newspaper, forgetting that in Japan there are women office workers and great modern newspapers getting practically no help from typewriters and typesetting machines.”⁸ In addition, even where technology seems unquestionably to play the critical role, an independent “social” element unavoidably enters the scene in the *design* of technology, which must take into account such facts as the level of education of the work force or its relative price. In this way the machine will reflect, as much as mould, the social relationships of work.

These caveats urge us to practice what William James called a “soft determinism” with regard to the influence of the machine on social relations. Nevertheless, I would say that our cautions qualify rather than invalidate the thesis that the prevailing level of technology imposes itself powerfully on the structural organization of the productive side of society. A foreknowledge of the shape of the technical core of society fifty years hence may not allow us to describe the political attributes of that society, and may perhaps only hint at its sociological character, but assuredly it presents us with a profile of requirements, both in labor skills and in supervisory needs, that differ considerably from those of today. We cannot say whether the society of the computer will give us the latter-day capitalist or the commissar, but it seems beyond question that it will give us the technician and the bureaucrat.

III

Frequently, during our efforts thus far to demonstrate what is valid and useful in the concept of technological determinism, we have been forced to defer certain aspects of the problem until later. It is time now to turn up the rug and to examine what has been swept under it. Let us try to systematize our qualifications and objections to the basic Marxian paradigm:

1. *Technological Progress Is Itself a Social Activity*

A theory of technological determinism must contend with the fact that the very activity of invention and innovation is an attribute of some societies and not of others. The Kalahari bushmen or the tribesmen of New Guinea, for instance, have persisted in a neolithic technology to the present day; the Arabs reached a high degree of technical proficiency in the past and have since suffered a decline; the classical Chinese developed technical expertise in some fields while unaccount-

⁸ Gilfillan (see n. 3), p. 202.

Do Machines Make History? 343

ably neglecting it in the area of production. What factors serve to encourage or discourage this technical thrust is a problem about which we know extremely little at the present moment.⁹

2. *The Course of Technological Advance Is Responsive to Social Direction*

Whether technology advances in the area of war, the arts, agriculture, or industry depends in part on the rewards, inducements, and incentives offered by society. In this way the direction of technological advance is partially the result of social policy. For example, the system of interchangeable parts, first introduced into France and then independently into England failed to take root in either country for lack of government interest or market stimulus. Its success in America is attributable mainly to government support and to its appeal in a society without guild traditions and with high labor costs.¹⁰ The general *level* of technology may follow an independently determined sequential path, but its areas of application certainly reflect social influences.

3. *Technological Change Must Be Compatible with Existing Social Conditions*

An advance in technology not only must be congruent with the surrounding technology but must also be compatible with the existing economic and other institutions of society. For example, labor-saving machinery will not find ready acceptance in a society where labor is abundant and cheap as a factor of production. Nor would a mass production technique recommend itself to a society that did not have a mass market. Indeed, the presence of slave labor seems generally to inhibit the use of machinery and the presence of expensive labor to accelerate it.¹¹

These reflections on the social forces bearing on technical progress tempt us to throw aside the whole notion of technological determinism as false or misleading.¹² Yet, to relegate technology from an undeserved position of *primum mobile* in history to that of a mediating factor, both acted upon by and acting on the body of society, is not to write off

⁹ An interesting attempt to find a line of social causation is found in E. Hagen, *The Theory of Social Change* (Homewood, Ill., 1962).

¹⁰ See K. R. Gilbert, "Machine-Tools," in Charles Singer, E. J. Holmyard, A. R. Hall, and Trevor I. Williams (eds.), *A History of Technology* (Oxford, 1958), IV, chap. xiv.

¹¹ See LaPiere (see n. 6), p. 284; also H. J. Habbakuk, *British and American Technology in the 19th Century* (Cambridge, 1962), *passim*.

¹² As, for example, in A. Hansen, "The Technological Determination of History," *Quarterly Journal of Economics* (1921), pp. 76-83.

344 *Robert L. Heilbroner*

its influence but only to specify its mode of operation with greater precision. Similarly, to admit we understand very little of the cultural factors that give rise to technology does not depreciate its role but focuses our attention on that period of history when technology is clearly a major historic force, namely Western society since 1700.

IV

What is the mediating role played by technology within modern Western society? When we ask this much more modest question, the interaction of society and technology begins to clarify itself for us:

1. *The Rise of Capitalism Provided a Major Stimulus for the Development of a Technology of Production*

Not until the emergence of a market system organized around the principle of private property did there also emerge an institution capable of systematically guiding the inventive and innovative abilities of society to the problem of facilitating production. Hence the environment of the eighteenth and nineteenth centuries provided both a novel and an extremely effective encouragement for the development of an *industrial* technology. In addition, the slowly opening political and social framework of late mercantilist society gave rise to social aspirations for which the new technology offered the best chance of realization. It was not only the steam-mill that gave us the industrial capitalist but the rising inventor-manufacturer who gave us the steam-mill.

2. *The Expansion of Technology within the Market System Took on a New "Automatic" Aspect*

Under the burgeoning market system not alone the initiation of technical improvement but its subsequent adoption and repercussion through the economy was largely governed by market considerations. As a result, both the rise and the proliferation of technology assumed the attributes of an impersonal diffuse "force" bearing on social and economic life. This was all the more pronounced because the political control needed to buffer its disruptive consequences was seriously inhibited by the prevailing *laissez-faire* ideology.

3. *The Rise of Science Gave a New Impetus to Technology*

The period of early capitalism roughly coincided with and provided a congenial setting for the development of an independent source of technological encouragement—the rise of the self-conscious activity of science. The steady expansion of scientific research, dedicated to the exploration of nature's secrets and to their harnessing for social use,

Do Machines Make History? 345

provided an increasingly important stimulus for technological advance from the middle of the nineteenth century. Indeed, as the twentieth century has progressed, science has become a major historical force in its own right and is now the indispensable precondition for an effective technology.

* * *

It is for these reasons that technology takes on a special significance in the context of capitalism—or, for that matter, of a socialism based on maximizing production or minimizing costs. For in these societies, both the continuous appearance of technical advance and its diffusion throughout the society assume the attributes of autonomous process, “mysteriously” generated by society and thrust upon its members in a manner as indifferent as it is imperious. This is why, I think, the problem of technological determinism—of how machines make history—comes to us with such insistence despite the ease with which we can disprove its more extreme contentions.

Technological determinism is thus peculiarly a problem of a certain historic epoch—specifically that of high capitalism and low socialism—in which the forces of technical change have been unleashed, but when the agencies for the control or guidance of technology are still rudimentary.

The point has relevance for the future. The surrender of society to the free play of market forces is now on the wane, but its subservience to the impetus of the scientific ethos is on the rise. The prospect before us is assuredly that of an undiminished and very likely accelerated pace of technical change. From what we can foretell about the direction of this technological advance and the structural alterations it implies, the pressures in the future will be toward a society marked by a much greater degree of organization and deliberate control. What other political, social, and existential changes the age of the computer will also bring we do not know. What seems certain, however, is that the problem of technological determinism—that is, of the impact of machines on history—will remain germane until there is forged a degree of public control over technology far greater than anything that now exists.

DO ARTIFACTS HAVE POLITICS?

[from Winner, L. (1986). *The whale and the reactor: a search for limits in an age of high technology*. Chicago, University of Chicago Press, 19-39.]

No idea is more provocative in controversies about technology and society than the notion that technical things have political qualities. At issue is the claim that the machines, structures, and systems of modern material culture can be accurately judged not only for their contributions to efficiency and productivity and their positive and negative environmental side effects, but also for the ways in which they can embody specific forms of power and authority. Since ideas of this kind are a persistent and troubling presence in discussions about the meaning of technology, they deserve explicit attention.

Writing in the early 1960s, Lewis Mumford gave classic statement to one version of the theme, arguing that “from late neolithic times in the Near East, right down to our own day, two technologies have recurrently existed side by side: one authoritarian, the other democratic, the first system-centered, immensely powerful, but inherently unstable, the other man-centered, relatively weak, but resourceful and durable.”¹ This thesis stands at the heart of Mumford’s studies of the city, architecture, and history of technics, and mirrors concerns voiced earlier in the works of Peter Kropotkin, William Morris, and other nineteenth-century critics of industrialism. During the 1970s, antinuclear and pro-solar energy movements in Europe and the United States adopted a similar notion as the centerpiece of their arguments. According to environmentalist Denis Hayes, “The increased deployment of nuclear power facilities must lead society toward authoritarianism. Indeed, safe reliance upon nuclear power as the principal source of energy may be possible only in a totalitarian state.” Echoing the views of many proponents of appropriate technology and the soft energy path, Hayes contends that “dispersed solar sources are more compatible than centralized technologies with social equity, freedom and cultural pluralism.”²

An eagerness to interpret technical artifacts in political language is by no means the exclusive property of critics of large-scale, high-technology systems. A long lineage of boosters has insisted that the biggest and best that

science and industry made available were the best guarantees of democracy, freedom, and social justice. The factory system, automobile, telephone, radio, television, space program, and of course nuclear power have all at one time or another been described as democratizing, liberating forces. David Lillienthal’s *TVA: Democracy on the March*, for example, found this promise in the phosphate fertilizers and electricity that technical progress was bringing to rural Americans during the 1940s.³ Three decades later Daniel Boorstin’s *The Republic of Technology* extolled television for “its power to disband armies, to cashier presidents, to create a whole new democratic world.”⁴ Scarcely a new invention comes along that someone doesn’t proclaim it as the salvation of a free society.

It is no surprise to learn that technical systems of various kinds are deeply interwoven in the conditions of modern politics. The physical arrangements of industrial production, warfare, communications, and the like have fundamentally changed the exercise of power and the experience of citizenship. But to go beyond this obvious fact and to argue that certain technologies *in themselves* have political properties seems, at first glance, completely mistaken. We all know that people have politics; things do not. To discover either virtues or evils in aggregates of steel, plastic, transistors, integrated circuits, chemicals, and the like seems just plain wrong, a way of mystifying human artifice and of avoiding the true sources, the human sources of freedom and oppression, justice and injustice. Blaming the hardware appears even more foolish than blaming the victims when it comes to judging conditions of public life.

Hence, the stern advice commonly given those who flirt with the notion that technical artifacts have political qualities: What matters is not technology itself, but the social or economic system in which it is embedded. This maxim, which in a number of variations is the central premise of a theory that can be called the social determination of technology, has an obvious wisdom. It serves as a needed corrective to those who focus uncritically upon such things as “the computer and its social impacts” but who fail to look behind technical devices to see the social circumstances of their development, deployment, and use. This view provides an antidote to naive technological determinism—the idea

that technology develops as the sole result of an internal dynamic and then, unmediated by any other influence, molds society to fit its patterns. Those who have not recognized the ways in which technologies are shaped by social and economic forces have not gotten very far.

But the corrective has its own shortcomings; taken literally, it suggests that technical *things* do not matter at all. Once one has done the detective work necessary to reveal the social origins—power holders behind a particular instance of technological change—one will have explained everything of importance. This conclusion offers comfort to social scientists. It validates what they had always suspected, namely, that there is nothing distinctive about the study of technology in the first place. Hence, they can return to their standard models of social power—those of interest-group politics, bureaucratic politics, Marxist models of class struggle, and the like—and have everything they need. The social determination of technology is, in this view, essentially no different from the social determination of, say, welfare policy or taxation.

There are, however, good reasons to believe that technology is politically significant in its own right, good reasons why the standard models of social science only go so far in accounting for what is most interesting and troublesome about the subject. Much of modern social and political thought contains recurring statements of what can be called a theory of technological politics, an odd mongrel of notions often crossbred with orthodox liberal, conservative, and socialist philosophies.⁵ The theory of technological politics draws attention to the momentum of large-scale sociotechnical systems, to the response of modern societies to certain technological imperatives, and to the ways human ends are powerfully transformed as they are adapted to technical means. This perspective offers a novel framework of interpretation and explanation for some of the more puzzling patterns that have taken shape in and around the growth of modern material culture. Its starting point is a decision to take technical artifacts seriously. Rather than insist that we immediately reduce everything to the interplay of social forces, the theory of technological politics suggests that we pay attention to the characteristics of technical objects and the meaning of those characteristics. A necessary complement to, rather than a replacement for, theories of the social determination of technology, this approach identifies certain technologies as political phenomena in their own right. It points us back, to

borrow Edmund Husserl's philosophical injunction, *to the things themselves*.

In what follows I will outline and illustrate two ways in which artifacts can contain political properties. First are instances in which the invention, design, or arrangement of a specific technical device or system becomes a way of settling an issue in the affairs of a particular community. Seen in the proper light, examples of this kind are fairly straightforward and easily understood. Second are cases of what can be called "inherently political technologies," man-made systems that appear to require or to be strongly compatible with particular kinds of political relationships. Arguments about cases of this kind are much more troublesome and closer to the heart of the matter. By the term "politics" I mean arrangements of power and authority in human associations as well as the activities that take place within those arrangements. For my purposes here, the term "technology" is understood to mean all of modern practical artifice, but to avoid confusion I prefer to speak of "technologies" plural, smaller or larger pieces or systems of hardware of a specific kind.⁶ My intention is not to settle any of the issues here once and for all, but to indicate their general dimensions and significance.

Technical Arrangements and Social Order

ANYONE WHO has traveled the highways of America and has gotten used to the normal height of overpasses may well find something a little odd about some of the bridges over the park ways on Long Island, New York. Many of the overpasses are extraordinarily low, having as little as nine feet of clearance at the curb. Even those who happened to notice this structural peculiarity would not be inclined to attach any special meaning to it. In our accustomed way of looking at things such as roads and bridges, we see the details of form as innocuous and seldom give them a second thought.

It turns out, however, that some two hundred or so low-hanging overpasses on Long Island are there for a reason. They were deliberately designed and built that way by someone who wanted to achieve a particular social effect. Robert Moses, the master builder of roads, parks, bridges, and other public works of the 1920s to the 1970s in New York, built his overpasses according to specifications that would discourage the presence of buses on his parkways. According to evidence provided by Moses' biographer, Robert A. Caro, the reasons reflect Moses' social class bias

and racial prejudice. Automobile-owning whites of “upper” and “comfortable middle” classes, as he called them, would be free to use the parkways for recreation and commuting. Poor people and blacks, who normally used public transit, were kept off the roads because the twelve-foot tall buses could not handle the overpasses. One consequence was to limit access of racial minorities and low-income groups to Jones Beach, Moses’ widely acclaimed public park. Moses made doubly sure of this result by vetoing a proposed extension of the Long Island Railroad to Jones Beach.

Robert Moses’ life is a fascinating story in recent U. S. political history. His dealings with mayors, governors, and presidents; his careful manipulation of legislatures, banks, labor unions, the press, and public opinion could be studied by political scientists for years. But the most important and enduring results of his work are his technologies, the vast engineering projects that give New York much of its present form. For generations after Moses’ death and the alliances he forged have fallen apart, his public works, especially the highways and bridges he built to favor the use of the automobile over the development of mass transit, will continue to shape that city. Many of his monumental structures of concrete and steel embody a systematic social inequality, a way of engineering relationships among people that, after a time, became just another part of the landscape. As New York planner Lee Koppleman told Caro about the low bridges on Wantagh Parkway, “The old son of a gun had made sure that buses would *never* be able to use his goddamned parkways.”⁷

Histories of architecture, city planning, and public works contain many examples of physical arrangements with explicit or implicit political purposes. One can point to Baron Haussmann’s broad Parisian thoroughfares, engineered at Louis Napoleon’s direction to prevent any recurrence of street fighting of the kind that took place during the revolution of 1848. Or one can visit any number of grotesque concrete buildings and huge plazas constructed on university campuses in the United States during the late 1960s and early 1970s to defuse student demonstrations. Studies of industrial machines and instruments also turn up interesting political stories, including some that violate our normal expectations about why technological innovations are made in the first place. If we suppose that new technologies are introduced to achieve increased efficiency, the history of technology shows that we will sometimes be

disappointed. Technological change expresses a panoply of human motives, not the least of which is the desire of some to have dominion over others even though it may require an occasional sacrifice of cost savings and some violation of the normal standard of trying to get more from less.

One poignant illustration can be found in the history of nineteenth-century industrial mechanization. At Cyrus McCormick’s reaper manufacturing plant in Chicago in the middle 1880s, pneumatic molding machines, a new and largely untested innovation, were added to the foundry at an estimated cost of \$500,000. The standard economic interpretation would lead us to expect that this step was taken to modernize the plant and achieve the kind of efficiencies that mechanization brings. But historian Robert Ozanne has put the development in a broader context. At the time, Cyrus McCormick II was engaged in a battle with the National Union of Iron Molders. He saw the addition of the new machines as a way to “weed out the bad element among the men,” namely, the skilled workers who had organized the union local in Chicago.⁸ The new machines, manned by unskilled laborers, actually produced inferior castings at a higher cost than the earlier process. After three years of use the machines were, in fact, abandoned, but by that time they had served their purpose—the destruction of the union. Thus, the story of these technical developments at the McCormick factory cannot be adequately understood outside the record of workers’ attempts to organize, police repression of the labor movement in Chicago during that period, and the events surrounding the bombing at Haymarket Square. Technological history and U.S. political history were at that moment deeply intertwined.

In the examples of Moses’ low bridges and McCormick’s molding machines, one sees the importance of technical arrangements that precede the *use* of the things in question. It is obvious that technologies can be used in ways that enhance the power, authority, and privilege of some over others, for example, the use of television to sell a candidate. In our accustomed way of thinking technologies are seen as neutral tools that can be used well or poorly, for good, evil, or something in between. But we usually do not stop to inquire whether a given device might have been designed and built in such a way that it produces a set of consequences logically and temporally *prior to any of its professed uses*. Robert Moses’ bridges, after all, were used to carry automobiles

from one point to another; McCormick's machines were used to make metal castings; both technologies, however, encompassed purposes far beyond their immediate use. If our moral and political language for evaluating technology includes only categories having to do with tools and uses, if it does not include attention to the meaning of the designs and arrangements of our artifacts, then we will be blinded to much that is intellectually and practically crucial.

Because the point is most easily understood in the light of particular intentions embodied in physical form, I have so far offered illustrations that seem almost conspiratorial. But to recognize the political dimensions in the shapes of technology does not require that we look for conscious conspiracies or malicious intentions. The organized movement of handicapped people in the United States during the 1970s pointed out the countless ways in which machines, instruments, and structures of common use—buses, buildings, sidewalks, plumbing fixtures, and so forth—made it impossible for many handicapped persons to move freely about, a condition that systematically excluded them from public life. It is safe to say that designs unsuited for the handicapped arose more from long-standing neglect than from anyone's active intention. But once the issue was brought to public attention, it became evident that justice required a remedy. A whole range of artifacts have been redesigned and rebuilt to accommodate this minority.

Indeed, many of the most important examples of technologies that have political consequences are those that transcend the simple categories "intended" and "unintended" altogether. These are instances in which the very process of technical development is so thoroughly biased in a particular direction that it regularly produces results heralded as wonderful breakthroughs by some social interests and crushing setbacks by others. In such cases it is neither correct nor insightful to say, "Someone intended to do somebody else harm." Rather one must say that the technological deck has been stacked in advance to favor certain social interests and that some people were bound to receive a better hand than others.

The mechanical tomato harvester, a remarkable device perfected by researchers at the University of California from the late 1940s to the present offers an illustrative tale. The machine is able to harvest tomatoes in a single pass through a row, cutting the plants from the ground,

shaking the fruit loose, and (in the newest models) sorting the tomatoes electronically into large plastic gondolas that hold up to twenty-five tons of produce headed for canning factories. To accommodate the rough motion of these harvesters in the field, agricultural researchers have bred new varieties of tomatoes that are hardier, sturdier, and less tasty than those previously grown. The harvesters replace the system of handpicking in which crews of farm workers would pass through the fields three or four times, putting ripe tomatoes in lug boxes and saving immature fruit for later harvest.⁹ Studies in California indicate that the use of the machine reduces costs by approximately five to seven dollars per ton as compared to hand harvesting.¹⁰ But the benefits are by no means equally divided in the agricultural economy. In fact, the machine in the garden has in this instance been the occasion for a thorough re-shaping of social relationships involved in tomato production in rural California.

By virtue of their very size and cost of more than \$50,000 each, the machines are compatible only with a highly concentrated form of tomato growing. With the introduction of this new method of harvesting, the number of tomato growers declined from approximately 4,000 in the early 1960s to about 600 in 1973, and yet there was a substantial increase in tons of tomatoes produced. By the late 1970s an estimated 32,000 jobs in the tomato industry had been eliminated as a direct consequence of mechanization.¹¹ Thus, a jump in productivity to the benefit of very large growers has occurred at the sacrifice of other rural agricultural communities.

The University of California's research on and development of agricultural machines such as the tomato harvester eventually became the subject of a lawsuit filed by attorneys for California Rural Legal Assistance, an organization representing a group of farm workers and other interested parties. The suit charged that university officials are spending tax monies on projects that benefit a handful of private interests to the detriment of farm workers, small farmers, consumers, and rural California generally and asks for a court injunction to stop the practice. The university denied these charges, arguing that to accept them "would require elimination of all research with any potential practical application."¹²

As far as I know, no one argued that the development of the tomato harvester was the result of a plot. Two

students of the controversy, William Friedland and Amy Barton, specifically exonerate the original developers of the machine and the hard tomato from any desire to facilitate economic concentration in that industry.¹³ What we see here instead is an ongoing social process in which scientific knowledge, technological invention, and corporate profit reinforce each other in deeply entrenched patterns, patterns that bear the unmistakable stamp of political and economic power. Over many decades agricultural research and development in U.S. land-grant colleges and universities has tended to favor the interests of large agribusiness concerns.¹⁴ It is in the face of such subtly ingrained patterns that opponents of innovations such as the tomato harvester are made to seem “antitechnology” or “antiprogess.” For the harvester is not merely the symbol of a social order that rewards some while punishing others; it is in a true sense an embodiment of that order.

Within a given category of technological change there are, roughly speaking, two kinds of choices that can affect the relative distribution of power, authority, and privilege in a community. Often the crucial decision is a simple “yes or no” choice—are we going to develop and adopt the thing or not? In recent years many local, national, and international disputes about technology have centered on “yes or no” judgments about such things as food additives, pesticides, the building of highways, nuclear reactors, dam projects, and proposed high-tech weapons. The fundamental choice about an antiballistic missile or supersonic transport is whether or not the thing is going to join society as a piece of its operating equipment. Reasons given for and against are frequently as important as those concerning the adoption of an important new law.

A second range of choices, equally critical in many instances, has to do with specific features in the design or arrangement of a technical system after the decision to go ahead with it has already been made. Even after a utility company wins permission to build a large electric power line, important controversies can remain with respect to the placement of its route and the design of its towers; even after an organization has decided to institute a system of computers, controversies can still arise with regard to the kinds of components, programs, modes of access, and other specific features the system will include. Once the mechanical tomato harvester had been developed in its basic form, a design alteration of critical social significance—the addition of electronic sorters, for

example—changed the character of the machine’s effects upon the balance of wealth and power in California agriculture. Some of the most interesting research on technology and politics at present focuses upon the attempt to demonstrate in a detailed, concrete fashion how seemingly innocuous design features in mass transit systems, water projects, industrial machinery, and other technologies actually mask social choices of profound significance. Historian David Noble has studied two kinds of automated machine tool systems that have different implications for the relative power of management and labor in the industries that might employ them. He has shown that although the basic electronic and mechanical components of the record/playback and numerical control systems are similar, the choice of one design over another has crucial consequences for social struggles on the shop floor. To see the matter solely in terms of cost cutting, efficiency, or the modernization of equipment is to miss a decisive element in the story.¹⁵

From such examples I would offer some general conclusions. These correspond to the interpretation of technologies as “forms of life” presented in the previous chapter, filling in the explicitly political dimensions of that point of view.

The things we call “technologies” are ways of building order in our world. Many technical devices and systems important in everyday life contain possibilities for many different ways of ordering human activity. Consciously or unconsciously, deliberately or inadvertently, societies choose structures for technologies that influence how people are going to work, communicate, travel, consume, and so forth over a very long time. In the processes by which structuring decisions are made, different people are situated differently and possess unequal degrees of power as well as unequal levels of awareness. By far the greatest latitude of choice exists the very first time a particular instrument, system, or technique is introduced. Because choices tend to become strongly fixed in material equipment, economic investment, and social habit, the original flexibility vanishes for all practical purposes once the initial commitments are made. In that sense technological innovations are similar to legislative acts or political foundings that establish a framework for public order that will endure over many generations. For that reason the same careful attention one would give to the rules, roles, and relationships of politics must also be given to such things as the building of highways, the creation of television networks, and the tailoring of

seemingly insignificant features on new machines. The issues that divide or unite people in society are settled not only in the institutions and practices of politics proper, but also, and less obviously, in tangible arrangements of steel and concrete, wires and semiconductors, nuts and bolts.

Inherently Political Technologies

NONE OF the arguments and examples considered thus far addresses a stronger, more troubling claim often made in writings about technology and society—the belief that some technologies are by their very nature political in a specific way. According to this view, the adoption of a given technical system unavoidably brings with it conditions for human relationships that have a distinctive political cast—for example, centralized or de-centralized, egalitarian or inegalitarian, repressive or liberating. This is ultimately what is at stake in assertions such as those of Lewis Mumford that two traditions of technology, one authoritarian, the other democratic, exist side-by-side in Western history. In all the cases cited above the technologies are relatively flexible in design and arrangement and variable in their effects. Although one can recognize a particular result produced in a particular setting, one can also easily imagine how a roughly similar device or system might have been built or situated with very much different political consequences. The idea we must now examine and evaluate is that certain kinds of technology do not allow such flexibility, and that to choose them is to choose unalterably a particular form of political life.

A remarkably forceful statement of one version of this argument appears in Friedrich Engels' little essay "On Authority" written in 1872. Answering anarchists who believed that authority is an evil that ought to be abolished altogether, Engels launches into a panegyric for authoritarianism, maintaining, among other things, that strong authority is a necessary condition in modern industry. To advance his case in the strongest possible way, he asks his readers to imagine that the revolution has already occurred. "Supposing a social revolution dethroned the capitalists, who now exercise their authority over the production and circulation of wealth. Supposing, to adopt entirely the point of view of the anti-authoritarians, that the land and the instruments of labour had become the collective property of the workers who use them. Will authority have disappeared or will it have only changed its form?"¹⁶

His answer draws upon lessons from three sociotechnical systems of his day, cotton-spinning mills, railways, and ships at sea. He observes that on its way to becoming finished thread, cotton moves through a number of different operations at different locations in the factory. The workers perform a wide variety of tasks, from running the steam engine to carrying the products from one room to another. Because these tasks must be coordinated and because the timing of the work is "fixed by the authority of the steam," laborers must learn to accept a rigid discipline. They must, according to Engels, work at regular hours and agree to subordinate their individual wills to the persons in charge of factory operations. If they fail to do so, they risk the horrifying possibility that production will come to a grinding halt. Engels pulls no punches. "The automatic machinery of a big factory," he writes, "is much more despotic than the small capitalists who employ workers ever have been."¹⁷

Similar lessons are adduced in Engels's analysis of the necessary operating conditions for railways and ships at sea. Both require the subordination of workers to an "imperious authority" that sees to it that things run according to plan. Engels finds that far from being an idiosyncrasy of capitalist social organization, relationships of authority and subordination arise "independently of all social organization, and are imposed upon us together with the material conditions under which we produce and make products circulate." Again, he intends this to be stern advice to the anarchists who, according to Engels, thought it possible simply to eradicate subordination and superordination at a single stroke. All such schemes are nonsense. The roots of unavoidable authoritarianism are, he argues, deeply implanted in the human involvement with science and technology. "If man, by dint of his knowledge and inventive genius, has subdued the forces of nature, the latter avenge themselves upon him by subjecting him, insofar as he employs them, to a veritable despotism independent of all social organization."¹⁸

Attempts to justify strong authority on the basis of supposedly necessary conditions of technical practice have an ancient history. A pivotal theme in the *Republic* is Plato's quest to borrow the authority of *technology* and employ it by analogy to buttress his argument in favor of authority in the state. Among the illustrations he chooses, like Engels, is that of a ship on the high seas. Because large sailing vessels by their very nature need to be steered with a firm hand, sailors must yield to their captain's commands; no reasonable person believes that

ships can be run democratically. Plato goes on to suggest that governing a state is rather like being captain of a ship or like practicing medicine as a physician. Much the same conditions that require central rule and decisive action in organized technical activity also create this need in government.

In Engels's argument, and arguments like it, the justification for authority is no longer made by Plato's classic analogy, but rather directly with reference to technology itself. If the basic case is as compelling as Engels believed it to be, one would expect that as a society adopted increasingly complicated technical systems as its material basis, the prospects for authoritarian ways of life would be greatly enhanced. Central control by knowledgeable people acting at the top of a rigid social hierarchy would seem increasingly prudent. In this respect his stand in "On Authority" appears to be at variance with Karl Marx's position in Volume I of *Capital*. Marx tries to show that increasing mechanization will render obsolete the hierarchical division of labor and the relationships of subordination that, in his view, were necessary during the early stages of modern manufacturing. "Modern Industry," he writes, "sweeps away by technical means the manufacturing division of labor, under which each man is bound hand and foot for life to a single detail operation. At the same time, the capitalistic form of that industry reproduces this same division of labour in a still more monstrous shape; in the factory proper, by converting the workman into a living appendage of the machine."¹⁹ In Marx's view the conditions that will eventually dissolve the capitalist division of labor and facilitate proletarian revolution are conditions latent in industrial technology itself. The differences between Marx's position in *Capital* and Engels's in his essay raise an important question for socialism: What, after all, does modern technology make possible or necessary in political life? The theoretical tension we see here mirrors many troubles in the practice of freedom and authority that had muddled the tracks of socialist revolution.

Arguments to the effect that technologies are in some sense inherently political have been advanced in a wide variety of contexts, far too many to summarize here. My reading of such notions, however, reveals there are two basic ways of stating the case. One version claims that the adoption of a given technical system actually requires the creation and maintenance of a particular set of social conditions as the operating environment of that system.

Engels's position is of this kind. A similar view is offered by a contemporary writer who holds that "if you accept nuclear power plants, you also accept a techno-scientific industrial-military elite. Without these people in charge, you could not have nuclear power."²⁰ In this conception some kinds of technology require their social environments to be structured in a particular way in much the same sense that an automobile requires wheels in order to move. The thing could not exist as an effective operating entity unless certain social as well as material conditions were met. The meaning of "required" here is that of practical (rather than logical) necessity. Thus, Plato thought it a practical necessity that a ship at sea have one captain and an unquestionably obedient crew.

A second, somewhat weaker, version of the argument holds that a given kind of technology is strongly compatible with, but does not strictly require, social and political relationships of a particular stripe. Many advocates of solar energy have argued that technologies of that variety are more compatible with a democratic, egalitarian society than energy systems based on coal, oil, and nuclear power; at the same time they do not maintain that anything about solar energy requires democracy. Their case is, briefly, that solar energy is decentralizing in both a technical and political sense: technically speaking, it is vastly more reasonable to build solar systems in a disaggregated, widely distributed manner than in large-scale centralized plants; politically speaking, solar energy accommodates the attempts of individuals and local communities to manage their affairs effectively because they are dealing with systems that are more accessible, comprehensible, and controllable than huge centralized sources. In this view solar energy is desirable not only for its economic and environmental benefits, but also for the salutary institutions it is likely to permit in other areas of public life.²¹

Within both versions of the argument there is a further distinction to be made between conditions that are internal to the workings of a given technical system and those that are external to it. Engels's thesis concerns internal social relations said to be required within cotton factories and railways, for example; what such relationships mean for the condition of society at large is, for him, a separate question. In contrast, the solar advocate's belief that solar technologies are compatible with democracy pertains to the way they complement aspects of society removed from the organization of those technologies as such.

There are, then, several different directions that arguments of this kind can follow. Are the social conditions predicated said to be required by, or strongly compatible with, the workings of a given technical system? Are those conditions internal to that system or external to it (or both)? Although writings that address such questions are often unclear about what is being asserted, arguments in this general category are an important part of modern political discourse. They enter into many attempts to explain how changes in social life take place in the wake of technological innovation. More important, they are often used to buttress attempts to justify or criticize proposed courses of action involving new technology. By offering distinctly political reasons for or against the adoption of a particular technology, arguments of this kind stand apart from more commonly employed, more easily quantifiable claims about economic costs and benefits, environmental impacts, and possible risks to public health and safety that technical systems may involve. The issue here does not concern how many jobs will be created, how much income generated, how many pollutants added, or how many cancers produced. Rather, the issue has to do with ways in which choices about technology have important consequences for the form and quality of human associations.

If we examine social patterns that characterize the environments of technical systems, we find certain devices and systems almost invariably linked to specific ways of organizing power and authority. The important question is: Does this state of affairs derive from an unavoidable social response to intractable properties in the things themselves, or is it instead a pattern imposed independently by a governing body, ruling class, or some other social or cultural institution to further its own purposes?

Taking the most obvious example, the atom bomb is an inherently political artifact. As long as it exists at all, its lethal properties demand that it be controlled by a centralized, rigidly hierarchical chain of command closed to all influences that might make its workings unpredictable. The internal social system of the bomb must be authoritarian; there is no other way. The state of affairs stands as a practical necessity independent of any larger political system in which the bomb is embedded, independent of the type of regime or character of its rulers. Indeed, democratic states must try to find ways to ensure that the social structures and mentality that

characterize the management of nuclear weapons do not “spin off” or “spill over” into the polity as a whole.

The bomb is, of course, a special case. The reasons very rigid relationships of authority are necessary in its immediate presence should be clear to anyone. If, however, we look for other instances in which particular varieties of technology are widely perceived to need the maintenance of a special pattern of power and authority, modern technical history contains a wealth of examples. Alfred D. Chandler in *The Visible Hand*, a monumental study of modern business enterprise, presents impressive documentation to defend the hypothesis that the construction and day-to-day operation of many systems of production, transportation, and communication in the nineteenth and twentieth centuries require the development of particular social form—a large-scale centralized, hierarchical organization administered by highly skilled managers. Typical of Chandler’s reasoning is his analysis of the growth of the railroads.²² Technology made possible fast, all-weather transportation; but safe, regular, reliable movement of goods and passengers, as well as the continuing maintenance and repair of locomotives, rolling stock, and track, roadbed, stations, roundhouses, and other equipment, required the creation of a sizable administrative organization. It meant the employment of a set of managers to supervise these functional activities over an extensive geographical area; and the appointment of an administrative command of middle and top executives to monitor, evaluate, and coordinate the work of managers responsible for the day-to-day operations. Throughout his book Chandler points to ways in which technologies used in the production and distribution of electricity, chemicals, and a wide range of industrial goods “demanded” or “required” this form of human association. “Hence, the operational requirements of railroads demanded the creation of the first administrative hierarchies in American business.”²³

Were there other conceivable ways of organizing these aggregates of people and apparatus? Chandler shows that a previously dominant social form, the small traditional family firm, simply could not handle the task in most cases. Although he does not speculate further, it is clear that he believes there is, to be realistic, very little latitude in the forms of power and authority appropriate within modern sociotechnical systems. The properties of many modern technologies.²⁴ But the weight of argument and

empirical evidence in *The Visible Hand* suggests that any significant departure from the basic pattern would be, at best, highly unlikely.

It may be that other conceivable arrangements of power and authority, for example, those of decentralized, democratic worker self-management, could prove capable of administering factories, refineries, communications systems, and railroads as well as or better than the organizations Chandler describes. Evidence from automobile assembly teams in Sweden and worker-managed plants in Yugoslavia and other countries is often presented to salvage these possibilities. Unable to settle controversies over this matter here, I merely point to what I consider to be their bone of contention. The available evidence tends to show that many large, sophisticated technological systems are in fact highly compatible with centralized, hierarchical managerial control. The interesting question, however, has to do with whether or not this pattern is in any sense a requirement of such systems, a question that is not solely empirical. The matter ultimately rests on our judgments about what steps, if any, are practically necessary in the workings of particular kinds of technology and what, if anything, such measures require of the structure of human associations. Was Plato right in saying that a ship at sea needs steering by a decisive hand and that this could only be accomplished by a single captain and an obedient crew? Is Chandler correct in saying that the properties of large-scale systems require centralized, hierarchical managerial control?

To answer such questions, we would have to examine in some detail the moral claims of practical necessity (including those advocated in the doctrines of economics) and weigh them against moral claims of other sorts, for example, the notion that it is good for sailors to participate in the command of a ship or that workers have a right to be involved in making and administering decisions in a factory. It is characteristic of societies based on large, complex technological systems, however, that moral reasons other than those of practical necessity appear increasingly obsolete, "idealistic," and irrelevant. Whatever claims one may wish to make on behalf of liberty, justice, or equality can be immediately neutralized when confronted with arguments to the effect, "Fine, but that's no way to run a railroad" (or steel mill, or airline, or communication system, and so on). Here we encounter an important quality in modern political discourse and in the way people commonly think about what measures are

justified in response to the possibilities technologies make available. In many instances, to say that some technologies are inherently political is to say that certain widely accepted reasons of practical necessity—especially the need to maintain crucial technological systems as smoothly working entities—have tended to eclipse other sorts of moral and political reasoning.

One attempt to salvage the autonomy of politics from the bind of practical necessity involves the notion that conditions of human association found in the internal workings of technological systems can easily be kept separate from the polity as a whole. Americans have long rested content in the belief that arrangements of power and authority inside industrial corporations, public utilities, and the like have little bearing on public institutions, practices, and ideas at large. That "democracy stops at the factory gates" was taken as a fact of life that had nothing to do with the practice of political freedom. But can the internal politics of technology and the politics of the whole community be so easily separated? A recent study of business leaders in the United States, contemporary exemplars of Chandler's "visible hand of management," found them remarkably impatient with such democratic scruples as "one man one vote. If democracy doesn't work for the firm, the most critical institution in all of society, American executives ask, how well can it be expected to work for the government of a nation—particularly when that government attempts to interfere with the achievements of the firm? The authors of the report observe that patterns of authority that work effectively in the corporation become for businessmen "the desirable model against which to compare political and economic relationships in the rest of society."²⁵ While such findings are far from conclusive, they do reflect a sentiment increasingly common in the land: what dilemmas such as the energy crisis require is not a redistribution of wealth or broader public participation but, rather, stronger, centralized public and private management.

An especially vivid case in which the operational requirements of a technical system might influence the quality of public life is the debate about the risks of nuclear power. As the supply of uranium for nuclear reactors runs out, a proposed alternative fuel is the plutonium generated as a byproduct in reactor cores. Well-known objections to plutonium recycling focus on its unacceptable economic costs, its risks of environmental contamination, and its dangers in regard

to the international proliferation of nuclear weapons. Beyond these concerns, however stands another less widely appreciated set of hazards—those that involve the sacrifice of civil liberties. The widespread use of plutonium as a fuel increases the chance that this toxic substance might be stolen by terrorists, organized crime, or other persons. This raises the prospect, and not a trivial one, that extraordinary measures would have to be taken to safeguard plutonium from theft and to recover it should the substance be stolen. Workers in the nuclear industry as well as ordinary citizens outside could well become subject to background security checks, covert surveillance, wiretapping, informers, and even emergency measures under martial law—all justified by the need to safeguard plutonium.

Russell W. Ayres's study of the legal ramifications of plutonium recycling concludes: "With the passage of time and the increase in the quantity of plutonium in existence will come pressure to eliminate the traditional checks the courts and legislatures place on the activities of the executive and to develop a powerful central authority better able to enforce strict safeguards." He avers that "once a quantity of plutonium had been stolen, the case for literally turning the country upside down to get it back would be overwhelming." Ayres anticipates and worries about the kinds of thinking that, I have argued, characterize inherently political technologies. It is still true that in a world in which human beings make and maintain artificial systems nothing is "required" in an absolute sense. Nevertheless, once a course of action is under way, once artifacts such as nuclear power plants have been built and put in operation, the kinds of reasoning that justify the adaptation of social life to technical requirements pop up as spontaneously as flowers in the spring. In Ayres's words, "Once recycling begins and the risks of plutonium theft become real rather than hypothetical, the case for governmental infringement of protected rights will seem compelling."²⁶ After a certain point, those who cannot accept the hard requirements and imperatives will be dismissed as dreamers and fools.

* * *

The two varieties of interpretation I have outlined indicate how artifacts can have political qualities. In the first instance we noticed ways in which specific features in the design or arrangement of a device or system could provide a convenient means of establishing patterns of

power and authority in a given setting. Technologies of this kind have a range of flexibility in the dimensions of their material form. It is precisely because they are flexible that their consequences for society must be understood with reference to the social actors able to influence which designs and arrangements are chosen. In the second instance we examined ways in which the intractable properties of certain kinds of technology are strongly, perhaps unavoidably, linked to particular institutionalized patterns of power and authority. Here the initial choice about whether or not to adopt something is decisive in regard to its consequences. There are no alternative physical designs or arrangements that would make a significant difference; there are, furthermore, no genuine possibilities for creative intervention by different social systems—capitalist or socialist—that could change the intractability of the entity or significantly alter the quality of its political effects.

To know which variety of interpretation is applicable in a given case is often what is at stake in disputes, some of them passionate ones, about the meaning of technology for how we live. I have argued a "both/and" position here, for it seems to me that both kinds of understanding are applicable in different circumstances. Indeed, it can happen that within a particular complex of technology—a system of communication or transportation, for example—some aspects may be flexible in their possibilities for society, while other aspects may be (for better or worse) completely intractable. The two varieties of interpretation I have examined here can overlap and intersect at many points.

These are, of course, issues on which people can disagree. Thus, some proponents of energy from renewable resources now believe they have at last discovered a set of intrinsically democratic, egalitarian, communitarian technologies. In my best estimation, however, the social consequences of building renewable energy systems will surely depend on the specific configurations of both hardware and the social institutions created to bring that energy to us. It may be that we will find ways to turn this silk purse into a sow's ear. By comparison, advocates of the further development of nuclear power seem to believe that they are working on a rather flexible technology whose adverse social effects can be fixed by changing the design parameters of reactors and nuclear waste disposal systems. For reasons indicated above, I believe them to be dead wrong in that faith. Yes, we may be able to

manage some of the “risks” to public health and safety that nuclear power brings. But as society adapts to the more dangerous and apparently indelible features of nuclear power, what will be the long-range toll in human freedom?

My belief that we ought to attend more closely to technical objects themselves is not to say that we can ignore the contexts in which those objects are situated. A ship at sea may well require, as Plato and Engels insisted, a single captain and obedient crew. But a ship out of

service, parked at the dock, needs only a caretaker. To understand which technologies and which contexts are important to us, and why, is an enterprise that must involve both the study of specific technical systems and their history as well as a thorough grasp of the concepts and controversies of political theory. In our times people are often willing to make drastic changes in the way they live to accommodate technological innovation while at the same time resisting similar kinds of changes justified on political grounds. If for no other reason than that, it is important for us to achieve a clearer view of these matters than has been our habit so far.

Notes.

1. Lewis Mumford, “Authoritarian and Democratic Technics,” *Technology and Culture* 5:1-8, 1964.
2. Denis Hayes, *Rays of Hope: The Transition to a Post-Petroleum World* (New York: W. W. Norton, 1977), 71, 159.
3. David Lillienthal, *T.V.A.: Democracy on the March* (New York: Harper and Brothers, 1944), 72-83.
4. Daniel J. Boorstin, *The Republic of Technology* (New York: Harper and Row, 1978), 7.
5. Langdon Winner, *Autonomous Technology: Technics-Out-of-Control as a Theme in Political Thought* (Cambridge: MIT Press, 1977).
6. The meaning of “technology” I employ in this essay does not encompass some of the broader definitions of that concept found in contemporary literature, for example, the notion of “technique” in the writings of Jacques Ellul. My purposes here are more limited. For a discussion of the difficulties that arise in attempts to define “technology,” see *Autonomous Technology*, 8-12.
7. Robert A. Caro, *The Power Broker: Robert Moses and the Fall of New York* (New York: Random House, 1974), 318, 481, 514, 546, 951-958, 952.
8. Robert Ozanne, *A Century of Labor-Management Relations at McCormick and International Harvester* (Madison: University of Wisconsin Press, 1967), 20.
9. The early history of the tomato harvester is told in Wayne D. Rasmussen, “Advances in American Agriculture: The Mechanical Tomato Harvester as a Case Study,” *Technology and Culture* 9:531-543, 1968.
10. Andrew Schmitz and David Seckler, “Mechanized Agriculture and Social Welfare: The Case of the Tomato Harvester,” *American Journal of Agricultural Economics* 52:569-577, 1970.
11. William H. Friedland and Amy Barton, “Tomato Technology,” *Society* 13:6, September/October 1976. See also William H. Friedland, *Social Sleepwalkers: Scientific and Technological Research in California Agriculture*, University of California, Davis, Department of Applied Behavioral Sciences, Research Monograph No. 13, 1974.
12. *University of California Clip Sheet* 54:36, May 1, 1979.

-
13. "Tomato Technology."
 14. A history and critical analysis of agricultural research in the land-grant colleges is given in James Hightower, *Hard Tomatoes, Hard Times* (Cambridge: Schenkman, 1978).
 15. David F. Noble, *Forces of Production: A Social History of Machine Tool Automation* (New York: Alfred A. Knopf, 1984).
 16. Friedrich Engels, "On Authority," in *The Marx-Engels Reader*, ed. 2, Robert Tucker (ed.) (New York: W. W. Norton, 1978), 731.
 17. Ibid.
 18. Ibid., 732, 731.
 19. Karl Marx, *Capital*, vol. 1, ed. 3, translated by Samuel Moore and Edward Aveling (New York: Modern Library, 1906), 530.
 20. Jerry Mander, *Four Arguments for the Elimination of Television* (New York: William Morrow, 1978), 44.
 21. See, for example, Robert Argue, Barbara Emanuel, and Stephen Graham, *The Sun Builders: A People's Guide to Solar, Wind and Wood Energy in Canada* (Toronto: Renewable Energy in Canada, 1978). "We think decentralization is an implicit component of renewable energy; this implies the de centralization of energy systems, communities and of power. Renewable energy doesn't require mammoth generation sources of disruptive transmission corridors. Our cities and towns, which have been dependent on centralized energy supplies, may be able to achieve some degree of autonomy, thereby controlling and administering their own energy needs." (16)
 22. Alfred D. Chandler, Jr., *The Visible Hand: The Managerial Revolution in American Business* (Cambridge: Belknap, 1977), 244.
 23. Ibid.
 24. Ibid., 500.
 25. Leonard Silk and David Vogel, *Ethics and Profits: The Crisis of Confidence in American Business* (New York: Simon and Schuster, 1976), 191.
 26. Russell W. Ayres, "Policing Plutonium: The Civil Liberties Fallout," *Harvard Civil Rights—Civil Liberties Law Review* 10 (1975): 443, 413-414, 374.

Five Things We Need to Know About Technological Change

by
Neil Postman

Talk delivered in Denver Colorado
March 28, 1998

Good morning your Eminences and Excellencies, ladies, and gentlemen.

The theme of this conference, “The New Technologies and the Human Person: Communicating the Faith in the New Millennium,” suggests, of course, that you are concerned about what might happen to faith in the new millennium, as well you should be. In addition to our computers, which are close to having a nervous breakdown in anticipation of the year 2000, there is a great deal of frantic talk about the 21st century and how it will pose for us unique problems of which we know very little but for which, nonetheless, we are supposed to carefully prepare. Everyone seems to worry about this—business people, politicians, educators, as well as theologians.

At the risk of sounding patronizing, may I try to put everyone’s mind at ease? I doubt that the 21st century will pose for us problems that are more stunning, disorienting or complex than those we faced in this century, or the 19th, 18th, 17th, or for that matter, many of the centuries before that. But for those who are excessively nervous about the new millennium, I can provide, right at the start, some good advice about how to confront it. The advice comes from people whom we can trust, and whose thoughtfulness, it’s safe to say, exceeds that of President Clinton, Newt Gingrich, or even Bill Gates. Here is what Henry David Thoreau told us: “All our inventions are but improved means to an unimproved end.” Here is what Goethe told us: “One should, each day, try to hear a little song, read a good poem, see a fine picture, and, if possible, speak a few reasonable words.” Socrates told us: “The unexamined life is not worth living.” Rabbi Hillel told us: “What is hateful to thee, do not do to another.” And here is the prophet Micah: “What does the Lord require of thee but to do justly, to love mercy and to walk humbly with thy God.” And I could say, if we had the time, (although you know it well enough) what Jesus, Isaiah, Mohammad, Spinoza, and Shakespeare told us. It is all the same: There is no escaping from ourselves. The human dilemma is as it has always been, and it is a delusion to believe that the technological changes of our era have rendered irrelevant the wisdom of the ages and the sages.

Nonetheless, having said this, I know perfectly well that because we do live in a technological age, we have some special problems that Jesus, Hillel, Socrates, and Micah did not and could not speak of. I do not have the wisdom to say what we ought to do about such problems, and so my contribution must confine itself to some things we need to know in order to address the problems. I call my talk *Five Things We Need to Know About Technological Change*. I base these ideas on my thirty years of studying the history of technological change but I do not think these are academic or esoteric ideas. They are to the sort of things everyone who is concerned with cultural stability and balance should know and I offer them to you in the hope that you will find them useful in thinking about the effects of technology on religious faith.

First Idea

The first idea is that all technological change is a trade-off. I like to call it a Faustian bargain. Technology giveth and technology taketh away. This means that for every advantage a new technology offers, there is always a corresponding disadvantage. The disadvantage may exceed in importance the advantage, or the advantage may well be worth the cost. Now, this may seem to be a rather obvious idea, but you would be surprised at how many people believe that new technologies are unmixed blessings. You need only think of the enthusiasms with which most people approach their understanding of computers. Ask anyone who knows something about computers to talk about them, and you will find that they will, unabashedly and relentlessly, extol the wonders of computers. You will also find that in most cases they will completely

neglect to mention any of the liabilities of computers. This is a dangerous imbalance, since the greater the wonders of a technology, the greater will be its negative consequences.

Think of the automobile, which for all of its obvious advantages, has poisoned our air, choked our cities, and degraded the beauty of our natural landscape. Or you might reflect on the paradox of medical technology which brings wondrous cures but is, at the same time, a demonstrable cause of certain diseases and disabilities, and has played a significant role in reducing the diagnostic skills of physicians. It is also well to recall that for all of the intellectual and social benefits provided by the printing press, its costs were equally monumental. The printing press gave the Western world prose, but it made poetry into an exotic and elitist form of communication. It gave us inductive science, but it reduced religious sensibility to a form of fanciful superstition. Printing gave us the modern conception of nationhood, but in so doing turned patriotism into a sordid if not lethal emotion. We might even say that the printing of the Bible in vernacular languages introduced the impression that God was an Englishman or a German or a Frenchman—that is to say, printing reduced God to the dimensions of a local potentate.

Perhaps the best way I can express this idea is to say that the question, “What will a new technology do?” is no more important than the question, “What will a new technology undo?” Indeed, the latter question is more important, precisely because it is asked so infrequently. One might say, then, that a sophisticated perspective on technological change includes one’s being skeptical of Utopian and Messianic visions drawn by those who have no sense of history or of the precarious balances on which culture depends. In fact, if it were up to me, I would forbid anyone from talking about the new information technologies unless the person can demonstrate that he or she knows something about the social and psychic effects of the alphabet, the mechanical clock, the printing press, and telegraphy. In other words, knows something about the costs of great technologies.

Idea Number One, then, is that culture always pays a price for technology.

Second Idea

This leads to the second idea, which is that the advantages and disadvantages of new technologies are never distributed evenly among the population. This means that every new technology benefits some and harms others. There are even some who are not affected at all. Consider again the case of the printing press in the 16th century, of which Martin Luther said it was “God’s highest and extremest act of grace, whereby the business of the gospel is driven forward.” By placing the word of God on every Christian’s kitchen table, the mass-produced book undermined the authority of the church hierarchy, and hastened the breakup of the Holy Roman See. The Protestants of that time cheered this development. The Catholics were enraged and distraught. Since I am a Jew, had I lived at that time, I probably wouldn’t have given a damn one way or another, since it would make no difference whether a pogrom was inspired by Martin Luther or Pope Leo X. Some gain, some lose, a few remain as they were.

Let us take as another example, television, although here I should add at once that in the case of television there are very few indeed who are not affected in one way or another. In America, where television has taken hold more deeply than anywhere else, there are many people who find it a blessing, not least those who have achieved high-paying, gratifying careers in television as executives, technicians, directors, newscasters and entertainers. On the other hand, and in the long run, television may bring an end to the careers of school teachers since school was an invention of the printing press and must stand or fall on the issue of how much importance the printed word will have in the future. There is no chance, of course, that television will go away but school teachers who are enthusiastic about its presence always call to my mind an image of some turn-of-the-century blacksmith who not only is singing the praises of the automobile but who also believes that his business will be enhanced by it. We know now that his business was not enhanced by it; it was rendered obsolete by it, as perhaps an intelligent blacksmith would have known.

The questions, then, that are never far from the mind of a person who is knowledgeable about technological change are these: Who specifically benefits from the development of a new technology? Which groups,

what type of person, what kind of industry will be favored? And, of course, which groups of people will thereby be harmed?

These questions should certainly be on our minds when we think about computer technology. There is no doubt that the computer has been and will continue to be advantageous to large-scale organizations like the military or airline companies or banks or tax collecting institutions. And it is equally clear that the computer is now indispensable to high-level researchers in physics and other natural sciences. But to what extent has computer technology been an advantage to the masses of people? To steel workers, vegetable store owners, automobile mechanics, musicians, bakers, bricklayers, dentists, yes, theologians, and most of the rest into whose lives the computer now intrudes? These people have had their private matters made more accessible to powerful institutions. They are more easily tracked and controlled; they are subjected to more examinations, and are increasingly mystified by the decisions made about them. They are more than ever reduced to mere numerical objects. They are being buried by junk mail. They are easy targets for advertising agencies and political institutions.

In a word, these people are losers in the great computer revolution. The winners, which include among others computer companies, multi-national corporations and the nation state, will, of course, encourage the losers to be enthusiastic about computer technology. That is the way of winners, and so in the beginning they told the losers that with personal computers the average person can balance a checkbook more neatly, keep better track of recipes, and make more logical shopping lists. Then they told them that computers will make it possible to vote at home, shop at home, get all the entertainment they wish at home, and thus make community life unnecessary. And now, of course, the winners speak constantly of the Age of Information, always implying that the more information we have, the better we will be in solving significant problems—not only personal ones but large-scale social problems, as well. But how true is this? If there are children starving in the world—and there are—it is not because of insufficient information. We have known for a long time how to produce enough food to feed every child on the planet. How is it that we let so many of them starve? If there is violence on our streets, it is not because we have insufficient information. If women are abused, if divorce and pornography and mental illness are increasing, none of it has anything to do with insufficient information. I dare say it is because something else is missing, and I don't think I have to tell this audience what it is. Who knows? This age of information may turn out to be a curse if we are blinded by it so that we cannot see truly where our problems lie. That is why it is always necessary for us to ask of those who speak enthusiastically of computer technology, why do you do this? What interests do you represent? To whom are you hoping to give power? From whom will you be withholding power?

I do not mean to attribute unsavory, let alone sinister motives to anyone. I say only that since technology favors some people and harms others, these are questions that must always be asked. And so, that there are always winners and losers in technological change is the second idea.

Third Idea

Here is the third. Embedded in every technology there is a powerful idea, sometimes two or three powerful ideas. These ideas are often hidden from our view because they are of a somewhat abstract nature. But this should not be taken to mean that they do not have practical consequences.

Perhaps you are familiar with the old adage that says: To a man with a hammer, everything looks like a nail. We may extend that truism: To a person with a pencil, everything looks like a sentence. To a person with a TV camera, everything looks like an image. To a person with a computer, everything looks like data. I do not think we need to take these aphorisms literally. But what they call to our attention is that every technology has a prejudice. Like language itself, it predisposes us to favor and value certain perspectives and accomplishments. In a culture without writing, human memory is of the greatest importance, as are the proverbs, sayings and songs which contain the accumulated oral wisdom of centuries. That is why Solomon was thought to be the wisest of men. In Kings I we are told he knew 3,000 proverbs. But in a culture with writing, such feats of memory are considered a waste of time, and proverbs are merely irrelevant fancies. The writing person favors logical organization and systematic analysis, not proverbs. The telegraphic person values speed, not introspection. The television person values immediacy, not history. And computer

people, what shall we say of them? Perhaps we can say that the computer person values information, not knowledge, certainly not wisdom. Indeed, in the computer age, the concept of wisdom may vanish altogether.

The third idea, then, is that every technology has a philosophy which is given expression in how the technology makes people use their minds, in what it makes us do with our bodies, in how it codifies the world, in which of our senses it amplifies, in which of our emotional and intellectual tendencies it disregards. This idea is the sum and substance of what the great Catholic prophet, Marshall McLuhan meant when he coined the famous sentence, "The medium is the message."

Fourth Idea

Here is the fourth idea: Technological change is not additive; it is ecological. I can explain this best by an analogy. What happens if we place a drop of red dye into a beaker of clear water? Do we have clear water plus a spot of red dye? Obviously not. We have a new coloration to every molecule of water. That is what I mean by ecological change. A new medium does not add something; it changes everything. In the year 1500, after the printing press was invented, you did not have old Europe plus the printing press. You had a different Europe. After television, America was not America plus television. Television gave a new coloration to every political campaign, to every home, to every school, to every church, to every industry, and so on.

That is why we must be cautious about technological innovation. The consequences of technological change are always vast, often unpredictable and largely irreversible. That is also why we must be suspicious of capitalists. Capitalists are by definition not only personal risk takers but, more to the point, cultural risk takers. The most creative and daring of them hope to exploit new technologies to the fullest, and do not much care what traditions are overthrown in the process or whether or not a culture is prepared to function without such traditions. Capitalists are, in a word, radicals. In America, our most significant radicals have always been capitalists--men like Bell, Edison, Ford, Carnegie, Sarnoff, Goldwyn. These men obliterated the 19th century, and created the 20th, which is why it is a mystery to me that capitalists are thought to be conservative. Perhaps it is because they are inclined to wear dark suits and grey ties.

I trust you understand that in saying all this, I am making no argument for socialism. I say only that capitalists need to be carefully watched and disciplined. To be sure, they talk of family, marriage, piety, and honor but if allowed to exploit new technology to its fullest economic potential, they may undo the institutions that make such ideas possible. And here I might just give two examples of this point, taken from the American encounter with technology. The first concerns education. Who, we may ask, has had the greatest impact on American education in this century? If you are thinking of John Dewey or any other education philosopher, I must say you are quite wrong. The greatest impact has been made by quiet men in grey suits in a suburb of New York City called Princeton, New Jersey. There, they developed and promoted the technology known as the standardized test, such as IQ tests, the SATs and the GREs. Their tests redefined what we mean by learning, and have resulted in our reorganizing the curriculum to accommodate the tests.

A second example concerns our politics. It is clear by now that the people who have had the most radical effect on American politics in our time are not political ideologues or student protesters with long hair and copies of Karl Marx under their arms. The radicals who have changed the nature of politics in America are entrepreneurs in dark suits and grey ties who manage the large television industry in America. They did not mean to turn political discourse into a form of entertainment. They did not mean to make it impossible for an overweight person to run for high political office. They did not mean to reduce political campaigning to a 30-second TV commercial. All they were trying to do is to make television into a vast and unsleeping money machine. That they destroyed substantive political discourse in the process does not concern them.

Fifth Idea

I come now to the fifth and final idea, which is that media tend to become mythic. I use this word in the sense in which it was used by the French literary critic, Roland Barthes. He used the word “myth” to refer to a common tendency to think of our technological creations as if they were God-given, as if they were a part of the natural order of things. I have on occasion asked my students if they know when the alphabet was invented. The question astonishes them. It is as if I asked them when clouds and trees were invented. The alphabet, they believe, was not something that was invented. It just is. It is this way with many products of human culture but with none more consistently than technology. Cars, planes, TV, movies, newspapers--they have achieved mythic status because they are perceived as gifts of nature, not as artifacts produced in a specific political and historical context.

When a technology become mythic, it is always dangerous because it is then accepted as it is, and is therefore not easily susceptible to modification or control. If you should propose to the average American that television broadcasting should not begin until 5 PM and should cease at 11 PM, or propose that there should be no television commercials, he will think the idea ridiculous. But not because he disagrees with your cultural agenda. He will think it ridiculous because he assumes you are proposing that something in nature be changed; as if you are suggesting that the sun should rise at 10 AM instead of at 6.

Whenever I think about the capacity of technology to become mythic, I call to mind the remark made by Pope John Paul II. He said, “Science can purify religion from error and superstition. Religion can purify science from idolatry and false absolutes.”

What I am saying is that our enthusiasm for technology can turn into a form of idolatry and our belief in its beneficence can be a false absolute. The best way to view technology is as a strange intruder, to remember that technology is not part of God’s plan but a product of human creativity and hubris, and that its capacity for good or evil rests entirely on human awareness of what it does for us and to us.

Conclusion

And so, these are my five ideas about technological change. First, that we always pay a price for technology; the greater the technology, the greater the price. Second, that there are always winners and losers, and that the winners always try to persuade the losers that they are really winners. Third, that there is embedded in every great technology an epistemological, political or social prejudice. Sometimes that bias is greatly to our advantage. Sometimes it is not. The printing press annihilated the oral tradition; telegraphy annihilated space; television has humiliated the word; the computer, perhaps, will degrade community life. And so on. Fourth, technological change is not additive; it is ecological, which means, it changes everything and is, therefore, too important to be left entirely in the hands of Bill Gates. And fifth, technology tends to become mythic; that is, perceived as part of the natural order of things, and therefore tends to control more of our lives than is good for us.

If we had more time, I could supply some additional important things about technological change but I will stand by these for the moment, and will close with this thought. In the past, we experienced technological change in the manner of sleep-walkers. Our unspoken slogan has been “technology über alles,” and we have been willing to shape our lives to fit the requirements of technology, not the requirements of culture. This is a form of stupidity, especially in an age of vast technological change. We need to proceed with our eyes wide open so that we may use technology rather than be used by it.

Marshall McLuhan Interview

Redacted (PR) from *The Playboy Interview: Marshall McLuhan*, which appeared in *Playboy Magazine*, March 1969, ©1994. This file was last edited 4/30/08.

In 1961, the name of Marshall McLuhan was unknown to everyone but his English students at the University of Toronto—and a coterie of academic admirers who followed his abstruse articles in small-circulation quarterlies. But then came two remarkable books—“The Gutenberg Galaxy” (1962) and “Understanding Media” (1964)—and the graying professor from Canada’s western hinterlands soon found himself characterized by the *San Francisco Chronicle* as “the hottest academic property around.” He has since won a world-wide following for his brilliant—and frequently baffling—theories about the impact of the media on man; and his name has entered the French language as *mucluhanisme*, a synonym for the world of pop culture.

Though his books are written in a difficult style—at once enigmatic, epigrammatic and overgrown with arcane literary and historic allusions—the revolutionary ideas lurking in them have made McLuhan a best-selling author. Despite protests from a legion of outraged scholastics and old-guard humanists who claim that McLuhan’s ideas range from demented to dangerous, his free-for-all theorizing has attracted the attention of top executives at General Motors (who paid him a handsome fee to inform them that automobiles were a thing of the past), Bell Telephone (to whom he explained that they didn’t really understand the function of the telephone) and a leading package-design house (which was told that packages will soon be obsolete). Anteing up \$5000, another huge corporation asked him to predict—via closed-circuit television—what their own products will be used for in the future; and Canada’s turned-on Prime Minister Pierre Trudeau engages him in monthly bull sessions designed to improve his television image.

McLuhan’s observations—“probes,” he prefers to call them—are riddled with such flamboyantly undecipherable aphorisms as “The electric light is pure information” and “People don’t actually read newspapers—they get into them every morning like a hot bath.” Of his own work, McLuhan has remarked: “I don’t pretend to understand it. After all, my stuff is very difficult.” Despite his convoluted syntax, flashy metaphors and word-playful one-liners, however, McLuhan’s basic thesis is relatively simple.

McLuhan contends that all media—in and of themselves and regardless of the messages they communicate—exert a compelling influence on man and society. Prehistoric, or tribal, man existed in a harmonious balance of the senses, perceiving the world equally through hearing, smell, touch, sight and taste. But technological innovations are extensions of human abilities and senses that alter this sensory balance—an alteration that, in turn, inexorably reshapes the society that created the technology. According to McLuhan, there have been three basic technological innovations: the invention of the phonetic alphabet, which jolted tribal man out of his sensory balance and gave dominance to the eye; the introduction of movable type in the 16th Century, which accelerated this process; and the invention of the telegraph in 1844, which heralded an electronics revolution that will ultimately retribalize man by restoring his sensory balance. McLuhan has made it his business to explain and extrapolate the repercussions of this electronic revolution.

For his efforts, critics have dubbed him “the Dr. Spock of pop culture,” “the guru of the boob tube,” a “Canadian Nkrumah who has joined the assault on reason,” a “metaphysical wizard possessed by a spatial sense of madness,” and “the high priest of popthink who conducts a Black Mass for dilettantes before the altar of historical determinism.” Amherst professor Benjamin De-Mott observed: “He’s swinging, switched on, with it and NOW. And wrong.”

But as Tom Wolfe has aptly inquired, “What if he is right? Suppose he is what he sounds like—the most important thinker since Newton, Darwin, Freud, Einstein and Pavlov?” Social historian Richard Kostelanetz contends that “the most extraordinary quality of McLuhan’s mind is that it discerns significance where others see only data, or nothing; he tells us how to measure phenomena previously unmeasurable.”

The unperturbed subject of this controversy was born in Edmonton, Alberta, on July 21, 1911. The son of a former actress and a real-estate salesman, McLuhan entered the University of Manitoba intending to become an engineer, but matriculated in 1934 with an M.A. in English literature. Next came a stint as an oarsman and graduate student at Cambridge, followed by McLuhan’s first teaching job—at the University of Wisconsin. It was a pivotal experience. “I was confronted with young Americans I was incapable of understanding,” he has since remarked. “I felt an urgent need to study their popular culture in order to

get through.” With the seeds sown, McLuhan let them germinate while earning a Ph.D., then taught at Catholic universities. (He is a devout Roman Catholic convert.)

His publishing career began with a number of articles on standard academic fare; but by the mid-Forties, his interest in popular culture surfaced, and true McLuhan efforts such as “The Psychopathology of Time and Life” began to appear. They hit book length for the first time in 1951 with the publication of “The Mechanical Bride”—an analysis of the social and psychological pressures generated by the press, radio, movies and advertising—and McLuhan was on his way. Though the book attracted little public notice, it won him the chairmanship of a Ford Foundation seminar on culture and communications and a \$40,000 grant, with part of which he started “Explorations,” a small periodical outlet for the seminar’s findings. By the late Fifties, his reputation had trickled down to Washington: In 1959, he became director of the Media Project of the National Association of Educational Broadcasters and the United States Office of Education, and the report resulting from this post became the first draft of “Understanding Media.” Since 1963, McLuhan has headed the University of Toronto’s Center for Culture and Technology, which until recently consisted entirely of McLuhan’s office, but now includes a six-room campus building.

Apart from his teaching, lecturing and administrative duties, McLuhan has become a sort of minor communication industry unto himself. Each month he issues to subscribers a mixed-media report called “The McLuhan Dew-Line”; and, punning on that title, he has also originated a series of recordings called “The Marshall McLuhan Dew-Line Plattertudes.” McLuhan contributed a characteristically mind-expanding essay about the media—“The Reversal of the Overheated-Image”—to our December 1968 issue. Also a compulsive collaborator, his literary efforts in tandem with colleagues have included a high school textbook and an analysis of the function of space in poetry and painting. “Counterblast,” his next book, is a manically graphic trip through the land of his theories.

In order to provide our readers with a map of this labyrinthine terra incognita, PLAYBOY assigned interviewer **Eric Norden** to visit McLuhan at his spacious new home in the wealthy Toronto suburb of Wychwood Park, where he lives with his wife, Corinne, and five of his six children. (His eldest son lives in New York, where he is completing a book on James Joyce, one of his father’s heroes.) Norden reports: “Tall, gray and gangly, with a thin but mobile mouth and an otherwise eminently forgettable face, McLuhan was dressed in an ill-fitting brown tweed suit, black shoes and a clip-on necktie. As we talked on into the night before a crackling fire, McLuhan expressed his reservations about the interview—indeed, about the printed word itself—as a means of communication, suggesting that the question-and-answer format might impede the in-depth flow of his ideas. I assured him that he would have as much time—and space—as he wished to develop his thoughts.”

The result has considerably more lucidity and clarity than McLuhan’s readers are accustomed to—perhaps because the Q. and A. format serves to pin him down by counteracting his habit of mercurially changing the subject in mid-stream of consciousness. It is also, we think, a protean and provocative distillation not only of McLuhan’s original theories about human progress and social institutions but of his almost immobilizingly intricate style—described by novelist George P. Elliott as “deliberately antilogical, circular, repetitious, unqualified, gnomic, outrageous” and, even less charitably, by critic Christopher Ricks as “a viscous fog through which loom stumbling metaphors.” But other authorities contend that McLuhan’s stylistic medium is part and parcel of his message—that the tightly structured “linear” modes of traditional thought and discourse are obsolescent in the new “postliterate” age of the electric media. Norden began the interview with an allusion to McLuhan’s favorite electric medium: television. *The Interview:*

Interviewer: To borrow Henry Gibson’s oft-repeated one-line poem on Rowan and Martin’s Laugh-In—“Marshall McLuhan, what are you doin’?”

McLuhan: Sometimes I wonder. I’m making explorations. I don’t know where they’re going to take me. My work is designed for the pragmatic purpose of trying to understand our technological environment and its psychic and social consequences. But my books constitute the process rather than the completed product of discovery; my purpose is to employ facts as tentative probes, as means of insight, of pattern recognition, rather than to use them in the traditional and sterile sense of classified data, categories, containers. I want to map new terrain rather than chart old landmarks.

But I’ve never presented such explorations as revealed truth. As an investigator, I have no fixed point of view, no commitment to any theory—my own or anyone else’s. As a matter of fact, I’m completely ready

to junk any statement I've ever made about any subject if events don't bear me out, or if I discover it isn't contributing to an understanding of the problem. The better part of my work on media is actually somewhat like a safe-cracker's. I don't know what's inside; maybe it's nothing. I just sit down and start to work. I grope, I listen, I test, I accept and discard; I try out different sequences—until the tumblers fall and the doors spring open.

Interviewer: Isn't such a methodology somewhat erratic and inconsistent—if not, as your critics would maintain, eccentric?

McLuhan: Any approach to environmental problems must be sufficiently flexible and adaptable to encompass the entire environmental matrix, which is in constant flux. I consider myself a generalist, not a specialist who has staked out a tiny plot of study as his intellectual turf and is oblivious to everything else. Actually, my work is a depth operation, the accepted practice in most modern disciplines from psychiatry to metallurgy and structural analysis. Effective study of the media deals not only with the content of the media but with the media themselves and the total cultural environment within which the media function. Only by standing aside from any phenomenon and taking an overview can you discover its operative principles and lines of force. There's really nothing inherently startling or radical about this study—except that for some reason few have had the vision to undertake it. For the past 3500 years of the Western world, the effects of media—whether it's speech, writing, printing, photography, radio or television—have been systematically overlooked by social observers. Even in today's revolutionary electronic age, scholars evidence few signs of modifying this traditional stance of ostrichlike disregard.

Interviewer: Why?

McLuhan: Because all media, from the phonetic alphabet to the computer, are extensions of man that cause deep and lasting changes in him and transform his environment. Such an extension is an intensification, an amplification of an organ, sense or function, and whenever it takes place, the central nervous system appears to institute a self-protective numbing of the affected area, insulating and anesthetizing it from conscious awareness of what's happening to it. It's a process rather like that which occurs to the body under shock or stress conditions, or to the mind in line with the Freudian concept of repression. I call this peculiar form of self-hypnosis *Narcissus narcosis*, a syndrome whereby man remains as unaware of the psychic and social effects of his new technology as a fish of the water it swims in. As a result, precisely at the point where a new media-induced environment becomes all pervasive and transmogrifies our sensory balance, it also becomes invisible.

This problem is doubly acute today because man must, as a simple survival strategy, become aware of what is happening to him, despite the attendant pain of such comprehension. The fact that he has not done so in this age of electronics is what has made this also the age of anxiety, which in turn has been transformed into its Doppelgänger—the therapeutically reactive age of anomie and apathy. But despite our self-protective escape mechanisms, the total-field awareness engendered by electronic media is enabling us—indeed, compelling us—to grope toward a consciousness of the unconscious, toward a realization that technology is an extension of our own bodies. We live in the first age when change occurs sufficiently rapidly to make such pattern recognition possible for society at large. Until the present era, this awareness has always been reflected first by the artist, who has had the power—and courage—of the seer to read the language of the outer world and relate it to the inner world.

Interviewer: Why should it be the artist rather than the scientist who perceives these relationships and foresees these trends?

McLuhan: Because inherent in the artist's creative inspiration is the process of subliminally sniffing out environmental change. It's always been the artist who perceives the alterations in man caused by a new medium, who recognizes that the future is the present, and uses his work to prepare the ground for it. But most people, from truck drivers to the literary Brahmins, are still blissfully ignorant of what the media do to them; unaware that because of their pervasive effects on man, it is the medium itself that is the message, not the content, and unaware that the medium is also the message—that, all puns aside, it literally works over and saturates and molds and transforms every sense ratio. The content or message of any particular medium has about as much importance as the stenciling on the casing of an atomic bomb. But the ability to perceive media-induced extensions of man, once the province of the artist, is now being expanded as the

new environment of electric information makes possible a new degree of perception and critical awareness by nonartists.

Interviewer: Is the public, then, at last beginning to perceive the “invisible” contours of these new technological environments

McLuhan: People are beginning to understand the nature of their new technology, but not yet nearly enough of them—and not nearly well enough. Most people, as I indicated, still cling to what I call the rearview-mirror view of their world. By this I mean to say that because of the invisibility of any environment during the period of its innovation, man is only consciously aware of the environment that has preceded it; in other words, an environment becomes fully visible only when it has been superseded by a new environment; thus we are always one step behind in our view of the world. Because we are benumbed by any new technology—which in turn creates a totally new environment—we tend to make the old environment more visible; we do so by turning it into an art form and by attaching ourselves to the objects and atmosphere that characterized it, just as we’ve done with jazz, and as we’re now doing with the garbage of the mechanical environment via pop art.

The present is always invisible because it’s environmental and saturates the whole field of attention so overwhelmingly; thus everyone but the artist, the man of integral awareness, is alive in an earlier day. In the midst of the electronic age of software, of instant information movement, we still believe we’re living in the mechanical age of hardware. At the height of the mechanical age, man turned back to earlier centuries in search of “pastoral” values. The Renaissance and the Middle Ages were completely oriented toward Rome; Rome was oriented toward Greece, and the Greeks were oriented toward the pre-Homeric primitives. We reverse the old educational dictum of learning by proceeding from the familiar to the unfamiliar by going from the unfamiliar to the familiar, which is nothing more or less than the numbing mechanism that takes place whenever new media drastically extend our senses.

Interviewer: If this “numbing” effect performs a beneficial role by protecting man from the psychic pain caused by the extensions of his nervous system that you attribute to the media, why are you attempting to dispel it and alert man to the changes in his environment?

McLuhan: In the past, the effects of media were experienced more gradually, allowing the individual and society to absorb and cushion their impact to some degree. Today, in the electronic age of instantaneous communication, I believe that our survival, and at the very least our comfort and happiness, is predicated on understanding the nature of our new environment, because unlike previous environmental changes, the electric media constitute a total and near-instantaneous transformation of culture, values and attitudes. This upheaval generates great pain and identity loss, which can be ameliorated only through a conscious awareness of its dynamics. If we understand the revolutionary transformations caused by new media, we can anticipate and control them; but if we continue in our self-induced subliminal trance, we will be their slaves.

Because of today’s terrific speed-up of information moving, we have a chance to apprehend, predict and influence the environmental forces shaping us—and thus win back control of our own destinies. The new extensions of man and the environment they generate are the central manifestations of the evolutionary process, and yet we still cannot free ourselves of the delusion that it is how a medium is used that counts, rather than what it does to us and with us. This is the zombie stance of the technological idiot. It’s to escape this Narcissus trance that I’ve tried to trace and reveal the impact of media on man, from the beginning of recorded time to the present.

Interviewer: Will you trace that impact for us—in condensed form?

McLuhan: It’s difficult to condense into the format of an interview such as this, but I’ll try to give you a brief rundown of the basic media breakthroughs. You’ve got to remember that my definition of media is broad; it includes any technology whatever that creates extensions of the human body and senses, from clothing to the computer. And a vital point I must stress again is that societies have always been shaped more by the nature of the media with which men communicate than by the content of the communication. All technology has the property of the Midas touch; whenever a society develops an extension of itself, all other functions of that society tend to be transmuted to accommodate that new form; once any new technology penetrates a society, it saturates every institution of that society. New technology is thus a revolutionizing agent. We see this today with the electric media and we saw it several thousand years ago

with the invention of the phonetic alphabet, which was just as far-reaching an innovation—and had just as profound consequences for man.

Interviewer: What were they?

McLuhan: Before the invention of the phonetic alphabet, man lived in a world where all the senses were balanced and simultaneous, a closed world of tribal depth and resonance, an oral culture structured by a dominant auditory sense of life. The ear, as opposed to the cool and neutral eye, is sensitive, hyperaesthetic and all-inclusive, and contributes to the seamless web of tribal kinship and interdependence in which all members of the group existed in harmony. The primary medium of communication was speech, and thus no man knew appreciably more or less than any other—which meant that there was little individualism and specialization, the hallmarks of “civilized” Western man. Tribal cultures even today simply cannot comprehend the concept of the individual or of the separate and independent citizen. Oral cultures act and react simultaneously, whereas the capacity to act without reacting, without involvement, is the special gift of “detached” literate man. Another basic characteristic distinguishing tribal man from his literate successors is that he lived in a world of acoustic space, which gave him a radically different concept of time-space relationships.

Interviewer: What do you mean by “acoustic space”?

McLuhan: I mean space that has no center and no margin, unlike strictly visual space, which is an extension and intensification of the eye. Acoustic space is organic and integral, perceived through the simultaneous interplay of all the senses; whereas “rational” or pictorial space is uniform, sequential and continuous and creates a closed world with none of the rich resonance of the tribal echoland. Our own Western time-space concepts derive from the environment created by the discovery of phonetic writing, as does our entire concept of Western civilization. The man of the tribal world led a complex, kaleidoscopic life precisely because the ear, unlike the eye, cannot be focused and is synaesthetic rather than analytical and linear. Speech is an utterance, or more precisely, an outering, of all our senses at once; the auditory field is simultaneous, the visual successive. The models of life of nonliterate people were implicit, simultaneous and discontinuous, and also far richer than those of literate man. By their dependence on the spoken word for information, people were drawn together into a tribal mesh; and since the spoken word is more emotionally laden than the written—conveying by intonation such rich emotions as anger, joy, sorrow, fear—tribal man was more spontaneous and passionately volatile. Audile-tactile tribal man partook of the collective unconscious, lived in a magical integral world patterned by myth and ritual, its values divine and unchallenged, whereas literate or visual man creates an environment that is strongly fragmented, individualistic, explicit, logical, specialized and detached.

Interviewer: Was it phonetic literacy alone that precipitated this profound shift of values from tribal involvement to “civilized” detachment?

McLuhan: Yes, it was. Any culture is an order of sensory preferences, and in the tribal world, the senses of touch, taste, hearing and smell were developed, for very practical reasons, to a much higher level than the strictly visual. Into this world, the phonetic alphabet fell like a bombshell, installing sight at the head of the hierarchy of senses. Literacy propelled man from the tribe, gave him an eye for an ear and replaced his integral in-depth communal interplay with visual linear values and fragmented consciousness. As an intensification and amplification of the visual function, the phonetic alphabet diminished the role of the senses of hearing and touch and taste and smell, permeating the discontinuous culture of tribal man and translating its organic harmony and complex synaesthesia into the uniform, connected and visual mode that we still consider the norm of “rational” existence. The whole man became fragmented man; the alphabet shattered the charmed circle and resonating magic of the tribal world, exploding man into an agglomeration of specialized and psychically impoverished “individuals,” or units, functioning in a world of linear time and Euclidean space.

Interviewer: But literate societies existed in the ancient world long before the phonetic alphabet. Why weren’t they detribalized?

McLuhan: The phonetic alphabet did not change or extend man so drastically just because it enabled him to read; as you point out, tribal culture had already coexisted with other written languages for thousands of years. But the phonetic alphabet was radically different from the older and richer hieroglyphic or

ideogrammic cultures. The writings of Egyptian, Babylonian, Mayan and Chinese cultures were an extension of the senses in that they gave pictorial expression to reality, and they demanded many signs to cover the wide range of data in their societies—unlike phonetic writing, which uses semantically meaningless letters to correspond to semantically meaningless sounds and is able, with only a handful of letters, to encompass all meanings and all languages. This achievement demanded the separation of both sights and sounds from their semantic and dramatic meanings in order to render visible the actual sound of speech, thus placing a barrier between men and objects and creating a dualism between sight and sound. It divorced the visual function from the interplay with the other senses and thus led to the rejection from consciousness of vital areas of our sensory experience and to the resultant atrophy of the unconscious. The balance of the sensorium—or Gestalt interplay of all the senses—and the psychic and social harmony it engendered was disrupted, and the visual function was overdeveloped. This was true of no other writing system.

Interviewer: How can you be so sure that this all occurred solely because of phonetic literacy—or, in fact, if it occurred at all?

McLuhan: You don't have to go back 3000 or 4000 years to see this process at work; in Africa today, a single generation of alphabetic literacy is enough to wrench the individual from the tribal web. When tribal man becomes phonetically literate, he may have an improved abstract intellectual grasp of the world, but most of the deeply emotional corporate family feeling is excised from his relationship with his social milieu. This division of sight and sound and meaning causes deep psychological effects, and he suffers a corresponding separation and impoverishment of his imaginative, emotional and sensory life. He begins reasoning in a sequential linear fashion; he begins categorizing and classifying data. As knowledge is extended in alphabetic form, it is localized and fragmented into specialties, creating division of function, of social classes, of nations and of knowledge—and in the process, the rich interplay of all the senses that characterized the tribal society is sacrificed.

Interviewer: But aren't there corresponding gains in insight, understanding and cultural diversity to compensate detribalized man for the loss of his communal values?

McLuhan: Your question reflects all the institutionalized biases of literate man. Literacy, contrary to the popular view of the “civilizing” process you've just echoed, creates people who are much less complex and diverse than those who develop in the intricate web of oral-tribal societies. Tribal man, unlike homogenized Western man, was not differentiated by his specialist talents or his visible characteristics, but by his unique emotional blends. The internal world of the tribal man was a creative mix of complex emotions and feelings that literate men of the Western world have allowed to wither or have suppressed in the name of efficiency and practicality. The alphabet served to neutralize all these rich divergencies of tribal cultures by translating their complexities into simple visual forms; and the visual sense, remember, is the only one that allows us to detach; all other senses involve us, but the detachment bred by literacy disinvolves and detribalizes man. He separates from the tribe as a predominantly visual man who shares standardized attitudes, habits and rights with other civilized men. But he is also given a tremendous advantage over the nonliterate tribal man who, today as in ancient times, is hamstrung by cultural pluralism, uniqueness and discontinuity—values that make the African as easy prey for the European colonialist as the barbarian was for the Greeks and Romans. Only alphabetic cultures have ever succeeded in mastering connected linear sequences as a means of social and psychic organization; the separation of all kinds of experiences into uniform and continuous units in order to generate accelerated action and alteration of form—in other words, applied knowledge—has been the secret of Western man's ascendancy over other men as well as over his environment.

Interviewer: Isn't the thrust of your argument, then, that the introduction of the phonetic alphabet was not progress, as has generally been assumed, but a psychic and social disaster?

McLuhan: It was both. I try to avoid value judgments in these areas, but there is much evidence to suggest that man may have paid too dear a price for his new environment of specialist technology and values. Schizophrenia and alienation may be the inevitable consequences of phonetic literacy. It's metaphorically significant, I suspect, that the old Greek myth has Cadmus, who brought the alphabet to man, sowing dragon's teeth that sprang up from the earth as armed men. Whenever the dragon's teeth of technological change are sown, we reap a whirlwind of violence. We saw this clearly in classical times, although it was somewhat moderated because phonetic literacy did not win an overnight victory over primitive values and institutions; rather, it permeated ancient society in a gradual, if inexorable, evolutionary process.

Interviewer: How long did the old tribal culture endure?

McLuhan: In isolated pockets, it held on until the invention of printing in the 16th Century, which was a vastly important qualitative extension of phonetic literacy. If the phonetic alphabet fell like a bombshell on tribal man, the printing press hit him like a 100-megaton H-bomb. The printing press was the ultimate extension of phonetic literacy: Books could be reproduced in infinite numbers; universal literacy was at last fully possible, if gradually realized; and books became portable individual possessions. Type, the prototype of all machines, ensured the primacy of the visual bias and finally sealed the doom of tribal man. The new medium of linear, uniform, repeatable type reproduced information in unlimited quantities and at hitherto-impossible speeds, thus assuring the eye a position of total predominance in man's sensorium. As a drastic extension of man, it shaped and transformed his entire environment, psychic and social, and was directly responsible for the rise of such disparate phenomena as nationalism, the Reformation, the assembly line and its offspring, the Industrial Revolution, the whole concept of causality, Cartesian and Newtonian concepts of the universe, perspective in art, narrative chronology in literature and a psychological mode of introspection or inner direction that greatly intensified the tendencies toward individualism and specialization engendered 2000 years before by phonetic literacy. The schism between thought and action was institutionalized, and fragmented man, first sundered by the alphabet, was at last diced into bite-sized tidbits. From that point on, Western man was Gutenberg man.

Interviewer: Even accepting the principle that technological innovations generate far-reaching environmental changes, many of your readers find it difficult to understand how you can hold the development of printing responsible for such apparently unrelated phenomena as nationalism and industrialism.

McLuhan: The key word is "apparently." Look a bit closer at both nationalism and industrialism and you'll see that both derived directly from the explosion of print technology in the 16th Century. Nationalism didn't exist in Europe until the Renaissance, when typography enabled every literate man to see his mother tongue analytically as a uniform entity. The printing press, by spreading mass-produced books and printed matter across Europe, turned the vernacular regional languages of the day into uniform closed systems of national languages—just another variant of what we call mass media—and gave birth to the entire concept of nationalism.

The individual newly homogenized by print saw the nation concept as an intense and beguiling image of group destiny and status. With print, the homogeneity of money, markets and transport also became possible for the first time, thus creating economic as well as political unity and triggering all the dynamic centralizing energies of contemporary nationalism. By creating a speed of information movement unthinkable before printing, the Gutenberg revolution thus produced a new type of visual centralized national entity that was gradually merged with commercial expansion until Europe was a network of states.

By fostering continuity and competition within homogeneous and contiguous territory, nationalism not only forged new nations but sealed the doom of the old corporate, noncompetitive and discontinuous medieval order of guilds and family-structured social organization; print demanded both personal fragmentation and social uniformity, the natural expression of which was the nation-state. Literate nationalism's tremendous speed-up of information movement accelerated the specialist function that was nurtured by phonetic literacy and nourished by Gutenberg, and rendered obsolete such generalist encyclopedic figures as Benvenuto Cellini, the goldsmith-cum-condottiere-cum-painter-cum-sculptor-cum-writer; it was the Renaissance that destroyed Renaissance Man.

Interviewer: Why do you feel that Gutenberg also laid the groundwork for the Industrial Revolution?

McLuhan: The two go hand in hand. Printing, remember, was the first mechanization of a complex handicraft; by creating an analytic sequence of step-by-step processes, it became the blue-print of all mechanization to follow. The most important quality of print is its repeatability; it is a visual statement that can be reproduced indefinitely, and repeatability is the root of the mechanical principle that has transformed the world since Gutenberg. Typography, by producing the first uniformly repeatable commodity, also created Henry Ford, the first assembly line and the first mass production. Movable type was archetype and prototype for all subsequent industrial development. Without phonetic literacy and the printing press, modern industrialism would be impossible. It is necessary to recognize literacy as typographic technology, shaping not only production and marketing procedures but all other areas of life, from education to city planning.

Interviewer: You seem to be contending that practically every aspect of modern life is a direct consequence of Gutenberg's invention of the printing press.

McLuhan: Every aspect of Western mechanical culture was shaped by print technology, but the modern age is the age of the electric media, which forge environments and cultures antithetical to the mechanical consumer society derived from print. Print tore man out of his traditional cultural matrix while showing him how to pile individual upon individual into a massive agglomeration of national and industrial power, and the typographic trance of the West has endured until today, when the electronic media are at last demesmerizing us. The Gutenberg Galaxy is being eclipsed by the constellation of Marconi.

Interviewer: You've discussed that constellation in general terms, but what precisely are the electric media that you contend have supplanted the old mechanical technology?

McLuhan: The electric media are the telegraph, radio, films, telephone, computer and television, all of which have not only extended a single sense or function as the old mechanical media did—i.e., the wheel as an extension of the foot, clothing as an extension of the skin, the phonetic alphabet as an extension of the eye—but have enhanced and externalized our entire central nervous systems, thus transforming all aspects of our social and psychic existence. The use of the electronic media constitutes a break boundary between fragmented Gutenberg man and integral man, just as phonetic literacy was a break boundary between oral-tribal man and visual man.

In fact, today we can look back at 3000 years of differing degrees of visualization, atomization and mechanization and at last recognize the mechanical age as an interlude between two great organic eras of culture. The age of print, which held sway from approximately 1500 to 1900, had its obituary tapped out by the telegraph, the first of the new electric media, and further obsequies were registered by the perception of "curved space" and non-Euclidean mathematics in the early years of the century, which revived tribal man's discontinuous time-space concepts—and which even Spengler dimly perceived as the death knell of Western literate values. The development of telephone, radio, film, television and the computer have driven further nails into the coffin. Today, television is the most significant of the electric media because it permeates nearly every home in the country, extending the central nervous system of every viewer as it works over and molds the entire sensorium with the ultimate message. It is television that is primarily responsible for ending the visual supremacy that characterized all mechanical technology, although each of the other electric media have played contributing roles.

Interviewer: But isn't television itself a primarily visual medium?

McLuhan: No, it's quite the opposite, although the idea that TV is a visual extension is an understandable mistake. Unlike film or photograph, television is primarily an extension of the sense of touch rather than of sight, and it is the tactile sense that demands the greatest interplay of all the senses. The secret of TV's tactile power is that the video image is one of low intensity or definition and thus, unlike either photograph or film, offers no detailed information about specific objects but instead involves the active participation of the viewer. The TV image is a mosaic mesh not only of horizontal lines but of millions of tiny dots, of which the viewer is physiologically able to pick up only 50 or 60 from which he shapes the image; thus he is constantly filling in vague and blurry images, bringing himself into in-depth involvement with the screen and acting out a constant creative dialog with the iconoscope. The contours of the resultant cartoonlike image are fleshed out within the imagination of the viewer, which necessitates great personal involvement and participation; the viewer, in fact, becomes the screen, whereas in film he becomes the camera. By requiring us to constantly fill in the spaces of the mosaic mesh, the iconoscope is tattooing its message directly on our skins. Each viewer is thus an unconscious pointillist painter like Seurat, limning new shapes and images as the iconoscope washes over his entire body. Since the point of focus for a TV set is the viewer, television is Orientalizing us by causing us all to begin to look within ourselves. The essence of TV viewing is, in short, intense participation and low definition—what I call a "cool" experience, as opposed to an essentially "hot," or high definition-low participation, medium like radio.

Interviewer: A good deal of the perplexity surrounding your theories is related to this postulation of hot and cool media. Could you give us a brief definition of each?

McLuhan: Basically, a hot medium excludes and a cool medium includes; hot media are low in participation, or completion, by the audience and cool media are high in participation. A hot medium is one that extends a

single sense with high definition. High definition means a complete filling in of data by the medium without intense audience participation. A photograph, for example, is high definition or hot; whereas a cartoon is low definition or cool, because the rough outline drawing provides very little visual data and requires the viewer to fill in or complete the image himself. The telephone, which gives the ear relatively little data, is thus cool, as is speech; both demand considerable filling in by the listener. On the other hand, radio is a hot medium because it sharply and intensely provides great amounts of high-definition auditory information that leaves little or nothing to be filled in by the audience. A lecture, by the same token, is hot, but a seminar is cool; a book is hot, but a conversation or bull session is cool.

In a cool medium, the audience is an active constituent of the viewing or listening experience. A girl wearing open-mesh silk stockings or glasses is inherently cool and sensual because the eye acts as a surrogate hand in filling in the low-definition image thus engendered. Which is why boys make passes at girls who wear glasses. In any case, the overwhelming majority of our technologies and entertainments since the introduction of print technology have been hot, fragmented and exclusive, but in the age of television we see a return to cool values and the inclusive in-depth involvement and participation they engender. This is, of course, just one more reason why the medium is the message, rather than the content; it is the participatory nature of the TV experience itself that is important, rather than the content of the particular TV image that is being invisibly and indelibly inscribed on our skins.

Interviewer: Even if, as you contend, the medium is the ultimate message, how can you entirely discount the importance of content? Didn't the content of Hitler's radio speeches, for example, have some effect on the Germans?

McLuhan: By stressing that the medium is the message rather than the content, I'm not suggesting that content plays no role—merely that it plays a distinctly subordinate role. Even if Hitler had delivered botany lectures, some other demagog would have used the radio to retribalize the Germans and rekindle the dark atavistic side of the tribal nature that created European fascism in the Twenties and Thirties. By placing all the stress on content and practically none on the medium, we lose all chance of perceiving and influencing the impact of new technologies on man, and thus we are always dumfounded by—and unprepared for—the revolutionary environmental transformations induced by new media. Buffeted by environmental changes he cannot comprehend, man echoes the last plaintive cry of his tribal ancestor, Tarzan, as he plummeted to earth: "Who greased my vine?" The German Jew victimized by the Nazis because his old tribalism clashed with their new tribalism could no more understand why his world was turned upside down than the American today can understand the reconfiguration of social and political institutions caused by the electric media in general and television in particular.

Interviewer: How is television reshaping our political institutions?

McLuhan: TV is revolutionizing every political system in the Western world. For one thing, it's creating a totally new type of national leader, a man who is much more of a tribal chieftain than a politician. Castro is a good example of the new tribal chieftain who rules his country by a mass-participational TV dialog and feedback; he governs his country on camera, by giving the Cuban people the experience of being directly and intimately involved in the process of collective decision making. Castro's adroit blend of political education, propaganda and avuncular guidance is the pattern for tribal chieftains in other countries. The new political showman has to literally as well as figuratively put on his audience as he would a suit of clothes and become a corporate tribal image—like Mussolini, Hitler and F.D.R. in the days of radio, and Jack Kennedy in the television era. All these men were tribal emperors on a scale theretofore unknown in the world, because they all mastered their media. . . . The overhauling of our traditional political system is only one manifestation of the retribalizing process wrought by the electric media, which is turning the planet into a global village.

Interviewer: Would you describe this retribalizing process in more detail?

McLuhan: The electronically induced technological extensions of our central nervous systems, which I spoke of earlier, are immersing us in a world-pool of information movement and are thus enabling man to incorporate within himself the whole of mankind. The aloof and dissociated role of the literate man of the Western world is succumbing to the new, intense depth participation engendered by the electronic media and bringing us back in touch with ourselves as well as with one another. But the instant nature of electric-information movement is decentralizing—rather than enlarging—the family of man into a new state

of multitudinous tribal existences. Particularly in countries where literate values are deeply institutionalized, this is a highly traumatic process, since the clash of the old segmented visual culture and the new integral electronic culture creates a crisis of identity, a vacuum of the self, which generates tremendous violence—violence that is simply an identity quest, private or corporate, social or commercial.

Interviewer: Do you relate this identity crisis to the current social unrest and violence in the United States?

McLuhan: Yes, and to the booming business psychiatrists are doing. All our alienation and atomization are reflected in the crumbling of such time-honored social values as the right of privacy and the sanctity of the individual; as they yield to the intensities of the new technology's electric circus, it seems to the average citizen that the sky is falling in. As man is tribally metamorphosed by the electric media, we all become Chicken Littles, scurrying around frantically in search of our former identities, and in the process unleash tremendous violence. As the preliterate confronts the literate in the postliterate arena, as new information patterns inundate and uproot the old, mental breakdowns of varying degrees—including the collective nervous breakdowns of whole societies unable to resolve their crises of identity—will become very common.

It is not an easy period in which to live, especially for the television-conditioned young who, unlike their literate elders, cannot take refuge in the zombie trance of Narcissus narcosis that numbs the state of psychic shock induced by the impact of the new media. From Tokyo to Paris to Columbia, youth mindlessly acts out its identity quest in the theater of the streets, searching not for goals but for roles, striving for an identity that eludes them.

Interviewer: Why do you think they aren't finding it within the educational system?

McLuhan: Because education, which should be helping youth to understand and adapt to their revolutionary new environments, is instead being used merely as an instrument of cultural aggression, imposing upon retribalized youth the obsolescent visual values of the dying literate age. Our entire educational system is reactionary, oriented to past values and past technologies, and will likely continue so until the old generation relinquishes power. The generation gap is actually a chasm, separating not two age groups but two vastly divergent cultures. I can understand the ferment in our schools, because our educational system is totally rearview mirror. It's a dying and outdated system founded on literate values and fragmented and classified data totally unsuited to the needs of the first television generation.

Interviewer: How do you think the educational system can be adapted to accommodate the needs of this television generation?

McLuhan: Well, before we can start doing things the right way, we've got to recognize that we've been doing them the wrong way—which most pedagogs and administrators and even most parents still refuse to accept. Today's child is growing up absurd because he is suspended between two worlds and two value systems, neither of which inclines him to maturity because he belongs wholly to neither but exists in a hybrid limbo of constantly conflicting values. The challenge of the new era is simply the total creative process of growing up—and mere teaching and repetition of facts are as irrelevant to this process as a dowsing to a nuclear power plant. To expect a "turned on" child of the electric age to respond to the old education modes is rather like expecting an eagle to swim. It's simply not within his environment, and therefore incomprehensible.

The TV child finds it difficult if not impossible to adjust to the fragmented, visual goals of our education after having had all his senses involved by the electric media; he craves in-depth involvement, not linear detachment and uniform sequential patterns. But suddenly and without preparation, he is snatched from the cool, inclusive womb of television and exposed—within a vast bureaucratic structure of courses and credits—to the hot medium of print. His natural instinct, conditioned by the electric media, is to bring all his senses to bear on the book he's instructed to read, and print resolutely rejects that approach, demanding an isolated visual attitude to learning rather than the Gestalt approach of the unified sensorium. The reading postures of children in elementary school are a pathetic testimonial to the effects of television; children of the TV generation separate book from eye by an average distance of four and a half inches, attempting psychomimetically to bring to the printed page the all-inclusive sensory experience of TV. They are becoming Cyclops, desperately seeking to wallow in the book as they do in the TV screen.

Interviewer: Might it be possible for the "TV child" to make the adjustment to his educational environment by synthesizing traditional literate-visual forms with the insights of his own electric culture—or must the medium of print be totally unassimilable for him?

McLuhan: Such a synthesis is entirely possible, and could create a creative blend of the two cultures—if the educational establishment was aware that there is an electric culture. In the absence of such elementary awareness, I'm afraid that the television child has no future in our schools. You must remember that the TV child has been relentlessly exposed to all the “adult” news of the modern world—war, racial discrimination, rioting, crime, inflation, sexual revolution. The war in Vietnam has written its bloody message on his skin; he has witnessed the assassinations and funerals of the nation's leaders; he's been orbited through the TV screen into the astronaut's dance in space, been inundated by information transmitted via radio, telephone, films, recordings and other people. His parents plopped him down in front of a TV set at the age of two to tranquilize him, and by the time he enters kindergarten, he's clocked as much as 4000 hours of television. As an IBM executive told me, “My children had lived several lifetimes compared to their grandparents when they began grade one.”

Interviewer: If you had children young enough to belong to the TV generation, how would you educate them?

McLuhan: Certainly not in our current schools, which are intellectual penal institutions. In today's world, to paraphrase Jefferson, the least education is the best education, since very few young minds can survive the intellectual tortures of our educational system. The mosaic image of the TV screen generates a depth-involving *nowness* and simultaneity in the lives of children that makes them scorn the distant visualized goals of traditional education as unreal, irrelevant and puerile. Another basic problem is that in our schools there is simply too much to learn by the traditional analytic methods; this is an age of information overload. The only way to make the schools other than prisons without bars is to start fresh with new techniques and values. . . .

Interviewer: [You say that personal freedom will still exist in the coming, retribalized world. What] about the political system most closely associated with individual freedom: democracy? Will it, too, survive the transition to your global village?

McLuhan: No, it will not. The day of political democracy as we know it today is finished. Let me stress again that individual freedom itself will not be submerged in the new tribal society, but it will certainly assume different and more complex dimensions. The ballot box, for example, is the product of literate Western culture—a hot box in a cool world—and thus obsolescent. The tribal will is consensually expressed through the simultaneous interplay of all members of a community that is deeply interrelated and involved, and would thus consider the casting of a “private” ballot in a shrouded polling booth a ludicrous anachronism. The TV networks' computers, by “projecting” a victor in a Presidential race while the polls are still open, have already rendered the traditional electoral process obsolescent.

In our software world of instant electric communications movement, politics is shifting from the old patterns of political representation by electoral delegation to a new form of spontaneous and instantaneous communal involvement in all areas of decision making. In a tribal all-at-once culture, the idea of the “public” as a differentiated agglomerate of fragmented individuals, all dissimilar but all capable of acting in basically the same way, like interchangeable mechanical cogs in a production line, is supplanted by a mass society in which personal diversity is encouraged while at the same time everybody reacts and interacts simultaneously to every stimulus. The election as we know it today will be meaningless in such a society.

Interviewer: How will the popular will be registered in the new tribal society if elections are *passé*?

McLuhan: The electric media open up totally new means of registering popular opinion. The old concept of the plebiscite, for example, may take on new relevance; TV could conduct daily plebiscites by presenting facts to 200,000,000 people and providing a computerized feedback of the popular will. But voting, in the traditional sense, is through as we leave the age of political parties, political issues and political goals, and enter an age where the collective tribal image and the iconic image of the tribal chieftain is the overriding political reality. But that's only one of countless new realities we'll be confronted with in the tribal village. We must understand that a totally new society is coming into being, one that rejects all our old values, conditioned responses, attitudes and institutions. If you have difficulty envisioning something as trivial as the imminent end of elections, you'll be totally unprepared to cope with the prospect of the forthcoming demise of spoken language and its replacement by a global consciousness.

Interviewer: You're right.

McLuhan: Let me help you. Tribal man is tightly sealed in an integral collective awareness that transcends conventional boundaries of time and space. As such, the new society will be one mythic integration, a resonating world akin to the old tribal echo chamber where magic will live again: a world of ESP. The current interest of youth in astrology, clairvoyance and the occult is no coincidence. Electric technology, you see, does not require words any more than a digital computer requires numbers. Electricity makes possible—and not in the distant future, either—an amplification of human consciousness on a world scale, without any verbalization at all.

Interviewer: Are you talking about global telepathy?

McLuhan: Precisely. Already, computers offer the potential of instantaneous translation of any code or language into any other code or language. If a data feedback is possible through the computer, why not a feed-forward of thought whereby a world consciousness links into a world computer? Via the computer, we could logically proceed from translating languages to bypassing them entirely in favor of an integral cosmic unconsciousness somewhat similar to the collective unconscious envisioned by Bergson. The computer thus holds out the promise of a technologically engendered state of universal understanding and unity, a state of absorption in the logos that could knit mankind into one family and create a perpetuity of collective harmony and peace. This is the real use of the computer, not to expedite marketing or solve technical problems but to speed the process of discovery and orchestrate terrestrial—and eventually galactic—environments and energies. Psychic communal integration, made possible at last by the electronic media, could create the universality of consciousness foreseen by Dante when he predicted that men would continue as no more than broken fragments until they were unified into an inclusive consciousness. In a Christian sense, this is merely a new interpretation of the mystical body of Christ; and Christ, after all, is the ultimate extension of man.

Interviewer: Isn't this projection of an electronically induced world consciousness more mystical than technological?

McLuhan: Yes—as mystical as the most advanced theories of modern nuclear physics. Mysticism is just tomorrow's science dreamed today.

Interviewer: You said a few minutes ago that all of contemporary man's traditional values, attitudes and institutions are going to be destroyed and replaced in and by the new electric age. That's a pretty sweeping generalization. Apart from the complex psychosocial metamorphoses you've mentioned, would you explain in more detail some of the specific changes you foresee?

McLuhan: The transformations are taking place everywhere around us. As the old value systems crumble, so do all the institutional clothing and garbage they fashioned. The cities, corporate extensions of our physical organs, are withering and being translated along with all other such extensions into information systems, as television and the jet—by compressing time and space—make all the world one village and destroy the old city-country dichotomy. New York, Chicago, Los Angeles—all will disappear like the dinosaur. The automobile, too, will soon be as obsolete as the cities it is currently strangling, replaced by new antigravitational technology. The marketing systems and the stock market as we know them today will soon be dead as the dodo, and automation will end the traditional concept of the job, replacing it with a role, and giving men the breath of leisure. The electric media will create a world of dropouts from the old fragmented society, with its neatly compartmentalized analytic functions, and cause people to drop in to the new integrated global-village community.

All these convulsive changes, as I've already noted, carry with them attendant pain, violence and war—the normal stigmata of the identity quest—but the new society is springing so quickly from the ashes of the old that I believe it will be possible to avoid the transitional anarchy many predict. Automation and cybernation can play an essential role in smoothing the transition to the new society.

Interviewer: How?

McLuhan: The computer can be used to direct a network of global thermostats to pattern life in ways that will optimize human awareness. Already, it's technologically feasible to employ the computer to program societies in beneficial ways.

Interviewer: How do you program an entire society—beneficially or otherwise?

McLuhan: There's nothing at all difficult about putting computers in the position where they will be able to conduct carefully orchestrated programing of the sensory life of whole populations. I know it sounds rather science-fictional, but if you understood cybernetics you'd realize we could do it today. The computer could program the media to determine the given messages a people should hear in terms of their overall needs, creating a total media experience absorbed and patterned by all the senses. We could program five hours less of TV in Italy to promote the reading of newspapers during an election, or lay on an additional 25 hours of TV in Venezuela to cool down the tribal temperature raised by radio the preceding month. By such orchestrated interplay of all media, whole cultures could now be programed in order to improve and stabilize their emotional climate, just as we are beginning to learn how to maintain equilibrium among the world's competing economies.

Interviewer: How does such environmental programing, however enlightened in intent, differ from Pavlovian brainwashing?

McLuhan: Your question reflects the usual panic of people confronted with unexplored technologies. I'm not saying such panic isn't justified, or that such environmental programing couldn't be brainwashing, or far worse—merely that such reactions are useless and distracting. Though I think the programing of societies could actually be conducted quite constructively and humanistically, I don't want to be in the position of a Hiroshima physicist extolling the potential of nuclear energy in the first days of August 1945. But an understanding of media's effects constitutes a civil defense against media fallout.

The alarm of so many people, however, at the prospect of corporate programing's creation of a complete service environment on this planet is rather like fearing that a municipal lighting system will deprive the individual of the right to adjust each light to his own favorite level of intensity. Computer technology can—and doubtless will—program entire environments to fulfill the social needs and sensory preferences of communities and nations. The content of that programing, however, depends on the nature of future societies—but that is in our own hands.

Interviewer: Is it really in our hands—or, by seeming to advocate the use of computers to manipulate the future of entire cultures, aren't you actually encouraging man to abdicate control over his destiny?

McLuhan: First of all—and I'm sorry to have to repeat this disclaimer—I'm not advocating anything; I'm merely probing and predicting trends. Even if I opposed them or thought them disastrous, I couldn't stop them, so why waste my time lamenting? As Carlyle said of author Margaret Fuller after she remarked, "I accept the Universe": "She'd better." I see no possibility of a worldwide Luddite rebellion that will smash all machinery to bits, so we might as well sit back and see what is happening and what will happen to us in a cybernetic world. Resenting a new technology will not halt its progress.

The point to remember here is that whenever we use or perceive any technological extension of ourselves, we necessarily embrace it. Whenever we watch a TV screen or read a book, we are absorbing these extensions of ourselves into our individual system and experiencing an automatic "closure" or displacement of perception; we can't escape this perpetual embrace of our daily technology unless we escape the technology itself and flee to a hermit's cave. By consistently embracing all these technologies, we inevitably relate ourselves to them as servomechanisms. Thus, in order to make use of them at all, we must serve them as we do gods. The Eskimo is a servomechanism of his kayak, the cowboy of his horse, the businessman of his clock, the cyberneticist—and soon the entire world—of his computer. In other words, to the spoils belongs the victor. This continuous modification of man by his own technology stimulates him to find continuous means of modifying it; man thus becomes the sex organs of the machine world just as the bee is of the plant world, permitting it to reproduce and constantly evolve to higher forms. The machine world reciprocates man's devotion by rewarding him with goods and services and bounty. Man's relationship with his machinery is thus inherently symbiotic. This has always been the case; it's only in the electric age that man has an opportunity to recognize this marriage to his own technology. Electric technology is a qualitative extension of this age-old man-machine relationship; 20th Century man's relationship to the computer is not by nature very different from prehistoric man's relationship to his boat or to his wheel—with the important difference that all previous technologies or extensions of man were partial and fragmentary, whereas the electric is total and inclusive. Now man is beginning to wear his brain outside his skull and his nerves outside his skin; new technology breeds new man. A recent cartoon portrayed a little boy telling his nonplused mother: "I'm going to be a computer when I grow up." Humor is often prophecy.

Interviewer: If man can't prevent this transformation of himself by technology—or into technology—how can he control and direct the process of change?

McLuhan: The first and most vital step of all, as I said at the outset, is simply to understand media and its revolutionary effects on all psychic and social values and institutions. Understanding is half the battle. The central purpose of all my work is to convey this message, that by understanding media as they extend man, we gain a measure of control over them. And this is a vital task, because the immediate interface between audile-tactile and visual perception is taking place everywhere around us. No civilian can escape this environmental blitzkrieg, for there is, quite literally, no place to hide. But if we diagnose what is happening to us, we can reduce the ferocity of the winds of change and bring the best elements of the old visual culture, during this transitional period, into peaceful coexistence with the new retribalized society.

If we persist, however, in our conventional rearview-mirror approach to these cataclysmic developments, all of Western culture will be destroyed and swept into the dustbin of history. If literate Western man were really interested in preserving the most creative aspects of his civilization, he would not cower in his ivory tower bemoaning change but would plunge himself into the vortex of electric technology and, by understanding it, dictate his new environment—turn ivory tower into control tower. But I can understand his hostile attitude, because I once shared his visual bias.

Interviewer: What changed your mind?

McLuhan: Experience. For many years, until I wrote my first book, *The Mechanical Bride*, I adopted an extremely moralistic approach to all environmental technology. I loathed machinery, I abominated cities, I equated the Industrial Revolution with original sin and mass media with the Fall. In short, I rejected almost every element of modern life in favor of a Rousseauvian utopianism. But gradually I perceived how sterile and useless this attitude was, and I began to realize that the greatest artists of the 20th Century—Yeats, Pound, Joyce, Eliot—had discovered a totally different approach, based on the identity of the processes of cognition and creation. I realized that artistic creation is the playback of ordinary experience—from trash to treasures. I ceased being a moralist and became a student.

As someone committed to literature and the traditions of literacy, I began to study the new environment that imperiled literary values, and I soon realized that they could not be dismissed by moral outrage or pious indignation. Study showed that a totally new approach was required, both to save what deserved saving in our Western heritage and to help man adopt a new survival strategy. I adapted some of this new approach in *The Mechanical Bride* by attempting to immerse myself in the advertising media in order to apprehend its impact on man, but even there some of my old literate “point of view” bias crept in. The book, in any case, appeared just as television was making all its major points irrelevant.

I soon realized that recognizing the symptoms of change was not enough; one must understand the cause of change, for without comprehending causes, the social and psychic effects of new technology cannot be counteracted or modified. But I recognized also that one individual cannot accomplish these self-protective modifications; they must be the collective effort of society, because they affect all of society; the individual is helpless against the pervasiveness of environmental change: the new garbage—or mess-age—induced by new technologies. Only the social organism, united and recognizing the challenge, can move to meet it.

Unfortunately, no society in history has ever known enough about the forces that shape and transform it to take action to control and direct new technologies as they extend and transform man. But today, change proceeds so instantaneously through the new media that it may be possible to institute a global education program that will enable us to seize the reins of our destiny—but to do this we must first recognize the kind of therapy that's needed for the effects of the new media. In such an effort, indignation against those who perceive the nature of those effects is no substitute for awareness and insight.

Interviewer: Are you referring to the critical attacks to which you've been subjected for some of your theories and predictions?

McLuhan: I am. But I don't want to sound uncharitable about my critics. Indeed, I appreciate their attention. After all, a man's detractors work for him tirelessly and for free. It's as good as being banned in Boston. But as I've said, I can understand their hostile attitude toward environmental change, having once shared it. Theirs is the customary human reaction when confronted with innovation: to flounder about attempting to adapt old responses to new situations or to simply condemn or ignore the harbingers of change—a practice refined by the Chinese emperors, who used to execute messengers bringing bad news.

The new technological environments generate the most pain among those least prepared to alter their old value structures. The literati find the new electronic environment far more threatening than do those less committed to literacy as a way of life. When an individual or social group feels that its whole identity is jeopardized by social or psychic change, its natural reaction is to lash out in defensive fury. But for all their lamentations, the revolution has already taken place.

Interviewer: You've explained why you avoid approving or disapproving of this revolution in your work, but you must have a private opinion. What is it?

McLuhan: I don't like to tell people what I think is good or bad about the social and psychic changes caused by new media, but if you insist on pinning me down about my own subjective reactions as I observe the reprimativization of our culture, I would have to say that I view such upheavals with total personal dislike and dissatisfaction. I do see the prospect of a rich and creative retribalized society—free of the fragmentation and alienation of the mechanical age—emerging from this traumatic period of culture clash; but I have nothing but distaste for the process of change. As a man molded within the literate Western tradition, I do not personally cheer the dissolution of that tradition through the electric involvement of all the senses: I don't enjoy the destruction of neighborhoods by high-rises or revel in the pain of identity quest. No one could be less enthusiastic about these radical changes than myself. I am not, by temperament or conviction, a revolutionary; I would prefer a stable, changeless environment of modest services and human scale. TV and all the electric media are unraveling the entire fabric of our society, and as a man who is forced by circumstances to live within that society, I do not take delight in its disintegration.

You see, I am not a crusader; I imagine I would be most happy living in a secure preliterate environment; I would never attempt to change my world, for better or worse. Thus I derive no joy from observing the traumatic effects of media on man, although I do obtain satisfaction from grasping their modes of operation. Such comprehension is inherently cool, since it is simultaneously involvement and detachment. This posture is essential in studying media. One must begin by becoming extraenvironmental, putting oneself beyond the battle in order to study and understand the configuration of forces. It's vital to adopt a posture of arrogant superiority; instead of scurrying into a corner and wailing about what media are doing to us, one should charge straight ahead and kick them in the electrodes. They respond beautifully to such resolute treatment and soon become servants rather than masters. But without this detached involvement, I could never objectively observe media; it would be like an octopus grappling with the Empire State Building. So I employ the greatest boon of literate culture: the power of man to act without reaction—the sort of specialization by dissociation that has been the driving motive force behind Western civilization.

The Western world is being revolutionized by the electric media as rapidly as the East is being Westernized, and although the society that eventually emerges may be superior to our own, the process of change is agonizing. I must move through this pain-wracked transitional era as a scientist would move through a world of disease; once a surgeon becomes personally involved and disturbed about the condition of his patient, he loses the power to help that patient. Clinical detachment is not some kind of haughty pose I affect—nor does it reflect any lack of compassion on my part; it's simply a survival strategy. The world we are living in is not one I would have created on my own drawing board, but it's the one in which I must live, and in which the students I teach must live. If nothing else, I owe it to them to avoid the luxury of moral indignation or the troglodytic security of the ivory tower and to get down into the junk yard of environmental change and steam-shovel my way through to a comprehension of its contents and its lines of force—in order to understand how and why it is metamorphosing man.

Interviewer: Despite your personal distaste for the upheavals induced by the new electric technology, you seem to feel that if we understand and influence its effects on us, a less alienated and fragmented society may emerge from it. Is it thus accurate to say that you are essentially optimistic about the future?

McLuhan: There are grounds for both optimism and pessimism. The extensions of man's consciousness induced by the electric media could conceivably usher in the millennium, but it also holds the potential for realizing the Anti-Christ—Yeats' rough beast, its hour come round at last, slouching toward Bethlehem to be born. Cataclysmic environmental changes such as these are, in and of themselves, morally neutral; it is how we perceive them and react to them that will determine their ultimate psychic and social consequences. If we refuse to see them at all, we will become their servants. It's inevitable that the world-pool of electronic information movement will toss us all about like corks on a stormy sea, but if we keep our cool during the

descent into the maelstrom, studying the process as it happens to us and what we can do about it, we can come through.

Personally, I have a great faith in the resiliency and adaptability of man, and I tend to look to our tomorrows with a surge of excitement and hope. I feel that we're standing on the threshold of a liberating and exhilarating world in which the human tribe can become truly one family and man's consciousness can be freed from the shackles of mechanical culture and enabled to roam the cosmos. I have a deep and abiding belief in man's potential to grow and learn, to plumb the depths of his own being and to learn the secret songs that orchestrate the universe. We live in a transitional era of profound pain and tragic identity quest, but the agony of our age is the labor pain of rebirth.

I expect to see the coming decades transform the planet into an art form; the new man, linked in a cosmic harmony that transcends time and space, will sensuously caress and mold and pattern every facet of the terrestrial artifact as if it were a work of art, and man himself will become an organic art form. There is a long road ahead, and the stars are only way stations, but we have begun the journey. To be born in this age is a precious gift, and I regret the prospect of my own death only because I will leave so many pages of man's destiny—if you will excuse the Gutenbergian image—tantalizingly unread. But perhaps, as I've tried to demonstrate in my examination of the postliterate culture, the story begins only when the book closes.

Technology and Social Justice

It is easy to find historical examples illustrating the thesis that technology may have something to contribute to social justice. In the fourteenth century the new technology of printing changed the face of Europe, bringing books and education out of the monasteries and spreading them far and wide among the people. Printing gave power to the Bible and led directly to the Protestant Reformation in Northern Europe. One may question whether Luther's Germany and Shakespeare's England enjoyed social justice, but they were certainly closer to it than the medieval Germany and England out of which they grew. Luther and Shakespeare brought at least the idea of justice—if not the reality—to ordinary citizens outside the nobility and the priesthood. The Protestant ethic, which took root in Germany, England, Holland, and Scandinavia with the help of printed books, carried with it a perpetual striving for social justice, even if the Utopian visions were seldom achieved.

More recent technologies that contributed in a practical way to social justice were those of public health—clean water supplies, sewage treatment, vaccination, antibiotics. These technologies could only be effective in protecting the rich from contagion and sickness if they were also available to the poor. Even if the rich and powerful receive preferential treatment, as they usually do, the benefits of public health technology are felt to some extent by everybody. In countries where public health technologies are enforced by law, there is no large gap in life expectancy between rich and poor.

The technology of synthetic materials has also helped to erase differences between rich and poor. Throughout history, until the nineteenth century, only the rich could afford to dress in brilliant colors, furs, and silk. Fine clothes were a badge of privilege and wealth. In the nineteenth century the chemical industry produced artificial dyestuffs. The twentieth century added artificial fur and silk and many other synthetic fabrics

TECHNOLOGY AND SOCIAL JUSTICE

Freeman Dyson

The fourth Louis Nizer Lecture on Public Policy

November 5, 1997

sponsored by



OCCIDENTAL PETROLEUM CORPORATION

and supported by

PHILLIPS NIZER BENJAMIN KRIM & BALLON LLP



CARNEGIE COUNCIL ON ETHICS
AND INTERNATIONAL AFFAIRS

but it was less painful than a civil war. It was a big step on the road to social justice.

I remember with great fondness the nursemaid, Ethel, who cared for me as a young child. She had left school, as girls of the servant class did in those days, at the age of fourteen. When my sister and I were safely in bed in the night nursery, we sometimes heard the "putt, putt, putt" of a motorbike approaching the house, stopping, and then driving away into the night. That was Ethel's young man taking her out for the evening. The motorbike was the first harbinger of the approaching social revolution. The motorbike was the technology of upward mobility. After Ethel left us and married the young man, she had three daughters of her own, and all of them went to college. One of her grandsons is now a university professor.

Those are enough examples to show that technology can be helpful in the struggle for social justice. But in each case, as Edward Tenner tells us in his book *Why Things Bite Back*, a step forward in technology tends to bring with it an unexpected step backward. A step forward for some people frequently brings with it a step backward for others. And it often happens that when an old privileged class of people is dispossessed and the blessings of wealth and power are spread more equally, the burdens of equalization fall disproportionately upon women. When the revolutions accompanying the technology of printing destroyed the wealth and power of the monasteries over much of Europe, both male and female orders were dispossessed, but the nuns lost more than the monks. Nuns in the old convents were in many ways more free than wives in the new Protestant communities. The old monastic society provided a refuge where women of outstanding ability—for example, Hildegard of Bingen—had access to higher education. Sheltered and supported by the monastic orders, women could follow their vocations as scholars and artists. When the monasteries were dissolved, nuns had to find shelter in other people's homes, either as wives or as servants. The new secular society replaced the monasteries with colleges and universities. In the universities

cheap enough for working-class women to afford. No longer can one tell a woman's social class by her clothes. It is a measure of social justice in modern societies that the children of the rich now dress down, imitating the style of the majority both in clothes and in behavior.

Household appliances are another technology with a tendency towards social justice. When I was a child in England in the 1920s, my mother employed four full-time servants: a cook, a housemaid, a nursemaid, and a gardener. We didn't consider ourselves rich. My father was a schoolteacher. We were an average middle-class family. In those days an average middle-class family needed four servants to do the hard manual work of cooking, cleaning, child care, and gardening. To do all this work a whole class of people existed who spent their lives as domestic servants. The professional and intellectual classes to which we belonged were riding on the backs of the servant class. Because of the servants, my mother had leisure to organize socially useful projects, such as a club for teenage girls and a birth control clinic. The birth control clinic was undoubtedly a godsend to the women who came to it for instruction in the art of not having unwanted babies. But it did not in any way narrow the gulf between her and them. She always spoke of her birth control clientele like a mistress speaking of servants.

My mother was a kind mistress and treated the servants well, according to the standards of the time, but the servants knew their place. They knew that if they disobeyed orders or answered back, they would be out on the street. Now, like the antebellum South, the servant class in England is gone with the wind, and the wind that blew it away was not the ravaging invasion of Sherman's army, but the peaceful invasion of an army of electric stoves, gas heaters, vacuum cleaners, refrigerators, washing machines, drying machines, garbage disposals, freezers, microwave ovens, juicers, choppers, and disposable diapers. The technology of household appliances made servants unnecessary, and, at the same time, the children of the servant class began to go to college and make the transition to the middle class. The transition was not painless,

mothers of the middle class, this was a big step backward. Appliances do not cook the dinner, clean the house, do the shopping, and mind the baby. The middle-class women of the 1950s were far less liberated than their mothers. The liberation that my mother's generation achieved had to be fought for all over again. Even now, in the 1990s, women are only partially liberated. To achieve partial liberation, they have replaced the old domestic servants with day care centers, cleaning ladies, and *au pair* girls imported from overseas. Electrical appliances help, but they only do a small part of the job.

The Institute for Advanced Study, where I have spent my working life, is a peculiar institution with a small permanent faculty. The faculty is supposed to be representative of the most distinguished men and women in academic life. Unfortunately, we have always found it difficult to appoint women to the faculty. The original faculty, appointed in the 1930s, contained one woman, the archaeologist Hetty Goldman. I remember her vividly. She was a formidable lady, small in stature and large in spirit, who led excavations of ancient sites in Turkey, ruling over small armies of Turkish laborers with an iron hand. Her colleagues used to say she was the equal of any two male archaeologists. There was never the slightest doubt that she had the "right stuff" to be an Institute professor. She was a natural leader in her own eyes and in ours. She belonged to my mother's generation of liberated women. She grew up, like my mother, in a society of women with servants. When she retired in 1947, she was not replaced. For almost forty years the Institute faculty was entirely male. In 1985, the sociologist Joan Scott became the second woman to join the faculty. And in 1997 the historian Patricia Crone became the third.

The history of our faculty encapsulates in a nutshell the history of women's liberation: a glorious beginning in the 1920s; a great backsliding in the 1950s; a gradual recovery in the 1980s. It is not altogether fanciful to blame the technology of household appliances for the backsliding. The advent of electrical appliances liberated the servants and shackled their mistresses.

men scholars could find shelter and security, but there was no place for women.

The technology of household appliances, likewise, brought a step backward to the stratum of society to which my mother belonged, the women of the middle class. My mother would be considered by the standards of today a thoroughly liberated woman. Trained as a lawyer, she helped to write the Act of Parliament that opened the professions in England to women. With the help of her servants, she could take care of her husband and children without being confined to the home. She was free to pursue her interests outside the home—her girls' club and birth control clinic. But she was by no means the most liberated of the women in our family. I had a collection of aunts who were in various ways more liberated than my mother. All of them had husbands and most of them had children, but this did not stop them from being liberated. All of them were more adventurous than their husbands. My Aunt Margaret was trained as a nurse and rose to become a matron, which meant that she was the managing administrator of a large hospital. My Aunt Ruth was a figure skater of international repute who kept an Olympic silver medal among her trophies. My Aunt Dulcibella was the first woman in England to receive an airplane pilot's license. She and her husband had an airplane which they used for traveling around in Africa. They loved Africa, and their lifestyle would have fit in very well with the group of adventurers that Michael Ondaatje describes in his novel *The English Patient*. My Aunt Dulcibella was also a professional actress, and if she had only been eighty years younger, she might have had a starring role in *The English Patient* movie. We did not consider these aunts of ours to be unusual. It was normal at that time for middle-class women to do something spectacular. My mother, with her birth control clinic, was the quiet one, the least daring of the four.

Now, consider what happened to the next generation of middle-class women in England and the United States. Thirty years later, in the 1950s, the servants were gone and the electrical appliances were taking their place. For wives and

High-Tech Medicine and Computers

I have discussed four technologies that led to large expansions of social justice. Although each of them had compensating negative effects, especially on women, the overall effects of all of them were predominantly positive. It will be just as easy to find examples of technologies that had predominantly negative effects. One could mention the technologies of gas chambers and nuclear weapons, useful for the convenient extermination of people to whom we do not wish to extend the benefits of social justice. But the more troubling examples are two of the technologies that are making the most rapid progress today: high-tech medicine and high-tech communication.

All of us who live in the United States are familiar with the ugly face that high-tech medicine presents to the patient: the initial telephone call answered by a machine rather than a human voice; the filling out of forms in the office; the repetitive answering of questions; the battery of routine chemical and physical tests carried out by technicians wearing rubber gloves; and finally, the abbreviated contact with the physician. It is all very different from the old-fashioned practice of medicine, when doctors were personal friends and advisers to patients and sometimes even made house calls. The face of high-tech medicine is ugly even when the patient is rich, and uglier still when the patient is poor. The ugliness results from many factors working together. First, the prevalence of malpractice litigation, which destroys trust, compelling doctors to conform to rigid rules and surrounding them with layers of bureaucratic documentation. Second, the high cost of the equipment that high-tech medicine demands, forcing medical centers to adopt elaborate cost-accounting systems. Third, the size of the staff needed to operate a high-tech center, with many doctors qualified in narrow specialties so that the patient rarely gets to see the same doctor twice. Fourth, the overwhelming cost of hospitalization, allowing patients a bare minimum of days for rest and recuperation after major illness or surgery. These factors, together, led to the situation

that confronts the patient today. What the patient needs most, but finds least, is personal attention.

Since personal attention has become the scarcest resource in high-tech medicine, it is inevitable that it should be distributed unequally. The majority of advanced countries have national health services that attempt, with varying degrees of success, to distribute medical attention fairly. In countries with national health services, medical attention is theoretically available to everybody. This is what the ethic of social justice demands. But the escalating cost of medical attention makes social justice more and more difficult to achieve. One way or another, as personal attention becomes scarcer, people of status tend to receive more of it and people without status to receive less. The national health services in countries where they exist make valiant efforts to preserve the ideal of social justice, but the march of medical technology and the concomitant increase of costs constantly erode the ideal. In the United States, which never had a national health service and does not pretend to distribute medical resources equally, the prospects for social justice are far worse. In the United States a medical system based on the ethic of the free market inevitably favors the rich over the poor, and the inequalities grow sharper as the costs increase.

I have seen in my own family a small example of the dilemma that the growth of high-tech medicine presents to physicians. One of my daughters is a cardiologist. For many years she worked in state-supported hospitals taking care of patients as they flowed through the system, working brutally long hours and still having little time for personal contact with her patients. Her patients in the public hospitals were predominantly poor and uninsured. Many of them had AIDS or gunshot wounds in addition to cardiac problems. The public health system, such as it was, was designed to get these patients out of the hospital and back on the streets as fast as possible. Last year my daughter was offered a job in a private cardiology practice with far shorter hours, better pay and working conditions, and an expectation of long, continued care of her patients. She accepted the offer without much

and pampering the top-end customer, the top-end customer usually wins.

The problem of unequal access to computers is only a small part of the problem of inequality in our society. Until the society is willing to attack the larger problems of inequality in housing, education, and health care, attempts to provide equal access to computers cannot be totally successful. Nevertheless, in attacking the general problems of inequality, computer access might be a good place to start. One of the virtues of the new technology of the Internet is that it has an inherent tendency to become global. The Internet easily infiltrates through barriers of language, custom, and culture. No technical barrier stops it from becoming universally accessible. To provide equality of access to the Internet is technically easier than providing equality of access to housing and health care. Universal access to the Internet would not solve all our social problems, but it would be a big step in the right direction. The Internet could then become an important tool for alleviating other kinds of inequality.

The Protestant Ethic and the Spirit of Capitalism

Up to now I have been talking as if technology came first and ethics second. I have been describing historical events in which technological changes occurred first and then increases or decreases of social justice occurred as a consequence. I depicted technological change as the cause of ethical improvement or deterioration. This view of history is opposed to the view propounded by Max Weber in his seminal book *The Protestant Ethic and the Spirit of Capitalism*. Weber argued that the Protestant ethic came first and the rise of capitalism and the technologies associated with it came second. Weber's view has become the prevailing view of modern historians. Weber said that ethics drove technology. I say that technology drives ethics.

I am not trying to prove Weber wrong. His historical vision remains profoundly true. It is true that the religious revolutions of the sixteenth century engendered an ethic of

hesitation. She is much happier in her new job. Now, for the first time, she knows her patients as individuals and can tailor their treatments to their individual histories and personalities. She feels that she is a better doctor, and her new job gave her the flexibility to take time off to have her first baby last July. From almost every point of view, her jump into private practice was a wise move. Her only problem was a small twinge of conscience for having abandoned the poor to take care of the rich. In the private practice her patients are not all rich, but they are all paying for the personal attention that she is now able to give them. She was forced to make a choice between social justice and professional satisfaction, and social justice lost. I don't blame her. But in a socially just society, physicians would not be forced to make such choices.

Similar dilemmas, not so stark as the dilemmas of medical practice but equally important, exist in the world of high-tech computing and communications. Here, too, there is a clash between the economic forces driving the technology and the needs of poor people. Access to personal computers and the Internet is like medical insurance: almost everybody needs it, but most poor people don't have it. The people who are wired, the people who browse the World Wide Web and conduct their daily lives and businesses on the Net, have tremendous economic and social advantages. Increasingly, jobs and business opportunities are offered through the Internet. Access to the Internet means access to well-paying jobs. People who are not wired in are in danger of becoming the new servant class. The gulf between the wired and the unwired is wide, and growing wider.

The computer and software industries are driven by two contradictory impulses. On the one hand, they sincerely wish to broaden their market by making computers accessible to everybody. On the other hand, they are forced by competitive pressures to upgrade their products constantly, increasing their power and speed and adding new features and new complications. The top end of the market drives the development of new products, and the new products remain out of the reach of the poor. In the tug of war between broadening the market

personal responsibility and restless inquiry, an ethic that encouraged the growth of capitalistic enterprise and technological innovation. It was no accident that Isaac Newton, the preeminent architect of modern science, was also a Protestant theologian. He took his theology as seriously as his science. It was no accident that King Henry VIII, the man who brought the Protestant revolution to England, also endowed the college where Newton lived and taught. Henry and Isaac were kindred spirits—both were rebels against authority, enemies of the Pope, tyrants, supreme egoists, suspicious to the point of paranoia, believers in the Protestant ethic, and in love with technology. Henry loved to build ships and Isaac loved to build telescopes. It is true that ethics can drive technology. I am only saying that this is not the whole truth, that technology can also drive ethics, that the chain of causation works in both directions. The technology of printing helped to cause the rise of the Protestant ethic just as much as the Protestant ethic helped to cause the rise of navigation and astronomy.

I am not the first to take issue with Weber on this question. The historian Richard Tawney also studied the interrelationships of religion and capitalism and came to conclusions similar to mine. He held Weber in high esteem and contributed a foreword to the English translation of *The Protestant Ethic and the Spirit of Capitalism*. Here are the concluding sentences of Tawney's foreword: "It is instructive to trace with Weber the influence of religious ideas on economic development. It is not less important to grasp the effect of economic arrangements accepted by an age on the opinion which it holds of the province of religion." Tawney's view is that technology influenced religion as strongly as religion influenced technology. Since my view of history is closer to Tawney's than to Weber's, I now ask the question: How can we push new technologies into directions conducive to social justice? How can we make ethics drive technology in such a way that the evil consequences are minimized and the good maximized? I shall hope to persuade you that the situation we are

in is not hopeless, that new technologies offer us real opportunities for making the world a happier place.

The Sun, the Genome, and the Internet

Finally, I turn to the positive side of my message. Technology guided by ethics has the power to help the billions of poor people all over the earth. My purpose is to help push technology in a new direction, away from toys for the rich and toward necessities for the poor. The time is ripe for this to happen. Three huge revolutionary forces are being harnessed just in time for the new century: the sun, the genome, and the Internet. These three forces are strong enough to reverse some of the worst evils of our time. The evils I am hoping to reverse are well known to you all. All over the earth, and especially in the poor countries to the south of us, millions of desperate people leave their villages and pour into overcrowded cities. There are now ten megacities in the world with populations twice as large as New York City. Soon there will be more. We all know that the increase of human population is one of the causes of the migration to cities. The other cause is the poverty and lack of jobs in villages. Both the population explosion and the poverty must be reversed if we are to have a decent future. Many experts on population say that if we can mitigate the poverty, the population will stabilize itself, as it has done in Europe and Japan. I am not an expert on population, so I won't say any more about that. I am saying that poverty can be reduced by a combination of solar energy, genetic engineering, and the Internet. Our task in the next century is to put the new technologies to work in the cause of social justice. Social justice means making the new technologies accessible to everyone, to bring wealth to poor countries and hope to poor people.

I have seen with my own eyes what happens to a village when the economic basis of life collapses, and I have seen how the economic basis of village life can be revived. My wife grew up in Westerhausen, a village in East Germany that was under Communist management. The Communist regime

took care of the village economy, selling the output of the farms to Russia at fixed prices, which gave the farmers economic security. The village remained beautiful and, on the whole, pleasant to live in. Nothing much had changed in the village since 1910. One thing the Communist regime did was organize a zoo, with a collection of animals maintained by a few professionals with a lot of help from the local school children. The village was justly proud of its zoo. The zoo was subsidized by the regime so it did not need to worry about being unprofitable. I visited the village under the old regime in 1975 and found it very friendly. Then came 1990 and the unification of Germany. Overnight, the economy of the village was wrecked. The farmers could no longer farm because nobody would buy their products. Russia could not buy because the price had to be paid in West German marks. German consumers would not buy because the local produce was not as good as that available in the supermarkets. The village farmers could not compete with the goods pouring in from France and Denmark. So the farmers were out of work. Most of the younger generation moved out of the village to compete for jobs in the cities, and most of the older generation remained. Many of them, both old and young, are still unemployed. The zoo, deprived of its subsidy, collapsed.

The sad exodus that I saw in the village of Westerhausen when I visited there in 1991 is the same exodus that is happening in villages all over the world. Everywhere the international market devalues the work of the village. Without work, the younger and the more enterprising people move out.

In the seven years since the unification, Westerhausen has slowly been recovering. Recovery is possible because of the process of gentrification. Wealthy people from the local towns move in and modernize the homes abandoned by the farmers. Cottages are demolished to make room for two-car garages. Ancient and narrow roads are widened. The village will survive as a community of nature lovers and commuters. Lying on the northern edge of the Harz Mountains, it is close to the big cities of northern Germany and

even closer to unspoiled mountain forests. Its permanent asset is natural beauty.

Two months ago my wife and I were back in the village. The change since we had last visited in 1991 was startling. We stayed in the elegant new home of a friend who had been in my wife's class in the village elementary school fifty years earlier. The village now looks well cared for and prosperous. The recovery from the disaster of 1990 has been slow and difficult, but it has been steady. The government did two things to mitigate the harshness of the free market: it allowed every homeowner to borrow money with almost zero interest from the government to modernize houses, and it allowed every farming cooperative to borrow money with almost zero interest to modernize farms. As a result, the houses that were not bought by outsiders are being modernized, and the few farmers who remained as farmers are flourishing. The zoo has been revived. In addition, there are some new enterprises. A Western immigrant has planted a large vineyard on a south-facing hillside and will soon be producing the first Westerhausen wines. My wife's family and many of her friends still live in the village. They gave us a warm and joyful welcome.

The probable future of Westerhausen can be seen in a thousand villages in England. The typical English village today is not primarily engaged in farming. The typical village remains beautiful and prosperous because of gentrification. Wealthy homeowners pay large sums of money for the privilege of living under a thatched roof. The thatching of roofs is one of the few ancient village crafts that still survives. The thatchers are mostly young, highly skilled, and well paid. The farmers who remain are either gentlemen amateurs, who run small farms as a hobby, or well-educated professionals, who run big farms as a business. The old population of peasant farmers, who used to live in the villages in poverty and squalor, disappeared long ago. Discreetly hidden in many of the villages are offices and factories engaged in high-tech industry. One of the head offices of IBM Europe is in the English village of Hursley not far from where I was born. In the

villages of France, at least in the area I know around Paris, the picture is much the same. Wealth came to the villages because they have what wealthy people seek: peace, security, and beauty.

What would it take to reverse the flow of jobs and people from villages to megacities all over the world? I believe the flow can be reversed by the same process of gentrification that is happening in Westerhausen. To make gentrification possible, the villages themselves must become sources of wealth. How can a godforsaken Mexican village become a source of wealth? Three facts can make it possible. First, solar energy is distributed equitably over the earth. Second, genetic engineering can make solar energy usable everywhere for the local creation of wealth. Third, the Internet can provide people in every village with the information and skills they need to develop their talents. The sun, the genome, and the Internet can work together to bring wealth to the villages of Mexico, just as the older technologies—electricity and automobiles—brought wealth to the villages of England. Let me talk briefly about each of the three new technologies, in turn.

Solar energy is most available where it is most needed—in the countryside rather than in cities, and in tropical countries, where most of the world's population lives, rather than in temperate latitudes. The quantity of solar energy is enormous compared with all other energy resources. Each square mile in the tropics receives about 1,000 megawatts averaged over day and night. This quantity of energy would be ample to support a dense population with all modern conveniences. Solar energy has not yet been used on a large scale for one simple reason: it is too expensive. It cannot compete in a free market with imported coal, oil, and natural gas. The country that has used solar energy on the largest scale is Brazil, where sugar was grown as an energy crop to make alcohol as a substitute for gasoline in cars and trucks. Brazil protected and subsidized the local alcohol industry. The experiment was technically successful, but the cost was high. Brazil has now reverted to free-market policies, and the

experiment is at an end. What the world needs is not high-cost subsidized solar energy, but solar energy cheap enough to compete with oil.

Solar energy is expensive today because it has to be collected from large areas and there is not yet a technology that covers large areas cheaply. One of the virtues of solar energy is the fact that it can be collected in many ways. It is adaptable to local conditions. The two main tools for collecting it are photoelectric panels, which convert sunlight directly into electricity, and energy crops, like the Brazilian sugar plantations, which convert sunlight into fuel. Roughly speaking, photoelectric collection is the method of choice for deserts, and energy crops are the method of choice for farmland and forests. Each method has its advantages and disadvantages. Photoelectric systems have high efficiency, typically between 10 percent and 15 percent, but are expensive to deploy and maintain. Energy crops have low efficiency, typically around 1 percent, and are expensive and messy to harvest. The electricity produced by photoelectric systems is intermittent and cannot be cheaply converted into storable forms of energy. Fuels produced from energy crops are storable and, therefore, more convenient.

To make solar energy cheap, we need a technology that combines the advantages of photovoltaic and biological systems. Two technical advances would make this possible. First, crop plants could be developed that convert sunlight into fuel with efficiency comparable to photovoltaic collectors, in the range of 10 percent rather than 1 percent. This would reduce the costs of land and harvesting by a large factor. Second, crop plants could be developed that do not need to be harvested at all. An energy crop could be a permanent forest with trees that convert sunlight to liquid fuel and deliver the fuel directly through their roots to a network of underground pipelines. If those two advantages could be combined, we would have a supply of solar energy that was cheap, abundant, ubiquitous, and environmentally benign.

The energy supply system of the future might be a large area of forest with species of trees varying from place to

place to suit the local climate and topography. We may hope that substantial parts of the forest would be nature reserves closed to human settlement and populated with wildlife so as to preserve the diversity of the natural ecologies. But the greater part could be open to human settlement, with teaming towns and villages under the trees. Landowners outside the nature reserves would be encouraged, but not compelled, to grow trees for energy. If the trees converted sunlight into fuel with 10 percent efficiency, landowners could sell the fuel for \$10,000 per acre per year and easily undercut the present price of gasoline. Owners of farmland and city lots alike would have a strong economic incentive to grow trees. The future energy plantation need not be a monotonous expanse of identical trees in regular rows. It could be as varied and as spontaneous as a natural woodland, interspersed with open spaces and houses, villages, towns, factories, and lakes.

To make this dream of a future landscape come true, the essential tool is genetic engineering. At present, large sums of money are being spent on sequencing the human genome. The Human Genome Project is motivated primarily by its medical applications. It will contribute enormously to the understanding and treatment of human diseases. It does not contribute directly to the engineering of trees. But alongside the human genome many other genomes are being sequenced—bacteria, yeast, worms, and fruit flies. For advancing the art of genetic engineering the genomes of simpler organisms are more useful than the human genome. Before long, we shall also have sequenced the genomes of the major crop plants—wheat, maize, and rice—and after that will come trees. Within a few decades, we shall have achieved a deep understanding of the genome, an understanding that will allow us to breed trees that will turn sunlight into fuel and still preserve the diversity that makes natural forests beautiful.

As soon as we can genetically engineer trees to use sunlight efficiently to make fuel, we shall also learn to breed trees that convert sunlight into useful chemicals of other kinds,

including silicon chips for computers and gasoline for cars. Economic forces will then move industries from cities to the country. Mining and manufacturing could be economically based on locally available solar energy, with genetically engineered creatures consuming and recycling the waste products. It might even become possible to build roads and buildings biologically, breeding little polyyps to lay down durable structures on land in the same way as their cousins build coral reefs in the ocean.

But the third, and most important, of the triad of new technologies is the Internet. The Internet is essential to enable businesses and farms in remote places to function as part of the modern global economy. The Internet will allow people in remote places to make business deals, buy and sell, keep in touch with their friends, continue their education, and follow their hobbies and avocations, with full knowledge of what is going on in the rest of the world.

This will not be the Internet of today, accessible only to computer-literate people in rich countries and to the wealthy elite in poor countries. It will be a truly global Internet, using a network of satellites in space for communication with places that fiber optics cannot reach and connected to local networks in every village. The new Internet will end the cultural isolation of poor countries and poor people.

Two technical problems have to be solved to make the Internet accessible to almost everybody on a global scale: large-scale architecture and the problem of the "last mile." Large-scale architecture means choosing the most efficient combination of landlines and satellite links to cover every corner of the globe. The Teledesic system of satellite communication now under development is intended to be a partial answer to this problem. The Teledesic system has 280 satellites in a dense network of low orbits, allowing any two points on the globe to be connected with minimum delay. If the Teledesic system fails, some other system will be designed to do the job. The problem of the "last mile" is more difficult. This is the problem of connecting individual homes and families, wherever they happen to be, with the nearest Internet termi-

be compelled by economic necessity to stay there. Many of us who have the freedom to choose, like the successful stockbrokers and business executives in England and Germany, will choose to live in villages.

So this is my dream: Solar energy, genetic engineering, and the Internet will work together to create a socially just world in which every Mexican village is as wealthy as Princeton. Of course, that is only a dream. Inequalities will persist. Poverty will not disappear. But I see a hope that the world will move far and fast in the directions I have been describing. Ethics must guide technology toward social justice. Let us all help to push the world in that direction as hard as we can. It does no harm to hope.

nal. The problem of the last mile has to be solved piecemeal, with methods depending on the local geography and the local culture. An ingenious method of solving the last-mile problem in urban American neighborhoods has been introduced recently by Paul Baran, the original inventor of the Internet. Baran's system is called Ricochet and consists of a multitude of small, wireless transmitters and receivers. Each user has a modem that communicates by radio with a local network. The feature that makes the system practical is that the transmitters constantly switch their frequencies so as not to interfere with one another. The system is flexible and cheap, avoiding the large expense of laying cable from the Internet terminal to every apartment and every house. It works well in the environment of urban America. It remains to be seen whether it is flexible and cheap enough to work well in the environment of a Mexican village or a Peruvian barrio.

Suppose, then, we can solve the technical problems of cheap solar energy, genetic engineering of industrial crop plants, and universal access to the Internet. What will follow? My thesis is that the solution of those three problems will bring about a worldwide social revolution, similar to the revolution we have seen in the villages of England and Germany. Cheap solar energy and genetic engineering will provide the basis for primary industries in the countryside. After that, the vast variety of secondary and tertiary economic activities that use the Internet for their coordination—food processing, publishing, education, entertainment, and health care—will follow the primary industries as they move from overgrown cities to country towns and villages. And as soon as the villages become rich, they will attract people and wealth back from the cities.

I am not suggesting that in the brave new world of the future everyone will be compelled to live in villages. Many of us will always prefer to live in large cities or in towns of moderate size. I am suggesting only that people should be free to choose. When wealth has moved back to the villages, people who live there will no longer be forced by economic necessity to move out, and people who live in megacities will no longer

Wired magazine, Issue 8.04, April 2000.

Why the future doesn't need us.

Our most powerful 21st-century technologies—robotics, genetic engineering, and nanotech—are threatening to make humans an endangered species.

By Bill Joy

From the moment I became involved in the creation of new technologies, their ethical dimensions have concerned me, but it was only in the autumn of 1998 that I became anxiously aware of how great are the dangers facing us in the 21st century. I can date the onset of my unease to the day I met Ray Kurzweil, the deservedly famous inventor of the first reading machine for the blind and many other amazing things.

Ray and I were both speakers at George Gilder's Telecosm conference, and I encountered him by chance in the bar of the hotel after both our sessions were over. I was sitting with John Searle, a Berkeley philosopher who studies consciousness. While we were talking, Ray approached and a conversation began, the subject of which haunts me to this day.

I had missed Ray's talk and the subsequent panel that Ray and John had been on, and they now picked right up where they'd left off, with Ray saying that the rate of improvement of technology was going to accelerate and that we were going to become robots or fuse with robots or something like that, and John countering that this couldn't happen, because the robots couldn't be conscious.

While I had heard such talk before, I had always felt sentient robots were in the realm of science fiction. But now, from someone I respected, I was hearing a strong argument that they were a near-term possibility. I was taken aback, especially given Ray's proven ability to imagine and create the future. I already knew that new technologies like genetic engineering and nanotechnology were giving us the power to remake the world, but a realistic and imminent scenario for intelligent robots surprised me.

It's easy to get jaded about such breakthroughs. We hear in the news almost every day of some kind of technological or scientific advance. Yet this was no ordinary prediction. In the hotel bar, Ray gave me a partial preprint of his then-forthcoming book *The Age of Spiritual Machines*, which outlined a utopia he foresaw—one in which humans gained near immortality by becoming one with robotic technology. On reading it, my sense of unease only intensified; I felt sure he had to be understating the dangers, understating the probability of a bad outcome along this path.

I found myself most troubled by a passage detailing a *dystopian* scenario:

THE NEW LUDDITE CHALLENGE

First let us postulate that the computer scientists succeed in developing intelligent machines that can do all things better than human beings can do them. In that case presumably all work will be done by vast, highly organized systems of machines and no human effort will be necessary. Either of two cases might occur. The machines might be permitted to make all of their own decisions without human oversight, or else human control over the machines might be retained.

If the machines are permitted to make all their own decisions, we can't make any conjectures as to the results, because it is impossible to guess how such machines might behave. We only point out

that the fate of the human race would be at the mercy of the machines. It might be argued that the human race would never be foolish enough to hand over all the power to the machines. But we are suggesting neither that the human race would voluntarily turn power over to the machines nor that the machines would willfully seize power. What we do suggest is that the human race might easily permit itself to drift into a position of such dependence on the machines that it would have no practical choice but to accept all of the machines' decisions. As society and the problems that face it become more and more complex and machines become more and more intelligent, people will let machines make more of their decisions for them, simply because machine-made decisions will bring better results than man-made ones. Eventually a stage may be reached at which the decisions necessary to keep the system running will be so complex that human beings will be incapable of making them intelligently. At that stage the machines will be in effective control. People won't be able to just turn the machines off, because they will be so dependent on them that turning them off would amount to suicide.

On the other hand it is possible that human control over the machines may be retained. In that case the average man may have control over certain private machines of his own, such as his car or his personal computer, but control over large systems of machines will be in the hands of a tiny elite—just as it is today, but with two differences. Due to improved techniques the elite will have greater control over the masses; and because human work will no longer be necessary the masses will be superfluous, a useless burden on the system. If the elite is ruthless they may simply decide to exterminate the mass of humanity. If they are humane they may use propaganda or other psychological or biological techniques to reduce the birth rate until the mass of humanity becomes extinct, leaving the world to the elite. Or, if the elite consists of soft-hearted liberals, they may decide to play the role of good shepherds to the rest of the human race. They will see to it that everyone's physical needs are satisfied, that all children are raised under psychologically hygienic conditions, that everyone has a wholesome hobby to keep him busy, and that anyone who may become dissatisfied undergoes "treatment" to cure his "problem." Of course, life will be so purposeless that people will have to be biologically or psychologically engineered either to remove their need for the power process or make them "sublimate" their drive for power into some harmless hobby. These engineered human beings may be happy in such a society, but they will most certainly not be free. They will have been reduced to the status of domestic animals.¹

In the book, you don't discover until you turn the page that the author of this passage is Theodore Kaczynski—the Unabomber. I am no apologist for Kaczynski. His bombs killed three people during a 17-year terror campaign and wounded many others. One of his bombs gravely injured my friend David Gelernter, one of the most brilliant and visionary computer scientists of our time. Like many of my colleagues, I felt that I could easily have been the Unabomber's next target.

Kaczynski's actions were murderous and, in my view, criminally insane. He is clearly a Luddite, but simply saying this does not dismiss his argument; as difficult as it is for me to acknowledge, I saw some merit in the reasoning in this single passage. I felt compelled to confront it.

Kaczynski's dystopian vision describes unintended consequences, a well-known problem with the design and use of technology, and one that is clearly related to Murphy's law—"Anything that can go wrong, will." (Actually, this is Finagle's law, which in itself shows that Finagle was right.) Our overuse of antibiotics has led to what may be the biggest such problem so far: the emergence of antibiotic-resistant and much more dangerous bacteria. Similar things happened when attempts to eliminate malarial mosquitoes using DDT caused them to acquire DDT resistance; malarial parasites likewise acquired multi-drug-resistant genes.²

The cause of many such surprises seems clear: The systems involved are complex, involving interaction among and feedback between many parts. Any changes to such a system will cascade in ways that are difficult to predict; this is especially true when human actions are involved.

I started showing friends the Kaczynski quote from *The Age of Spiritual Machines*; I would hand them Kurzweil's book, let them read the quote, and then watch their reaction as they discovered who had written it. At around the same time, I found Hans Moravec's book *Robot: Mere Machine to Transcendent Mind*. Moravec is one of the leaders in robotics research, and was a founder of the world's largest robotics research program, at Carnegie Mellon University. *Robot* gave me more material to try out on my friend material surprisingly supportive of Kaczynski's argument. For example:

The Short Run (Early 2000s)

Biological species almost never survive encounters with superior competitors. Ten million years ago, South and North America were separated by a sunken Panama isthmus. South America, like Australia today, was populated by marsupial mammals, including pouched equivalents of rats, deers, and tigers. When the isthmus connecting North and South America rose, it took only a few thousand years for the northern placental species, with slightly more effective metabolisms and reproductive and nervous systems, to displace and eliminate almost all the southern marsupials.

In a completely free marketplace, superior robots would surely affect humans as North American placentals affected South American marsupials (and as humans have affected countless species). Robotic industries would compete vigorously among themselves for matter, energy, and space, incidentally driving their price beyond human reach. Unable to afford the necessities of life, biological humans would be squeezed out of existence.

There is probably some breathing room, because we do not live in a completely free marketplace. Government coerces nonmarket behavior, especially by collecting taxes. Judiciously applied, governmental coercion could support human populations in high style on the fruits of robot labor, perhaps for a long while.

A textbook dystopia—and Moravec is just getting wound up. He goes on to discuss how our main job in the 21st century will be “ensuring continued cooperation from the robot industries” by passing laws decreeing that they be “nice,”³ and to describe how seriously dangerous a human can be “once transformed into an unbounded superintelligent robot.” Moravec's view is that the robots will eventually succeed us—that humans clearly face extinction.

I decided it was time to talk to my friend Danny Hillis. Danny became famous as the cofounder of Thinking Machines Corporation, which built a very powerful parallel supercomputer. Despite my current job title of Chief Scientist at Sun Microsystems, I am more a computer architect than a scientist, and I respect Danny's knowledge of the information and physical sciences more than that of any other single person I know. Danny is also a highly regarded futurist who thinks long-term — four years ago he started the Long Now Foundation, which is building a clock designed to last 10,000 years, in an attempt to draw attention to the pitifully short attention span of our society. (See “[Test of Time](#),” *Wired* 8.03, page 78.)

So I flew to Los Angeles for the express purpose of having dinner with Danny and his wife, Pati. I went through my now-familiar routine, trotting out the ideas and passages that I found so disturbing. Danny's answer—directed specifically at Kurzweil's scenario of humans merging with robots—came swiftly, and quite surprised me. He said, simply, that the changes would come gradually, and that we would get used to them.

But I guess I wasn't totally surprised. I had seen a quote from Danny in Kurzweil's book in which he said, “I'm as fond of my body as anyone, but if I can be 200 with a body of silicon, I'll take it.” It seemed that he was at peace with this process and its attendant risks, while I was not.

While talking and thinking about Kurzweil, Kaczynski, and Moravec, I suddenly remembered a novel I had read almost 20 years ago—*The White Plague*, by Frank Herbert—in which a molecular biologist is driven insane by the senseless murder of his family. To seek revenge he constructs and disseminates a new and

highly contagious plague that kills widely but selectively. (We're lucky Kaczynski was a mathematician, not a molecular biologist.) I was also reminded of the Borg of *Star Trek*, a hive of partly biological, partly robotic creatures with a strong destructive streak. Borg-like disasters are a staple of science fiction, so why hadn't I been more concerned about such robotic dystopias earlier? Why weren't other people more concerned about these nightmarish scenarios?

Part of the answer certainly lies in our attitude toward the new—in our bias toward instant familiarity and unquestioning acceptance. Accustomed to living with almost routine scientific breakthroughs, we have yet to come to terms with the fact that the most compelling 21st-century technologies—robotics, genetic engineering, and nanotechnology—pose a different threat than the technologies that have come before. Specifically, robots, engineered organisms, and nanobots share a dangerous amplifying factor: They can self-replicate. A bomb is blown up only once—but one bot can become many, and quickly get out of control.

Much of my work over the past 25 years has been on computer networking, where the sending and receiving of messages creates the opportunity for out-of-control replication. But while replication in a computer or a computer network can be a nuisance, at worst it disables a machine or takes down a network or network service. Uncontrolled self-replication in these newer technologies runs a much greater risk: a risk of substantial damage in the physical world.

Each of these technologies also offers untold promise: The vision of near immortality that Kurzweil sees in his robot dreams drives us forward; genetic engineering may soon provide treatments, if not outright cures, for most diseases; and nanotechnology and nanomedicine can address yet more ills. Together they could significantly extend our average life span and improve the quality of our lives. Yet, with each of these technologies, a sequence of small, individually sensible advances leads to an accumulation of great power and, concomitantly, great danger.

What was different in the 20th century? Certainly, the technologies underlying the weapons of mass destruction (WMD)—nuclear, biological, and chemical (NBC)—were powerful, and the weapons an enormous threat. But building nuclear weapons required, at least for a time, access to both rare—indeed, effectively unavailable—raw materials and highly protected information; biological and chemical weapons programs also tended to require large-scale activities.

The 21st-century technologies—genetics, nanotechnology, and robotics (GNR)—are so powerful that they can spawn whole new classes of accidents and abuses. Most dangerously, for the first time, these accidents and abuses are widely within the reach of individuals or small groups. They will not require large facilities or rare raw materials. Knowledge alone will enable the use of them.

Thus we have the possibility not just of weapons of mass destruction but of knowledge-enabled mass destruction (KMD), this destructiveness hugely amplified by the power of self-replication.

I think it is no exaggeration to say we are on the cusp of the further perfection of extreme evil, an evil whose possibility spreads well beyond that which weapons of mass destruction bequeathed to the nation-states, on to a surprising and terrible empowerment of extreme individuals.

Nothing about the way I got involved with computers suggested to me that I was going to be facing these kinds of issues.

My life has been driven by a deep need to ask questions and find answers. When I was 3, I was already reading, so my father took me to the elementary school, where I sat on the principal's lap and read him a story. I started school early, later skipped a grade, and escaped into books—I was incredibly motivated to learn. I asked lots of questions, often driving adults to distraction.

As a teenager I was very interested in science and technology. I wanted to be a ham radio operator but didn't have the money to buy the equipment. Ham radio was the Internet of its time: very addictive and quite solitary. Money issues aside, my mother put her foot down—I was not to be a ham; I was antisocial enough already.

I may not have had many close friends, but I was awash in ideas. By high school, I had discovered the great science fiction writers. I remember especially Heinlein's *Have Spacesuit Will Travel* and Asimov's *I, Robot*, with its Three Laws of Robotics. I was enchanted by the descriptions of space travel, and wanted to have a telescope to look at the stars; since I had no money to buy or make one, I checked books on telescope-making out of the library and read about making them instead. I soared in my imagination.

Thursday nights my parents went bowling, and we kids stayed home alone. It was the night of Gene Roddenberry's original *Star Trek*, and the program made a big impression on me. I came to accept its notion that humans had a future in space, Western-style, with big heroes and adventures. Roddenberry's vision of the centuries to come was one with strong moral values, embodied in codes like the Prime Directive: to not interfere in the development of less technologically advanced civilizations. This had an incredible appeal to me; ethical humans, not robots, dominated this future, and I took Roddenberry's dream as part of my own.

I excelled in mathematics in high school, and when I went to the University of Michigan as an undergraduate engineering student I took the advanced curriculum of the mathematics majors. Solving math problems was an exciting challenge, but when I discovered computers I found something much more interesting: a machine into which you could put a program that attempted to solve a problem, after which the machine quickly checked the solution. The computer had a clear notion of correct and incorrect, true and false. Were my ideas correct? The machine could tell me. This was very seductive.

I was lucky enough to get a job programming early supercomputers and discovered the amazing power of large machines to numerically simulate advanced designs. When I went to graduate school at UC Berkeley in the mid-1970s, I started staying up late, often all night, inventing new worlds inside the machines. Solving problems. Writing the code that argued so strongly to be written.

In *The Agony and the Ecstasy*, Irving Stone's biographical novel of Michelangelo, Stone described vividly how Michelangelo released the statues from the stone, "breaking the marble spell," carving from the images in his mind.⁴ In my most ecstatic moments, the software in the computer emerged in the same way. Once I had imagined it in my mind I felt that it was already there in the machine, waiting to be released. Staying up all night seemed a small price to pay to free it—to give the ideas concrete form.

After a few years at Berkeley I started to send out some of the software I had written—an instructional Pascal system, Unix utilities, and a text editor called vi (which is still, to my surprise, widely used more than 20 years later)—to others who had similar small PDP-11 and VAX minicomputers. These adventures in software eventually turned into the Berkeley version of the Unix operating system, which became a personal "success disaster"—so many people wanted it that I never finished my PhD. Instead I got a job working for DARPA putting Berkeley Unix on the Internet and fixing it to be reliable and to run large research applications well. This was all great fun and very rewarding. And, frankly, I saw no robots here, or anywhere near.

Still, by the early 1980s, I was drowning. The Unix releases were very successful, and my little project of one soon had money and some staff, but the problem at Berkeley was always office space rather than money—there wasn't room for the help the project needed, so when the other founders of Sun Microsystems showed up I jumped at the chance to join them. At Sun, the long hours continued into the early days of workstations and personal computers, and I have enjoyed participating in the creation of advanced microprocessor technologies and Internet technologies such as Java and Jini.

From all this, I trust it is clear that I am not a Luddite. I have always, rather, had a strong belief in the value of the scientific search for truth and in the ability of great engineering to bring material progress. The Industrial Revolution has immeasurably improved everyone's life over the last couple hundred years, and I always expected my career to involve the building of worthwhile solutions to real problems, one problem at a time.

I have not been disappointed. My work has had more impact than I had ever hoped for and has been more widely used than I could have reasonably expected. I have spent the last 20 years still trying to figure out how to make computers as reliable as I want them to be (they are not nearly there yet) and how to make them simple to use (a goal that has met with even less relative success). Despite some progress, the problems that remain seem even more daunting.

But while I was aware of the moral dilemmas surrounding technology's consequences in fields like weapons research, I did not expect that I would confront such issues in my own field, or at least not so soon.

Perhaps it is always hard to see the bigger impact while you are in the vortex of a change. Failing to understand the consequences of our inventions while we are in the rapture of discovery and innovation seems to be a common fault of scientists and technologists; we have long been driven by the overarching desire to know that is the nature of science's quest, not stopping to notice that the progress to newer and more powerful technologies can take on a life of its own.

I have long realized that the big advances in information technology come not from the work of computer scientists, computer architects, or electrical engineers, but from that of physical scientists. The physicists Stephen Wolfram and Brosl Hasslacher introduced me, in the early 1980s, to chaos theory and nonlinear systems. In the 1990s, I learned about complex systems from conversations with Danny Hillis, the biologist Stuart Kauffman, the Nobel-laureate physicist Murray Gell-Mann, and others. Most recently, Hasslacher and the electrical engineer and device physicist Mark Reed have been giving me insight into the incredible possibilities of molecular electronics.

In my own work, as codesigner of three microprocessor architectures—SPARC, picoJava, and MAJC—and as the designer of several implementations thereof, I've been afforded a deep and firsthand acquaintance with Moore's law. For decades, Moore's law has correctly predicted the exponential rate of improvement of semiconductor technology. Until last year I believed that the rate of advances predicted by Moore's law might continue only until roughly 2010, when some physical limits would begin to be reached. It was not obvious to me that a new technology would arrive in time to keep performance advancing smoothly.

But because of the recent rapid and radical progress in molecular electronics—where individual atoms and molecules replace lithographically drawn transistors—and related nanoscale technologies, we should be able to meet or exceed the Moore's law rate of progress for another 30 years. By 2030, we are likely to be able to build machines, in quantity, a million times as powerful as the personal computers of today—sufficient to implement the dreams of Kurzweil and Moravec.

As this enormous computing power is combined with the manipulative advances of the physical sciences and the new, deep understandings in genetics, enormous transformative power is being unleashed. These combinations open up the opportunity to completely redesign the world, for better or worse: The replicating and evolving processes that have been confined to the natural world are about to become realms of human endeavor.

In designing software and microprocessors, I have never had the feeling that I was designing an intelligent machine. The software and hardware is so fragile and the capabilities of the machine to "think" so clearly absent that, even as a possibility, this has always seemed very far in the future.

But now, with the prospect of human-level computing power in about 30 years, a new idea suggests itself: that I may be working to create tools which will enable the construction of the technology that may replace our species. How do I feel about this? Very uncomfortable. Having struggled my entire career to build reliable software systems, it seems to me more than likely that this future will not work out as well as some people may imagine. My personal experience suggests we tend to overestimate our design abilities.

Given the incredible power of these new technologies, shouldn't we be asking how we can best coexist with them? And if our own extinction is a likely, or even possible, outcome of our technological development, shouldn't we proceed with great caution?

The dream of robotics is, first, that intelligent machines can do our work for us, allowing us lives of leisure, restoring us to Eden. Yet in his history of such ideas, *Darwin Among the Machines*, George Dyson warns: "In the game of life and evolution there are three players at the table: human beings, nature, and machines. I am firmly on the side of nature. But nature, I suspect, is on the side of the machines." As we have seen, Moravec agrees, believing we may well not survive the encounter with the superior robot species.

How soon could such an intelligent robot be built? The coming advances in computing power seem to make it possible by 2030. And once an intelligent robot exists, it is only a small step to a robot species—to an intelligent robot that can make evolved copies of itself.

A second dream of robotics is that we will gradually replace ourselves with our robotic technology, achieving near immortality by downloading our consciousnesses; it is this process that Danny Hillis thinks we will gradually get used to and that Ray Kurzweil elegantly details in *The Age of Spiritual Machines*. (We are beginning to see intimations of this in the implantation of computer devices into the human body, as illustrated on the [cover](#) of *Wired* 8.02.)

But if we are downloaded into our technology, what are the chances that we will thereafter be ourselves or even human? It seems to me far more likely that a robotic existence would not be like a human one in any sense that we understand, that the robots would in no sense be our children, that on this path our humanity may well be lost.

Genetic engineering promises to revolutionize agriculture by increasing crop yields while reducing the use of pesticides; to create tens of thousands of novel species of bacteria, plants, viruses, and animals; to replace reproduction, or supplement it, with cloning; to create cures for many diseases, increasing our life span and our quality of life; and much, much more. We now know with certainty that these profound changes in the biological sciences are imminent and will challenge all our notions of what life is.

Technologies such as human cloning have in particular raised our awareness of the profound ethical and moral issues we face. If, for example, we were to reengineer ourselves into several separate and unequal species using the power of genetic engineering, then we would threaten the notion of equality that is the very cornerstone of our democracy.

Given the incredible power of genetic engineering, it's no surprise that there are significant safety issues in its use. My friend Amory Lovins recently cowrote, along with Hunter Lovins, an editorial that provides an ecological view of some of these dangers. Among their concerns: that "the new botany aligns the development of plants with their economic, not evolutionary, success." (See "[A Tale of Two Botanies](#)," page 247.) Amory's long career has been focused on energy and resource efficiency by taking a whole-system view of human-made systems; such a whole-system view often finds simple, smart solutions to otherwise seemingly difficult problems, and is usefully applied here as well.

After reading the Lovins' editorial, I saw an op-ed by Gregg Easterbrook in *The New York Times* (November 19, 1999) about genetically engineered crops, under the headline: "Food for the Future: Someday, rice will have built-in vitamin A. Unless the Luddites win."

Are Amory and Hunter Lovins Luddites? Certainly not. I believe we all would agree that golden rice, with its built-in vitamin A, is probably a good thing, if developed with proper care and respect for the likely dangers in moving genes across species boundaries.

Awareness of the dangers inherent in genetic engineering is beginning to grow, as reflected in the Lovins' editorial. The general public is aware of, and uneasy about, genetically modified foods, and seems to be rejecting the notion that such foods should be permitted to be unlabeled.

But genetic engineering technology is already very far along. As the Lovins note, the USDA has already approved about 50 genetically engineered crops for unlimited release; more than half of the world's soybeans and a third of its corn now contain genes spliced in from other forms of life.

While there are many important issues here, my own major concern with genetic engineering is narrower: that it gives the power—whether militarily, accidentally, or in a deliberate terrorist act—to create a White Plague.

The many wonders of nanotechnology were first imagined by the Nobel-laureate physicist Richard Feynman in a speech he gave in 1959, subsequently published under the title "There's Plenty of Room at the Bottom." The book that made a big impression on me, in the mid-'80s, was Eric Drexler's *Engines of Creation*, in which he described beautifully how manipulation of matter at the atomic level could create a utopian future of abundance, where just about everything could be made cheaply, and almost any imaginable disease or physical problem could be solved using nanotechnology and artificial intelligences.

A subsequent book, *Unbounding the Future: The Nanotechnology Revolution*, which Drexler cowrote, imagines some of the changes that might take place in a world where we had molecular-level "assemblers." Assemblers could make possible incredibly low-cost solar power, cures for cancer and the common cold by augmentation of the human immune system, essentially complete cleanup of the environment, incredibly inexpensive pocket supercomputers—in fact, any product would be manufacturable by assemblers at a cost no greater than that of wood—spaceflight more accessible than transoceanic travel today, and restoration of extinct species.

I remember feeling good about nanotechnology after reading *Engines of Creation*. As a technologist, it gave me a sense of calm—that is, nanotechnology showed us that incredible progress was possible, and indeed perhaps inevitable. If nanotechnology was our future, then I didn't feel pressed to solve so many problems in the present. I would get to Drexler's utopian future in due time; I might as well enjoy life more in the here and now. It didn't make sense, given his vision, to stay up all night, all the time.

Drexler's vision also led to a lot of good fun. I would occasionally get to describe the wonders of nanotechnology to others who had not heard of it. After teasing them with all the things Drexler described I would give a homework assignment of my own: "Use nanotechnology to create a vampire; for extra credit create an antidote."

With these wonders came clear dangers, of which I was acutely aware. As I said at a nanotechnology conference in 1989, "We can't simply do our science and not worry about these ethical issues."⁵ But my subsequent conversations with physicists convinced me that nanotechnology might not even work—or, at least, it wouldn't work anytime soon. Shortly thereafter I moved to Colorado, to a skunk works I had set up, and the focus of my work shifted to software for the Internet, specifically on ideas that became Java and Jini.

Then, last summer, Brosl Hasslacher told me that nanoscale molecular electronics was now practical. This was *new* news, at least to me, and I think to many people—and it radically changed my opinion about nanotechnology. It sent me back to *Engines of Creation*. Rereading Drexler's work after more than 10 years, I was dismayed to realize how little I had remembered of its lengthy section called "Dangers and Hopes," including a discussion of how nanotechnologies can become "engines of destruction." Indeed, in my rereading of this cautionary material today, I am struck by how naive some of Drexler's safeguard proposals seem, and how much greater I judge the dangers to be now than even he seemed to then. (Having anticipated and described many technical and political problems with nanotechnology, Drexler started the Foresight Institute in the late 1980s "to help prepare society for anticipated advanced technologies"—most important, nanotechnology.)

The enabling breakthrough to assemblers seems quite likely within the next 20 years. Molecular electronics—the new subfield of nanotechnology where individual molecules are circuit elements—should mature quickly and become enormously lucrative within this decade, causing a large incremental investment in all nanotechnologies.

Unfortunately, as with nuclear technology, it is far easier to create destructive uses for nanotechnology than constructive ones. Nanotechnology has clear military and terrorist uses, and you need not be suicidal to release a massively destructive nanotechnological device—such devices can be built to be selectively destructive, affecting, for example, only a certain geographical area or a group of people who are genetically distinct.

An immediate consequence of the Faustian bargain in obtaining the great power of nanotechnology is that we run a grave risk—the risk that we might destroy the biosphere on which all life depends.

As Drexler explained:

"Plants" with "leaves" no more efficient than today's solar cells could out-compete real plants, crowding the biosphere with an inedible foliage. Tough omnivorous "bacteria" could out-compete real bacteria: They could spread like blowing pollen, replicate swiftly, and reduce the biosphere to dust in a matter of days. Dangerous replicators could easily be too tough, small, and rapidly spreading to stop—at least if we make no preparation. We have trouble enough controlling viruses and fruit flies.

Among the cognoscenti of nanotechnology, this threat has become known as the "gray goo problem." Though masses of uncontrolled replicators need not be gray or gooey, the term "gray goo" emphasizes that replicators able to obliterate life might be less inspiring than a single species of crabgrass. They might be superior in an evolutionary sense, but this need not make them valuable.

The gray goo threat makes one thing perfectly clear: We cannot afford certain kinds of accidents with replicating assemblers.

Gray goo would surely be a depressing ending to our human adventure on Earth, far worse than mere fire or ice, and one that could stem from a simple laboratory accident.[6](#) Oops.

It is most of all the power of destructive self-replication in genetics, nanotechnology, and robotics (GNR) that should give us pause. Self-replication is the modus operandi of genetic engineering, which uses the machinery of the cell to replicate its designs, and the prime danger underlying gray goo in nanotechnology. Stories of run-amok robots like the Borg, replicating or mutating to escape from the ethical constraints imposed on them by their creators, are well established in our science fiction books and movies. It is even possible that self-replication may be more fundamental than we thought, and hence harder—or even impossible—to control. A recent article by Stuart Kauffman in *Nature* titled "Self-Replication: Even

Peptides Do It” discusses the discovery that a 32-amino-acid peptide can “autocatalyze its own synthesis.” We don't know how widespread this ability is, but Kauffman notes that it may hint at “a route to self-reproducing molecular systems on a basis far wider than Watson-Crick base-pairing.”⁷

In truth, we have had in hand for years clear warnings of the dangers inherent in widespread knowledge of GNR technologies—of the possibility of knowledge alone enabling mass destruction. But these warnings haven't been widely publicized; the public discussions have been clearly inadequate. There is no profit in publicizing the dangers.

The nuclear, biological, and chemical (NBC) technologies used in 20th-century weapons of mass destruction were and are largely military, developed in government laboratories. In sharp contrast, the 21st-century GNR technologies have clear commercial uses and are being developed almost exclusively by corporate enterprises. In this age of triumphant commercialism, technology—with science as its handmaiden—is delivering a series of almost magical inventions that are the most phenomenally lucrative ever seen. We are aggressively pursuing the promises of these new technologies within the now-unchallenged system of global capitalism and its manifold financial incentives and competitive pressures.

This is the first moment in the history of our planet when any species, by its own voluntary actions, has become a danger to itself—as well as to vast numbers of others.

It might be a familiar progression, transpiring on many worlds—a planet, newly formed, placidly revolves around its star; life slowly forms; a kaleidoscopic procession of creatures evolves; intelligence emerges which, at least up to a point, confers enormous survival value; and then technology is invented. It dawns on them that there are such things as laws of Nature, that these laws can be revealed by experiment, and that knowledge of these laws can be made both to save and to take lives, both on unprecedented scales. Science, they recognize, grants immense powers. In a flash, they create world-altering contrivances. Some planetary civilizations see their way through, place limits on what may and what must not be done, and safely pass through the time of perils. Others, not so lucky or so prudent, perish.

That is Carl Sagan, writing in 1994, in *Pale Blue Dot*, a book describing his vision of the human future in space. I am only now realizing how deep his insight was, and how sorely I miss, and will miss, his voice. For all its eloquence, Sagan's contribution was not least that of simple common sense—an attribute that, along with humility, many of the leading advocates of the 21st-century technologies seem to lack.

I remember from my childhood that my grandmother was strongly against the overuse of antibiotics. She had worked since before the first World War as a nurse and had a commonsense attitude that taking antibiotics, unless they were absolutely necessary, was bad for you.

It is not that she was an enemy of progress. She saw much progress in an almost 70-year nursing career; my grandfather, a diabetic, benefited greatly from the improved treatments that became available in his lifetime. But she, like many levelheaded people, would probably think it greatly arrogant for us, now, to be designing a robotic “replacement species,” when we obviously have so much trouble making relatively simple things work, and so much trouble managing—or even understanding—ourselves.

I realize now that she had an awareness of the nature of the order of life, and of the necessity of living with and respecting that order. With this respect comes a necessary humility that we, with our early-21st-century chutzpah, lack at our peril. The commonsense view, grounded in this respect, is often right, in advance of the scientific evidence. The clear fragility and inefficiencies of the human-made systems we have built should give us all pause; the fragility of the systems I have worked on certainly humbles me.

We should have learned a lesson from the making of the first atomic bomb and the resulting arms race. We didn't do well then, and the parallels to our current situation are troubling.

The effort to build the first atomic bomb was led by the brilliant physicist J. Robert Oppenheimer. Oppenheimer was not naturally interested in politics but became painfully aware of what he perceived as the grave threat to Western civilization from the Third Reich, a threat surely grave because of the possibility that Hitler might obtain nuclear weapons. Energized by this concern, he brought his strong intellect, passion for physics, and charismatic leadership skills to Los Alamos and led a rapid and successful effort by an incredible collection of great minds to quickly invent the bomb.

What is striking is how this effort continued so naturally after the initial impetus was removed. In a meeting shortly after V-E Day with some physicists who felt that perhaps the effort should stop, Oppenheimer argued to continue. His stated reason seems a bit strange: not because of the fear of large casualties from an invasion of Japan, but because the United Nations, which was soon to be formed, should have foreknowledge of atomic weapons. A more likely reason the project continued is the momentum that had built up—the first atomic test, Trinity, was nearly at hand.

We know that in preparing this first atomic test the physicists proceeded despite a large number of possible dangers. They were initially worried, based on a calculation by Edward Teller, that an atomic explosion might set fire to the atmosphere. A revised calculation reduced the danger of destroying the world to a three-in-a-million chance. (Teller says he was later able to dismiss the prospect of atmospheric ignition entirely.) Oppenheimer, though, was sufficiently concerned about the result of Trinity that he arranged for a possible evacuation of the southwest part of the state of New Mexico. And, of course, there was the clear danger of starting a nuclear arms race.

Within a month of that first, successful test, two atomic bombs destroyed Hiroshima and Nagasaki. Some scientists had suggested that the bomb simply be demonstrated, rather than dropped on Japanese cities—saying that this would greatly improve the chances for arms control after the war—but to no avail. With the tragedy of Pearl Harbor still fresh in Americans' minds, it would have been very difficult for President Truman to order a demonstration of the weapons rather than use them as he did—the desire to quickly end the war and save the lives that would have been lost in any invasion of Japan was very strong. Yet the overriding truth was probably very simple: As the physicist Freeman Dyson later said, “The reason that it was dropped was just that nobody had the courage or the foresight to say no.”

It's important to realize how shocked the physicists were in the aftermath of the bombing of Hiroshima, on August 6, 1945. They describe a series of waves of emotion: first, a sense of fulfillment that the bomb worked, then horror at all the people that had been killed, and then a convincing feeling that on no account should another bomb be dropped. Yet of course another bomb was dropped, on Nagasaki, only three days after the bombing of Hiroshima.

In November 1945, three months after the atomic bombings, Oppenheimer stood firmly behind the scientific attitude, saying, “It is not possible to be a scientist unless you believe that the knowledge of the world, and the power which this gives, is a thing which is of intrinsic value to humanity, and that you are using it to help in the spread of knowledge and are willing to take the consequences.”

Oppenheimer went on to work, with others, on the Acheson-Lilienthal report, which, as Richard Rhodes says in his recent book *Visions of Technology*, “found a way to prevent a clandestine nuclear arms race without resorting to armed world government”; their suggestion was a form of relinquishment of nuclear weapons work by nation-states to an international agency.

This proposal led to the Baruch Plan, which was submitted to the United Nations in June 1946 but never adopted (perhaps because, as Rhodes suggests, Bernard Baruch had “insisted on burdening the plan with conventional sanctions,” thereby inevitably dooming it, even though it would “almost certainly have been rejected by Stalinist Russia anyway”). Other efforts to promote sensible steps toward internationalizing nuclear power to prevent an arms race ran afoul either of US politics and internal distrust, or distrust by the Soviets. The opportunity to avoid the arms race was lost, and very quickly.

Two years later, in 1948, Oppenheimer seemed to have reached another stage in his thinking, saying, “In some sort of crude sense which no vulgarity, no humor, no overstatement can quite extinguish, the physicists have known sin; and this is a knowledge they cannot lose.”

In 1949, the Soviets exploded an atom bomb. By 1955, both the US and the Soviet Union had tested hydrogen bombs suitable for delivery by aircraft. And so the nuclear arms race began.

Nearly 20 years ago, in the documentary *The Day After Trinity*, Freeman Dyson summarized the scientific attitudes that brought us to the nuclear precipice:

“I have felt it myself. The glitter of nuclear weapons. It is irresistible if you come to them as a scientist. To feel it’s there in your hands, to release this energy that fuels the stars, to let it do your bidding. To perform these miracles, to lift a million tons of rock into the sky. It is something that gives people an illusion of illimitable power, and it is, in some ways, responsible for all our troubles—this, what you might call technical arrogance, that overcomes people when they see what they can do with their minds.”⁸

Now, as then, we are creators of new technologies and stars of the imagined future, driven—this time by great financial rewards and global competition—despite the clear dangers, hardly evaluating what it may be like to try to live in a world that is the realistic outcome of what we are creating and imagining.

In 1947, *The Bulletin of the Atomic Scientists* began putting a Doomsday Clock on its cover. For more than 50 years, it has shown an estimate of the relative nuclear danger we have faced, reflecting the changing international conditions. The hands on the clock have moved 15 times and today, standing at nine minutes to midnight, reflect continuing and real danger from nuclear weapons. The recent addition of India and Pakistan to the list of nuclear powers has increased the threat of failure of the nonproliferation goal, and this danger was reflected by moving the hands closer to midnight in 1998.

In our time, how much danger do we face, not just from nuclear weapons, but from all of these technologies? How high are the extinction risks?

The philosopher John Leslie has studied this question and concluded that the risk of human extinction is at least 30 percent,⁹ while Ray Kurzweil believes we have “a better than even chance of making it through,” with the caveat that he has “always been accused of being an optimist.” Not only are these estimates not encouraging, but they do not include the probability of many horrid outcomes that lie short of extinction.

Faced with such assessments, some serious people are already suggesting that we simply move beyond Earth as quickly as possible. We would colonize the galaxy using von Neumann probes, which hop from star system to star system, replicating as they go. This step will almost certainly be necessary 5 billion years from now (or sooner if our solar system is disastrously impacted by the impending collision of our galaxy with the Andromeda galaxy within the next 3 billion years), but if we take Kurzweil and Moravec at their word it might be necessary by the middle of this century.

What are the moral implications here? If we must move beyond Earth this quickly in order for the species to survive, who accepts the responsibility for the fate of those (most of us, after all) who are left behind? And even if we scatter to the stars, isn’t it likely that we may take our problems with us or find, later, that they have followed us? The fate of our species on Earth and our fate in the galaxy seem inextricably linked.

Another idea is to erect a series of shields to defend against each of the dangerous technologies. The Strategic Defense Initiative, proposed by the Reagan administration, was an attempt to design such a shield against the threat of a nuclear attack from the Soviet Union. But as Arthur C. Clarke, who was privy to discussions about the project, observed: “Though it might be possible, at vast expense, to construct local

defense systems that would ‘only’ let through a few percent of ballistic missiles, the much touted idea of a national umbrella was nonsense. Luis Alvarez, perhaps the greatest experimental physicist of this century, remarked to me that the advocates of such schemes were ‘very bright guys with no common sense.’”

Clarke continued: “Looking into my often cloudy crystal ball, I suspect that a total defense might indeed be possible in a century or so. But the technology involved would produce, as a by-product, weapons so terrible that no one would bother with anything as primitive as ballistic missiles.” [10](#)

In *Engines of Creation*, Eric Drexler proposed that we build an active nanotechnological shield—a form of immune system for the biosphere—to defend against dangerous replicators of all kinds that might escape from laboratories or otherwise be maliciously created. But the shield he proposed would itself be extremely dangerous—nothing could prevent it from developing autoimmune problems and attacking the biosphere itself. [11](#)

Similar difficulties apply to the construction of shields against robotics and genetic engineering. These technologies are too powerful to be shielded against in the time frame of interest; even if it were possible to implement defensive shields, the side effects of their development would be at least as dangerous as the technologies we are trying to protect against.

These possibilities are all thus either undesirable or unachievable or both. The only realistic alternative I see is relinquishment: to limit development of the technologies that are too dangerous, by limiting our pursuit of certain kinds of knowledge.

Yes, I know, knowledge is good, as is the search for new truths. We have been seeking knowledge since ancient times. Aristotle opened his *Metaphysics* with the simple statement: “All men by nature desire to know.” We have, as a bedrock value in our society, long agreed on the value of open access to information, and recognize the problems that arise with attempts to restrict access to and development of knowledge. In recent times, we have come to revere scientific knowledge.

But despite the strong historical precedents, if open access to and unlimited development of knowledge henceforth puts us all in clear danger of extinction, then common sense demands that we reexamine even these basic, long-held beliefs.

It was Nietzsche who warned us, at the end of the 19th century, not only that God is dead but that “faith in science, which after all exists undeniably, cannot owe its origin to a calculus of utility; it must have originated *in spite of* the fact that the disutility and dangerousness of the ‘will to truth,’ of ‘truth at any price’ is proved to it constantly.” It is this further danger that we now fully face—the consequences of our truth-seeking. The truth that science seeks can certainly be considered a dangerous substitute for God if it is likely to lead to our extinction.

If we could agree, as a species, what we wanted, where we were headed, and why, then we would make our future much less dangerous—then we might understand what we can and should relinquish. Otherwise, we can easily imagine an arms race developing over GNR technologies, as it did with the NBC technologies in the 20th century. This is perhaps the greatest risk, for once such a race begins, it’s very hard to end it. This time—unlike during the Manhattan Project—we aren’t in a war, facing an implacable enemy that is threatening our civilization; we are driven, instead, by our habits, our desires, our economic system, and our competitive need to know.

I believe that we all wish our course could be determined by our collective values, ethics, and morals. If we had gained more collective wisdom over the past few thousand years, then a dialogue to this end would be more practical, and the incredible powers we are about to unleash would not be nearly so troubling.

One would think we might be driven to such a dialogue by our instinct for self-preservation. Individuals clearly have this desire, yet as a species our behavior seems to be not in our favor. In dealing with the

nuclear threat, we often spoke dishonestly to ourselves and to each other, thereby greatly increasing the risks. Whether this was politically motivated, or because we chose not to think ahead, or because when faced with such grave threats we acted irrationally out of fear, I do not know, but it does not bode well.

The new Pandora's boxes of genetics, nanotechnology, and robotics are almost open, yet we seem hardly to have noticed. Ideas can't be put back in a box; unlike uranium or plutonium, they don't need to be mined and refined, and they can be freely copied. Once they are out, they are out. Churchill remarked, in a famous left-handed compliment, that the American people and their leaders "invariably do the right thing, after they have examined every other alternative." In this case, however, we must act more presciently, as to do the right thing only at last may be to lose the chance to do it at all.

As Thoreau said, "We do not ride on the railroad; it rides upon us"; and this is what we must fight, in our time. The question is, indeed, Which is to be master? Will we survive our technologies?

We are being propelled into this new century with no plan, no control, no brakes. Have we already gone too far down the path to alter course? I don't believe so, but we aren't trying yet, and the last chance to assert control—the fail-safe point—is rapidly approaching. We have our first pet robots, as well as commercially available genetic engineering techniques, and our nanoscale techniques are advancing rapidly. While the development of these technologies proceeds through a number of steps, it isn't necessarily the case—as happened in the Manhattan Project and the Trinity test—that the last step in proving a technology is large and hard. The breakthrough to wild self-replication in robotics, genetic engineering, or nanotechnology could come suddenly, reprising the surprise we felt when we learned of the cloning of a mammal.

And yet I believe we do have a strong and solid basis for hope. Our attempts to deal with weapons of mass destruction in the last century provide a shining example of relinquishment for us to consider: the unilateral US abandonment, without preconditions, of the development of biological weapons. This relinquishment stemmed from the realization that while it would take an enormous effort to create these terrible weapons, they could from then on easily be duplicated and fall into the hands of rogue nations or terrorist groups.

The clear conclusion was that we would create additional threats to ourselves by pursuing these weapons, and that we would be more secure if we did not pursue them. We have embodied our relinquishment of biological and chemical weapons in the 1972 Biological Weapons Convention (BWC) and the 1993 Chemical Weapons Convention (CWC).¹²

As for the continuing sizable threat from nuclear weapons, which we have lived with now for more than 50 years, the US Senate's recent rejection of the Comprehensive Test Ban Treaty makes it clear relinquishing nuclear weapons will not be politically easy. But we have a unique opportunity, with the end of the Cold War, to avert a multipolar arms race. Building on the BWC and CWC relinquishments, successful abolition of nuclear weapons could help us build toward a habit of relinquishing dangerous technologies. (Actually, by getting rid of all but 100 nuclear weapons worldwide—roughly the total destructive power of World War II and a considerably easier task—we could eliminate this extinction threat. ¹³)

Verifying relinquishment will be a difficult problem, but not an unsolvable one. We are fortunate to have already done a lot of relevant work in the context of the BWC and other treaties. Our major task will be to apply this to technologies that are naturally much more commercial than military. The substantial need here is for transparency, as difficulty of verification is directly proportional to the difficulty of distinguishing relinquished from legitimate activities.

I frankly believe that the situation in 1945 was simpler than the one we now face: The nuclear technologies were reasonably separable into commercial and military uses, and monitoring was aided by the nature of atomic tests and the ease with which radioactivity could be measured. Research on military applications

could be performed at national laboratories such as Los Alamos, with the results kept secret as long as possible.

The GNR technologies do not divide clearly into commercial and military uses; given their potential in the market, it's hard to imagine pursuing them only in national laboratories. With their widespread commercial pursuit, enforcing relinquishment will require a verification regime similar to that for biological weapons, but on an unprecedented scale. This, inevitably, will raise tensions between our individual privacy and desire for proprietary information, and the need for verification to protect us all. We will undoubtedly encounter strong resistance to this loss of privacy and freedom of action.

Verifying the relinquishment of certain GNR technologies will have to occur in cyberspace as well as at physical facilities. The critical issue will be to make the necessary transparency acceptable in a world of proprietary information, presumably by providing new forms of protection for intellectual property.

Verifying compliance will also require that scientists and engineers adopt a strong code of ethical conduct, resembling the Hippocratic oath, and that they have the courage to whistleblow as necessary, even at high personal cost. This would answer the call—50 years after Hiroshima—by the Nobel laureate Hans Bethe, one of the most senior of the surviving members of the Manhattan Project, that all scientists “cease and desist from work creating, developing, improving, and manufacturing nuclear weapons and other weapons of potential mass destruction.”¹⁴ In the 21st century, this requires vigilance and personal responsibility by those who would work on both NBC and GNR technologies to avoid implementing weapons of mass destruction and knowledge-enabled mass destruction.

Thoreau also said that we will be “rich in proportion to the number of things which we can afford to let alone.” We each seek to be happy, but it would seem worthwhile to question whether we need to take such a high risk of total destruction to gain yet more knowledge and yet more things; common sense says that there is a limit to our material needs—and that certain knowledge is too dangerous and is best forgone.

Neither should we pursue near immortality without considering the costs, without considering the commensurate increase in the risk of extinction. Immortality, while perhaps the original, is certainly not the only possible utopian dream.

I recently had the good fortune to meet the distinguished author and scholar Jacques Attali, whose book *Lignes d'horizons* (*Millennium*, in the English translation) helped inspire the Java and Jini approach to the coming age of pervasive computing, as previously described in this magazine. In his new book *Fraternités*, Attali describes how our dreams of utopia have changed over time:

“At the dawn of societies, men saw their passage on Earth as nothing more than a labyrinth of pain, at the end of which stood a door leading, via their death, to the company of gods and to *Eternity*. With the Hebrews and then the Greeks, some men dared free themselves from theological demands and dream of an ideal City where *Liberty* would flourish. Others, noting the evolution of the market society, understood that the liberty of some would entail the alienation of others, and they sought *Equality*.”

Jacques helped me understand how these three different utopian goals exist in tension in our society today. He goes on to describe a fourth utopia, *Fraternity*, whose foundation is altruism. Fraternity alone associates individual happiness with the happiness of others, affording the promise of self-sustainment.

This crystallized for me my problem with Kurzweil's dream. A technological approach to Eternity—near immortality through robotics—may not be the most desirable utopia, and its pursuit brings clear dangers. Maybe we should rethink our utopian choices.

Where can we look for a new ethical basis to set our course? I have found the ideas in the book *Ethics for the New Millennium*, by the Dalai Lama, to be very helpful. As is perhaps well known but little heeded, the Dalai Lama argues that the most important thing is for us to conduct our lives with love and compassion for others, and that our societies need to develop a stronger notion of universal responsibility and of our interdependency; he proposes a standard of positive ethical conduct for individuals and societies that seems consonant with Attali's Fraternity utopia.

The Dalai Lama further argues that we must understand what it is that makes people happy, and acknowledge the strong evidence that neither material progress nor the pursuit of the power of knowledge is the key—that there are limits to what science and the scientific pursuit alone can do.

Our Western notion of happiness seems to come from the Greeks, who defined it as “the exercise of vital powers along lines of excellence in a life affording them scope.” [15](#)

Clearly, we need to find meaningful challenges and sufficient scope in our lives if we are to be happy in whatever is to come. But I believe we must find alternative outlets for our creative forces, beyond the culture of perpetual economic growth; this growth has largely been a blessing for several hundred years, but it has not brought us unalloyed happiness, and we must now choose between the pursuit of unrestricted and undirected growth through science and technology and the clear accompanying dangers.

It is now more than a year since my first encounter with Ray Kurzweil and John Searle. I see around me cause for hope in the voices for caution and relinquishment and in those people I have discovered who are as concerned as I am about our current predicament. I feel, too, a deepened sense of personal responsibility—not for the work I have already done, but for the work that I might yet do, at the confluence of the sciences.

But many other people who know about the dangers still seem strangely silent. When pressed, they trot out the “this is nothing new” riposte—as if awareness of what could happen is response enough. They tell me, There are universities filled with bioethicists who study this stuff all day long. They say, All this has been written about before, and by experts. They complain, Your worries and your arguments are already old hat.

I don't know where these people hide their fear. As an architect of complex systems I enter this arena as a generalist. But should this diminish my concerns? I am aware of how much has been written about, talked about, and lectured about so authoritatively. But does this mean it has reached people? Does this mean we can discount the dangers before us?

Knowing is not a rationale for not acting. Can we doubt that knowledge has become a weapon we wield against ourselves?

The experiences of the atomic scientists clearly show the need to take personal responsibility, the danger that things will move too fast, and the way in which a process can take on a life of its own. We can, as they did, create insurmountable problems in almost no time flat. We must do more thinking up front if we are not to be similarly surprised and shocked by the consequences of our inventions.

My continuing professional work is on improving the reliability of software. Software is a tool, and as a toolbuilder I must struggle with the uses to which the tools I make are put. I have always believed that making software more reliable, given its many uses, will make the world a safer and better place; if I were to come to believe the opposite, then I would be morally obligated to stop this work. I can now imagine such a day may come.

This all leaves me not angry but at least a bit melancholic. Henceforth, for me, progress will be somewhat bittersweet.

Do you remember the beautiful penultimate scene in Manhattan where Woody Allen is lying on his couch and talking into a tape recorder? He is writing a short story about people who are creating unnecessary, neurotic problems for themselves, because it keeps them from dealing with more unsolvable, terrifying problems about the universe.

He leads himself to the question, “Why is life worth living?” and to consider what makes it worthwhile for him: Groucho Marx, Willie Mays, the second movement of the Jupiter Symphony, Louis Armstrong’s recording of “Potato Head Blues,” Swedish movies, Flaubert’s *Sentimental Education*, Marlon Brando, Frank Sinatra, the apples and pears by Cézanne, the crabs at Sam Wo’s, and, finally, the showstopper: his love Tracy’s face.

Each of us has our precious things, and as we care for them we locate the essence of our humanity. In the end, it is because of our great capacity for caring that I remain optimistic we will confront the dangerous issues now before us.

My immediate hope is to participate in a much larger discussion of the issues raised here, with people from many different backgrounds, in settings not predisposed to fear or favor technology for its own sake.

As a start, I have twice raised many of these issues at events sponsored by the Aspen Institute and have separately proposed that the American Academy of Arts and Sciences take them up as an extension of its work with the Pugwash Conferences. (These have been held since 1957 to discuss arms control, especially of nuclear weapons, and to formulate workable policies.)

It’s unfortunate that the Pugwash meetings started only well after the nuclear genie was out of the bottle—roughly 15 years too late. We are also getting a belated start on seriously addressing the issues around 21st - century technologies—the prevention of knowledge-enabled mass destruction—and further delay seems unacceptable.

So I’m still searching; there are many more things to learn. Whether we are to succeed or fail, to survive or fall victim to these technologies, is not yet decided. I’m up late again—it’s almost 6 am. I’m trying to imagine some better answers, to break the spell and free them from the stone.

1 The passage Kurzweil quotes is from Kaczynski's Unabomber Manifesto, which was published jointly, under duress, by *The New York Times* and *The Washington Post* to attempt to bring his campaign of terror to an end. I agree with David Gelernter, who said about their decision:

"It was a tough call for the newspapers. To say yes would be giving in to terrorism, and for all they knew he was lying anyway. On the other hand, to say yes might stop the killing. There was also a chance that someone would read the tract and get a hunch about the author; and that is exactly what happened. The suspect's brother read it, and it rang a bell.

"I would have told them not to publish. I'm glad they didn't ask me. I guess."

(*Drawing Life: Surviving the Unabomber*. Free Press, 1997: 120.)

2 Garrett, Laurie. *The Coming Plague: Newly Emerging Diseases in a World Out of Balance*. Penguin, 1994: 47-52, 414, 419, 452.

3 Isaac Asimov described what became the most famous view of ethical rules for robot behavior in his book *I, Robot* in 1950, in his Three Laws of Robotics: 1. A robot may not injure a human being, or, through inaction, allow a human being to come to harm. 2. A robot must obey the orders given it by human beings, except where such orders would conflict with the First Law. 3. A robot must protect its own existence, as long as such protection does not conflict with the First or Second Law.

4 Michelangelo wrote a sonnet that begins:

*Non ha l' ottimo artista alcun concetto
Ch' un marmo solo in sè non circonscriva
Col suo soverchio; e solo a quello arriva
La man che ubbidisce all' intelletto.*

Stone translates this as:

*The best of artists hath no thought to show
which the rough stone in its superfluous shell
doth not include; to break the marble spell
is all the hand that serves the brain can do.*

Stone describes the process: "He was not working from his drawings or clay models; they had all been put away. He was carving from the images in his mind. His eyes and hands knew where every line, curve, mass must emerge, and at what depth in the heart of the stone to create the low relief."

(*The Agony and the Ecstasy*. Doubleday, 1961: 6, 144.)

5 First Foresight Conference on Nanotechnology in October 1989, a talk titled "The Future of Computation." Published in Crandall, B. C. and James Lewis, editors. *Nanotechnology: Research and Perspectives*. MIT Press, 1992: 269. See also www.foresight.org/Conferences/MNT01/Nano1.html.

6 In his 1963 novel *Cat's Cradle*, Kurt Vonnegut imagined a gray-goo-like accident where a form of ice called ice-nine, which becomes solid at a much higher temperature, freezes the oceans.

7 Kauffman, Stuart. "Self-replication: Even Peptides Do It." *Nature*, 382, August 8, 1996: 496. See www.santafe.edu/sfi/People/kauffman/sak-peptides.html.

8 Else, Jon. *The Day After Trinity: J. Robert Oppenheimer and The Atomic Bomb* (available at www.pyramiddirect.com).

9 This estimate is in Leslie's book *The End of the World: The Science and Ethics of Human Extinction*, where he notes that the probability of extinction is substantially higher if we accept Brandon Carter's Doomsday Argument, which is, briefly, that "we ought to have some reluctance to believe that we are very exceptionally early, for instance in the earliest 0.001 percent, among all humans who will ever have lived. This would be some reason for thinking that humankind will not survive for many more centuries, let alone colonize the galaxy. Carter's doomsday argument doesn't generate any risk estimates just by itself. It is an argument for *revising* the estimates which we generate when we consider various possible dangers." (Routledge, 1996: 1, 3, 145.)

10 Clarke, Arthur C. "Presidents, Experts, and Asteroids." *Science*, June 5, 1998. Reprinted as "Science and Society" in *Greetings, Carbon-Based Biped!* *Collected Essays, 1934-1998*. St. Martin's Press, 1999: 526.

11 And, as David Forrest suggests in his paper "Regulating Nanotechnology Development," available at www.foresight.org/NanoRev/Forrest1989.html, "If we used strict liability as an alternative to regulation it would be impossible for any developer to internalize the cost of the risk (destruction of the biosphere), so theoretically the activity of developing nanotechnology should never be undertaken." Forrest's analysis leaves us with only government regulation to protect us - not a comforting thought.

12 Meselson, Matthew. "The Problem of Biological Weapons." Presentation to the 1,818th Stated Meeting of the American Academy of Arts and Sciences, January 13, 1999. (minerva.amacad.org/archive/bulletin4.htm)

13 Doty, Paul. "The Forgotten Menace: Nuclear Weapons Stockpiles Still Represent the Biggest Threat to Civilization." *Nature*, 402, December 9, 1999: 583.

14 See also Hans Bethe's 1997 letter to President Clinton, at www.fas.org/bethecr.htm.

15 Hamilton, Edith. *The Greek Way*. W. W. Norton & Co., 1942: 35.

Bill Joy, cofounder and Chief Scientist of Sun Microsystems, was cochair of the presidential commission on the future of IT research, and is coauthor of The Java Language Specification. His work on the Jini pervasive computing technology was featured in Wired 6.08.

[Copyright](#) © 1993-2004 The Condé Nast Publications Inc. All rights reserved.

[Copyright](#) © 1994-2003 Wired Digital, Inc. All rights reserved.

CHAPTER 5

THE TECHNOLOGICAL SUBVERSION OF ENVIRONMENTAL ETHICS

The unexamined life is not worth living.

—Socrates

We don't want to exaggerate our degrees of freedom. But they are not zero.

—Charles Taylor¹

How do we get beyond this other story of our age, this regressive story of technological impoverishment? We need to clarify and deepen our understanding of this fundamental choice our culture is facing. Things will play a key role, but we also need to understand technology differently and better than we have so far. To feel uncomfortable amidst all our comfort is not enough. We want to know what is the point, of getting everything under control. Humans are not merely puppets in the unfolding of the age. Coming to terms with why we do what we do has often been cathartic in the maturation of humans. Understanding how technological culture hangs together will enable us to foresee its ultimate shape, intelligently evaluate it, and find its pivots of reform.

I

The Underlying Ethic of Technology

Some have argued that we live in an invisible iron cage.² Indeed, technological forces are shaping people's lives in ways that they have little or no control over, especially if the basic framework of technology goes unchallenged; but, as Charles Taylor points out, the conquest of nature had a

benevolent point to it. It was to serve humanity. So, he finds that, along with other forces, there are moral forces at work here shaping our lives. We live neither in an iron cage nor in an arena of unconstrained choice; we inhabit a possibility space where some moral choices are being made.³ We did see earlier that there is a kind of ethical appeal to not letting our resources go to waste. So what ethical forces might be called upon to reform technology in a deep way? How should we understand the basic choice we face? For developing what I call the vision and underlying ethic of technology, I will draw heavily upon Albert Borgmann's theory of technology, the best account of the character of the technological culture we have so far. Then we will use this vision of technology to show that the concerns of environmental ethics and people's better concerns for nature generally will be subverted by technology unless we as a culture come to grips with the irony of this vision and begin to make a fundamentally different choice, that is, choose things over consumption.

Making the Appeal of Technology Intelligible: The Promise of Technology

Neither Heidegger nor Thoreau makes clear what it is about technology that is attractive to people. Claiming that we delight in the exercise of power seems correct enough when we think of the enormous amount of power we wield with technology, yet this view does not address our more intelligible motives and, therefore, does not really address many of the proponents of technology without trivializing their concerns. In one way or another, most of us, if not all, see technology as good. What is at the heart of our petty homocentrism? What good is technology?

Typically people articulate what good technology is when they say that something is better or improved and demonstrate that "that's progress." Advertisements are continually pointing out what is better about the product advertised, even if the chief "advantage" is two for the price of one. Although they may well dupe us, these advertisements normally appeal to standards that at least on a deep and general level are already in place and widely shared in consumer culture. We hear everywhere around us, not just in advertisements, what better is. "It means less work." "It's more comfortable." "It's convenient." "It's healthier." "It's faster and more productive." "It's less of a hassle." "Sleeker looks better." "It's lighter." "It doesn't get in your way." "You don't have to wait on anyone else." "It's exciting." When we see the very latest devices, often our expressions are on the order of "Wow!" or "That's great!" or "Look at that, would you?" So, at deeper levels, there seems to be a good deal of likemindedness about what constitutes better in our culture.

Another approach is to consider what people think of as clear exam-

ples of progress. Television today is far different from what it was in the past. In the early 50s, one was lucky to own a television. Reception was poor, the picture rough and in black and white; the screen was small, the set large; the number of programs was very limited. In addition to other obvious improvements, now the sets come on instantly, are controlled from the couch, can be found in all sizes and nearly everywhere, and have access to a vast number and variety of programs, especially with video cassette recorders. Even if they are not willing to pay the price for all of them, most count these changes as improvements, and rarely do we find people watching a black and white set any longer. What are the standards which make these changes count as improvements?

Television as a clear example of technology will play a key role in our understanding of the nature of the fundamental choice we face, but Borgmann uses another paradigmatic example of technology, the central heating system, to disclose most of these standards of technology.⁴ We can easily trace the development of central heating systems back to the wood-burning stove or the hearth. The chief advantages of the heating system over these latter two are various. Central heating is *easier*: We do not have to gather, stack, chop or carry the wood. An automatic thermostat means that we do not have to trouble ourselves in the morning or evening with setting a thermostat. Central heating is more *instantaneous*. We do not have to wait for the house to warm up. It's *ubiquitous*. Warmth is provided to each corner of the room, to every room, and everywhere equally well. Finally, a central heating system is *safer* than a hearth. My grandmother was born in a newly built chicken coop because three weeks earlier her family's house burnt down from a chimney fire. So the standards by which people judge central heating to be better than a woodburning stove are ease, instantaneity, ubiquity and safety or some combination of these, for example, convenience. These four "technological standards" can be collected under the more general notion of technological availability. To be more *available* is to be an improvement, then, in terms of one or more of the four above standards.

Why does it seem to people that this availability is good? From one perspective, this availability relieves people of burdens: less effort, less time and less learning skills are required. Available anytime and anywhere, they are disburdened of the constraints of time and place. They are disburdened of having to take risks. Historically, modern technology was envisioned as enabling people not just to subjugate nature, but to do so for the purpose of freeing humanity from misery and toil. To be relieved of these burdens then, fulfills this vision of technology. *To the degree people personally share this vision*, they will also see its concrete manifestations, such as central heating, as unquestionably good. Compared to older versions, the

latest portable computers exemplify this relief from burdens and are attractive to many for this reason.

By overcoming nature, technology would, as some in the seventeenth century foresaw, not only relieve humans of burdens, but it would make available to them—easier, safer, quicker, and more ubiquitous—all the goods of the Earth. So, technological availability negatively disburdens people of misery and toil, and positively enriches their lives, makes them happy, it seems. So seen, technology has an attractive glow about it.

Technology promises to bring the forces of nature and culture under control, to liberate us from misery and toil, and to enrich our lives. . . . [More accurately], implied in the technological mode of taking up with the world there is a promise that this approach to reality will, by way of the domination of nature, yield liberation and enrichment.⁵

Borgmann calls this “the promise of technology.”

Clearly those below the middle-class of advanced industrialized countries and those outside those countries do not derive the benefits of technology, although many do feel the pull of its promise. As I outlined in chapter 1 and will argue later here, the claims of social justice will not likely be met until the more privileged ones, the middle and higher classes of these industrialized countries, come to terms with the questionable character of technology's promise. So, the critique of technology I am developing here does not apply to those in poverty. It applies only to those who have too much.

For these latter, technology has made good on its promise in important ways. My grandmother's father died from what she believes was pneumonia when she was eleven, leaving her and her younger sister to perform heroic feats to save the cattle from starvation in the rough times of an extended winter that followed. Often hitching the team up before dawn and returning hours after dark, especially in winter, her family took an entire day to get to and from town sixteen miles away. For the privileged, then, many past hardships have now been conquered. Although we may have legitimate concerns about whether there is too much medical technology, no one could reasonably refuse every advance of modern medical technology. The weather will never be brought under control, but, via comfortable structures, nature's heat and cold, rain and snow are controlled as well as darkness and drought. Toilsome labor is largely eliminated within the culture of technology.

A reasonable person may reject motorcycles in favor of horses to do ranch work, but that person still rides to town on paved roads in a car, has parts shipped by air, reads a newspaper and books, transacts business over

the phone, and owns at least a radio. No thoughtful person will want to turn her back on technology entirely. Thus, technology, by conquering nature, has relieved humans of severe burdens. Today we are still working to overcome those, such as cancer and AIDS, that remain. So, if technology does not saddle us in the long run with more than it has relieved us from, it will have made good on this aspect of what at first seemed and still does seem promising about it. It could turn out that ozone depletion, global warming, ecosystem destruction, the population explosion, polluted land, air, streams and oceans, and human and mechanical errors will impose burdens far greater than those we were relieved from in the first place. To meet these problems certainly calls for a reform of present practices. We read or hear of these calls for reform nearly every day. More common critiques of technology, such as David Ehrenfeld's *The Arrogance of Humanism*, attempt to show that technology will fail by its own standards, bringing disaster upon us.⁶

Much as reform in these areas is needed and much as these pessimistic critiques deserve thoughtful consideration, the present work will turn to a uniquely different task. It grants and, in fact, seeks to have the reader appreciate, the genuine success of modern technology. Technology has relieved, and technology will, I assume for the purposes at hand, continue to relieve, humans of many hardships of the human condition.

So what is wrong with technology for those within the realm of its benefits? Underlying these standards of availability is really a vision of a good life that is free and prosperous. What is at the bottom of concern with technological availability is an aspiration for freedom and happiness. Most people, at least in the Western tradition, are concerned with liberty and prosperity. For Aristotle only the Greek free man was able to have sufficient time and sufficient wherewithal to develop the moral and intellectual virtues he thought to be required for happiness or eudaemonia. The Hebrew people's understanding of the covenant centered on an idea of prosperity. Jesus preached of a free and abundant life. The Enlightenment, as we see its results in “the pursuit of happiness” in our Declaration of Independence, is fully within this tradition. But to find agreement at this high level of abstraction is not to see that the crucial differences lie with the particular versions of freedom and prosperity. For Socrates living well had to do with human excellence and living a just life, not with materialism. The blessed life and the abundant life of the Hebrews and of Jesus was not commodity happiness. So, too, we must look carefully at the particular idea of freedom and prosperity governing people's attraction to technology, for only at this level of particularity will its misleading and harmful features begin to show. In other words, one can criticize the trivialized forms of freedom and prosperity on which the technological society is centered

without, at the same time, criticizing freedom and prosperity more generally as a vision of the good life. Quite the contrary, we can call technology into question even more sharply by showing that technology fails to provide the free, prosperous, and good life we want in our waking moments. Technological society offers a flattened vision of freedom and prosperity. The more disburdened, the better off I am according to this vision. So, the technological idea of freedom is really one of disburdenment. What about prosperity? Cellular phones are currently a status symbol. These devices which disburden people of the constraints of place are taken to be a sign of affluence because, generally, only the more prosperous have them. So, in part, to be prosperous is to have the latest, most refined device. A sign of affluence, too, is to be able to go to an undiscovered exotic place, have the most channels and compact disks, own specially designed clothing, own what no one else has yet. Thus, in part, to be prosperous is to own the most varied, the widest assortment of commodities. Finally, when people buy a product on sale they get both the commodity they purchased and still have money left over. Why is that attractive? Because they can buy something else with the money saved. They are better off that way, they think, because they get more items for the money. Thus people pursue prosperity through the standards of owning the *most numerous, widest variety, and the very latest (most refined)* commodities. The powers that be in the technological society own and control the most of these items. Such is the picture of the good life envied by those keeping up with the Joneses. Our culture's vision of the good life is the goods life.

Does this vision really deliver a good life? If we say no merely because it differs from the blessed life according to Abraham, Moses, and the prophets, or from the Greeks' eudaemonia, our analysis would be dogmatic and presumptuous. Technology must be thought through; it will not be met by simply reacting against it. So, if we answer "no," as I will, then we must be able to provide good reasons.

The Technological Means to Freedom and Happiness: The Device

The ironic consequences of this vision of freedom and prosperity can be drawn out through a careful analysis of the peculiar way technology transforms or, more specifically, dominates nature and culture. Technology does not dominate these in the traditional manner of lordship it over them; rather, as Albert Borgmann shows, technology follows a pattern, unique to the modern era, in the way it gets everything under control. We can expose this pattern by examining instances of it.

The central heating system dominates warmth; it brings warmth un-

der control in ways that the wood-burning stoves do not. To show its unique form of domination, Borgmann distinguishes between "things" and "devices." A thing in his sense

is inseparable from its context, namely its world, and from our commerce with the thing and its world, namely, engagement. The experience of a thing is always and also a bodily and social engagement with the thing's world . . . Thus a stove used to furnish more than mere warmth. It was a *focus*, a hearth, a place that gathered the work and leisure of a family and gave the house a center. Its coldness marked the morning, and the spreading of its warmth marked the beginning of the day. It assigned to various family members tasks that defined their place in the household . . . It provided the entire family a regular and bodily engagement with the rhythm of the seasons that was woven together with the threat of cold and the solace of warmth, the smell of wood smoke, the exertion of sawing and carrying, the teaching of skills, and the fidelity to daily tasks . . . Physical engagement is not simply physical contact but the experience of the world through the manifold sensibility of the body. That sensibility is sharpened and strengthened in skill. Skill is intensive and refined world engagement.⁷

Here, in his retrieval of the thing's world and our engagement with the thing, Borgmann has been influenced by Heidegger's fourfold account of things. Obviously, Earth and sky are woven together with mortals. He points out that in Roman times the hearth was the abode of household gods, though he does not make much of it.⁸ Borgmann's account goes beyond Heidegger in emphasizing social and bodily engagement to a degree to which Heidegger seems insensitive. He also steps beyond Heidegger by highlighting the way things focus practices.⁹ Practices call for skills and the development of character; the diversity of different characters is joined to each other through participating in a world of practices. In our terms developed earlier, the hearth is the correlational coexistent thing which establishes the world of the household and, correlatively, calls forth its members and calls on their deeper capacities.

Today the hearth, if it exists at all, is no longer the central location in the house although the mantel still remains a place of honor. What has replaced the thing is the "device." The device (the central heating system) provides a commodity, one element of the original thing (warmth alone) and disburdens people of all the elements that compose the world and engaging character of the thing. This world of the thing, its ties to the natural and cultural world and our engagement with that many-dimensional world on bodily, cerebral and social levels, is taken over by the *machinery* (the central heating plant itself) of the device.

The machinery makes no demands on our skill, strength, or attention, and it is less demanding the less it makes its presence felt. In the progress of technology, the machinery of the device has therefore a tendency to become concealed or to shrink. Of all the physical properties of a device, those alone are crucial and prominent which constitute the commodity that the device procures.¹⁰

To make the commodity even more technologically available, the machinery varies radically in the history of technology (wood or coal or oil or electricity or gas). Owing to this radical variability and to this concealment, the machinery becomes necessarily *unfamiliar*. I probably do not know by what means the water is heated in a building. But the device is not just machinery or even most importantly machinery. The device makes available a commodity—warmth. Warmth is what the central heating system is for. Just the opposite of the machinery, the commodity tends to *expand* (become ubiquitous in the house), to remain relatively *fixed* as the means change (from coal to electricity) and to be *familiar*. It follows that—unlike with things—there is a wide division between what a device provides, the commodity, and how it provides this commodity, the machinery. Hence, and this is Borgmann's central insight we saw illustrated earlier with second homes, devices *split* means and ends into mere means and mere ends.

Even though these claims that a thing makes on people are not always experienced as burdensome (as we see from the above account), this very world of the thing and the engagement it calls for can be felt at times as a burden or hassle. The technological device and its refinement *disburdens* people of all these problems by expanding the commodity, so that the world of the thing no longer determines when, in what way, and where it is available. Thus, it disburdens them of the claims that call for engagement. In short, the technological device disburdens people of the thing's world and its claims upon them. The device is considered the more refined the more it lifts these burdens from them. The ideal device is one where, from an experiential standpoint, a commodity can be enjoyed unencumbered by means. A reliable self-regulating central heating system whose maintenance and energy bill are taken care of by a management agency can be taken as a paradigmatic example.

The peculiar way technology dominates things is not limited, of course, to the central heating system. Considering how household technologies have changed, Witold Rybczynski in *Home: A Short History of An Idea* writes.

The evolution of domestic technology . . . demonstrates that the history of physical amenities can be divided into two major phases: all the years

leading up to 1890, and the three following decades. If this sounds outlandish, it is worth reminding ourselves that all the "modern" devices that contribute to our domestic comfort—central heating, indoor plumbing, running hot and cold water, electric light and power and elevators—were unavailable before 1890, and were well known by 1920. We live, like it or not, on the far side of a great technological divide. As John Lukacs reminds us, although the home of 1930 would be familiar to us, it would have been unrecognizable to the citizen of 1885.¹¹

Just as with household technologies, so too with other features of our surroundings and our cultural and natural environment generally. This thing-to-device example is representative of the pattern of the technological transformations of the Earth. Generally then, this transformation is one in which:

Devices . . . dissolve the coherent and engaging character of the pre-technological world of things. In a device, the relatedness of the world is replaced by a machinery, but the machinery is concealed, and the commodities, which are made available by a device, are enjoyed without the encumbrance of or the engagement with a context [that is, the world of the thing].¹²

Borgmann calls this pattern the *device paradigm*. At times I will call it the separation pattern of technology. (See diagram on the following page.)

In chapter 4, we characterized our age as one in which we reduce everything to resources that we want to control. Now we can see that the device pattern is used to get control of these resources. The purpose of the device is to supply people with unencumbered commodities. So now we can develop this picture of our age further. The fuller vision is one in which everything gets reduced to resources, machinery and commodities.

Ironic Consequences

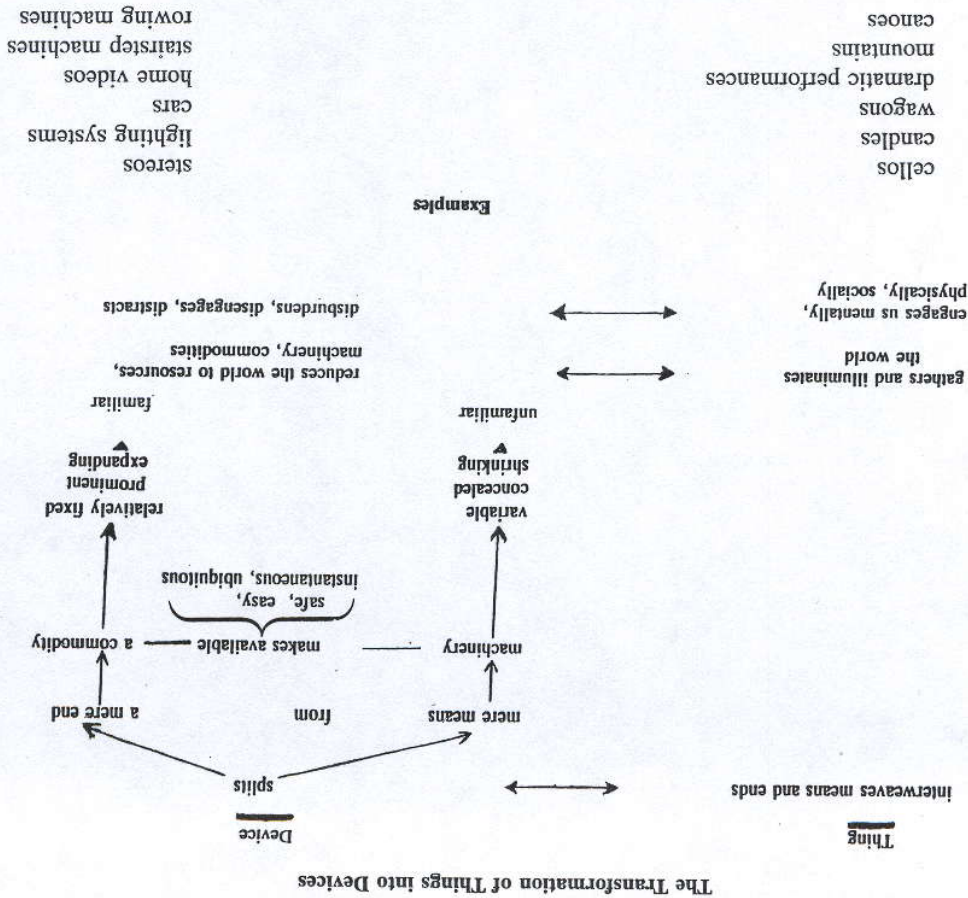
So far we have developed a theory by which we can interpret what has taken and is taking place with regard to the technological transformation in our time. Using this theory we can pass from technological object to technological object, seeing how they more or less fit the pattern. The illustration of the pattern does not commit us as yet to an evaluation of the good or bad of what has taken place. Now we are in a position to begin that task. What are the consequences of this change from things to devices?

Don Ihde finds that technologies transform experience in a "non-neutral" manner. A tool always amplifies in some way certain aspects of nor-

mal embodied experience while simultaneously reducing other aspects.¹³ A dentist's probe shows the hardness and cavities in a tooth to a degree fingers miss, while the wetness and warmth of the tooth felt by the fingers go undetected by probe. This change Idhe finds is non-neutral because the amplified features are heightened, drawing our attention, while the reduced features tend to go unheeded and are overlooked and often forgotten. Asked what a hearth is for, we find it logical, after having experienced central heating, to answer that it supplies heat, ignoring or not even seeing its other aspects. Extending Idhe's insight makes it more intelligible why we become fascinated with commodities, heedless of what has been reduced. Yet pointing out that this change is non-neutral is not enough. We now want to comprehend what exactly has been hidden from us. We need a language which articulates what is overlooked and forgotten, for then we can see in what ways this change is non-neutral. Our language of things retrieves and focuses this loss. It reveals the general pattern that things are transformed into devices, detaching people from things, their world and each other.

Idhe further argues that these lost features only tend to recede, thus, implying that they are retrievable. With certain kinds of instruments (not all technological objects are devices, splitting means and ends), this is true. We can easily retrieve the features missed by the dentist's probe. With devices, however, these features do not just tend to withdraw, so that a change of attitude, perception or act of will could retrieve them. Notice that mere warmth, no matter how expanded the commodity has become, is not a substitute for the thing of the wood-burning stove. Mere warmth could not be the essence of a household; it does not warrant that kind of attention or care, of heeding. Indeed the source itself is concealed and the warmth is suffused throughout the house so that it fails to provide a focus. So warmth is no substitute for the thing because it lacks a world with which to become engaged. More than this, because it is impossible to recover in the mere warmth of the central heating plant the full-bodied experience of the hearth, the machinery of devices ineluctably withdraws the world from people. A device is necessarily unfamiliar in the ways that the context of the thing was familiar. Thus the transformation of the thing into a device does not merely tend to obscure possibilities of experience, but its very material structure makes the rich experience of the thing impossible.

Another way of putting this is that devices allow the possibility of only slim points of contact with "narrowly defined aspects of what used to be things of depth."¹⁴ Devices force people to take them as commodity bearers; they leave them no choice. So our way of taking commodities is not a psychological matter, but a real matter. Technology is not only a way



of seeing (and for this reason characterizing technology as a vision is perhaps misleading), it is more importantly a way of *shaping*. The very material structure of a device is such that it can be experienced only as holding up a commodity calling for consumption and nothing more.

The implication of this change of shape is alienation. What seemed promising at the outset—relieving people of burdens—leads ironically to disengagement, diversion, distraction and loneliness. In short, we become not-at-home in the universe. But clearly, simply finding ourselves free from the exclusive use of candles and outhouses, does place us in this alienated position. So how can such positive events as electric lighting and indoor plumbing lead to these ironic results?

To be relieved of famine, cold, darkness, confinement and other genuine adversities of the human condition was an intelligible and urgent demand for the early phase of modern technology. For the middle-classes of advanced industrialized countries, most, not all, of these kinds of challenges have been met for some time. At the stage of mature technology, the challenges can be quite frivolous. Food processors, electric pencil sharpeners, prepared fishing leaders, automatic cameras, electric knives, and some pain relievers are typical. The basic question here is: Do we need to be relieved of every last and least burden? Aren't some of these burdens actually good in senses that touch our very humanity? When people reflect on these questions they may answer them differently, but when they act, they tend to act in agreement with a vision that seeks to bring everything under control. Ironically, in the wake of such technological success, in the wake of the initial excitement over owning the latest item, the item falls back into the ordinary every day and they become bored. Being bored, they become disengaged and alienated from what may have been a vital practice, such as preparing meals or gardening or photography. Accordingly, they seek diversion. Thus, ironic consequences follow from the disburdenment of every hassle, problem or felt demand. If we pursue disburdenment in this unchecked and unreflective manner, as people are doing in the stage of mature technology, then these are the results we should expect.

However, it may seem as though we have been just too nostalgic. The disburdenment devices yield "frees us up for other things" as people commonly say. Yet this perspective makes us think that technology is mostly about freedom, as Charles Taylor thinks, when the promise and vision of technology are mostly a promise of happiness. The most unique and devastating critique of technology is not centered on technological freedom, but on the fact that technology *fails most where it succeeds most* at procuring happiness, at procuring the good things of life. As a culture, we think not only that we can use technology to liberate us, but also that we can use it to fill that new possibility space with technologically available goods. In

short, what people are freed up for are not other things, but *more commodities*. Then too when people imagine what they are doing as they throw food into the microwave as freeing themselves up for other and more important things, they ignore how pervasive the technological order is. The totality of technological devices is far more consequential than any particular device. The former point can be advanced best by developing the latter first.

Extensively yet unobtrusively this technological way of taking up with the world pervades and informs what people think, say and do. We need an account of technology as *correlational environment*. Organizations, institutions, the ways nature and culture are arranged and accessible all become modeled on the device. As people make more decisions for consumption against engagement, our average, everyday world is stamped more deeply with the pattern of the device. In other words, devices do not simply liberate people from some things and free them for other and better things. We are surrounded. The things enabling correlational coexistence have nearly disappeared. As the totality of our daily environment changes from an environment of things to an environment of devices, from an environment making demands on people to an environment that is more at their finger tips, this change necessarily entails heedlessness and evokes an attitude of cultural petty homocentrism. So it is important to consider not just the appropriateness of this or that device in a particular context, but to consider what the consequences are of the totality of these devices and people's *typical* use of them. We would expect the consumptive ways of life in such surroundings to be disburdened and disengaged.

So what are people finally freed up for? How do they attempt to use technological means to positively enrich their lives? Typically people use devices to procure entertainment commodities.¹⁵ Hence our culture treats tradition, culture and nature as resources to be mined. Just as ubiquitously available warmth is not a substitute for the hearth, entertainment commodities are at best insubstantial aspects of the original things. Because they use devices here to procure the delights that matter, final things, these entertainment commodities can be thought of as final commodities. In this respect, to consume a final commodity is no different from consuming an instrumental one that disburdens us of a chore.

The Ironic Consequences of Final Consumption

Television is a clear example of a final commodity. Its refinements from the first sets to those of today fit the same pattern as the refinements in central heating and the refinements of devices in accordance with Borgmann's device paradigm in general. So television is an instance of the

vision of domination, liberation, and enrichment. It does not make demands on people and is a window of the world, making all the goods of the Earth available, technologically available, to them in their living rooms.

Understandably, television has tremendous appeal to us as a culture. It's where technology comes home to people. The amount of time they spend watching it indicates its power.

The A.C. Nielsen Company (1989) currently estimates that people in the United States view upwards of 4 hours of television each day. Given the likelihood that such estimates are inflated, let us assume a more conservative estimate of 2 1/2 hours of television viewing per day over the period of a lifetime. Even at this more conservative rate, a typical American would spend more than 7 full years watching television out of the approximately 47 waking years each of us lives by age 70—this assuming an average of 8 hours of sleep per day. Such a figure is even more striking when we consider that Americans have about 5 1/2 hours a day of free time, or approximately 16 years available for leisure of the same 47-year span. From this point of view and based on a conservative estimate, Americans are spending nearly half of their available free time watching television.¹⁶

Since it is the most popular way people enjoy final consumption, it is worthwhile to examine in detail the *experience* of this form of consumption as we develop the ironic consequences of final consumption. *Television and the Quality of Life* by Robert Kubey and Mihaly Csikszentmihalyi (cited above) does just that, examining systematically the reported experience of television in contrast to the reported experience of other activities people spend time on in their daily life. So we need to look closely at their findings with a view to showing how technology in the form TV does not fulfill its promise.

Kubey and Csikszentmihalyi find that television is inexpensive and is easily and quickly available for those who have time for it.¹⁷ It helps people to relax and, at times, may help some to retreat before gathering themselves to face a difficulty.¹⁸ People can watch it for prolonged periods without wearing themselves out physically. It tends to bring families together and family members normally do talk with each other while watching. They also feel better than when watching alone.¹⁹ It is used for news programs, for nature shows, and to present dramas such as *Death of a Salesman*. It could be used to present lectures in chemistry and Plato.²⁰ It helps connect us with our culture and some of its common stories. On the other hand, "viewing is almost always mildly rewarding in that it provides relaxation, distraction, and escape with minimal effort."²¹ It gives people something to do with their time and most report they do want to watch it.²²

With so much of our leisure time taken up with television and with so many benefits, one might mistakenly conclude that people choose to watch because it is better than anything else they could be doing. Yet the actual cumulative benefits they receive from television are rather low and often negative.

Ironically, the reported experience of people viewing television often turns out to be one of disappointment.²³ Not only are chemistry courses not aired because few would watch them, not only do most people gravitate toward watching movies with light and escapist content rather than challenging dramas, but, just as important, half the people who watch television do not use television guides to help them decide what to watch.²⁴ The stories viewers share, then, are not those shown on public television. The shows tend to support existing beliefs. As Stu Silverman told Kubey, "Television reassures us, it's 'nice,' it doesn't offend or challenge an audience. It is designed to do just the opposite of art, to reassure rather than excite. That often is what people want."²⁵ Although television does help people to relax, it does not do so any more than other activities such as reading. Moreover, it helps people relax only while watching it and not later as sports and other activities do.²⁶ Although this study found that television is not a completely passive activity, it is comparatively so. It is not usually challenging, requires little mental alertness, and is reported to exact fewer skills than eating. Only idling was reported to be more passive.²⁷ Unlike activities that gather and restore a person, a "passive spillover effect" tended to follow watching television, making people feel duller, more passive, and less able to concentrate.²⁸ Families for which television provides a center also experience this spillover effect carried over into other family activities.²⁹ Finally, the positive benefits one receives from television tend to be enjoyed less the more one watches.³⁰ Heavy viewers are not made happy by watching it; they generally feel worse than light viewers both before and after viewing.³¹ Even light viewers do not report themselves to be any happier than average while watching.³²

Aristotle found amusement, like sleep, to be therapeutic as long as life is oriented around exertion. So, too, Kubey and Csikszentmihalyi find that those who stand to benefit most from television use and need it least.³³ More often television is used to disburden people of problems in ways that do not go to the roots of the problems, are only marginally effective, and, hence, are entirely inappropriate.³⁴ People disburden themselves of the problem of leisure time that their time-saving and labor-saving devices have created by killing time watching television. They disburden themselves of the problem of loneliness when devices leave them isolated by turning on a device, the television set. Heavy use is higher among singles.³⁵ Such an answer to loneliness is only a diversion from genuine forms of

social engagement. On the other hand, television is often used as a way for family members, usually fathers, to avoid talking with other family members and avoid dealing with family problems.³⁶ Television resolves the problem of independently ordering ones life, of giving shape to the day. It takes care of boredom.³⁷ Heavy television viewing is likely driven by a wish to escape, to be disburdened of bad days and bad moods, of personal problems and of alienation from self.³⁸ Diverting ones attention, it tends to mask the deeper and more real problems a person is having and, hence, leaves these problems unintelligently resolved. Does it meet the task of leading a more rewarding and meaningful life? No. "Happiness is a more complex state than relaxation. It requires a more elusive set of conditions, and is therefore more difficult to obtain."³⁹ Television seems to "encourage a false sense of well-being in some people," distracting them from and becoming an obstacle to the hard work it takes to realize ones potentials.⁴⁰

More indirectly, we can ask what people are missing when they watch television. When viewers are not pleased with the amount of time they watch television, the entire reason is not only that it is a low-grade activity, one that many think best fits the phrase "Am I lazy!"⁴¹ Part of the reason, too, has to do with what television is displacing. Many report that they feel as though they should have been doing something else. College educated viewers felt this way more often than other people because "they should have been doing something more productive."⁴² Television rearranges life through decreasing the amount of time spent involved with other activities. It at times provides a center for the household, but such a center seems flimsy at best, especially in comparison with other potential centers or centers of the past. In another context, Kubey and Csikszentmihalyi speak of these kinds of centers of life.

When people are asked what they enjoy most, and enough time is left for a genuine answer to emerge, we often find that the most enjoyable things involve doing something, and usually something rather complex and demanding. Rarely does watching television get mentioned, or any other passive or consummatory activity . . . The first reflex for many people is to say that one most enjoys going on vacations, going to movies or restaurants—the typical "leisure" responses in our culture. But as people think more deeply about their real feelings they will mention enjoyable times with their families, and then there is often a point when their faces light up and they say something like: "Actually, the best times in my life have been . . ." and start talking with great enthusiasm about designing and sewing quilts, rock climbing, playing music, working on a basement lathe, or about other activities that require concentrated skill, that do not separate the individual from the end result of his or her effort, and that provide the kind of exhilaration and high focused attention of flow. So . . .

we are still able to keep in sight those *vivid signposts* that show what it is that makes life worth living . . . [On] reflecting on such occasions, people often say that not only was the experience enjoyable at the time, but that it helped them grow and become more than they had been. Compared to such optimal experiences, much television watching could be deemed a waste of time . . . wasting it amounts to wasting life.⁴³

Casting television in terms of the symmetrical relationship of correlational coexistence developed earlier, we can see that the medium is just not enough for humans to make the center of their lives. It does not call forth their humanity in any depth. Hence, Kubey and Csikszentmihalyi worry that by spending so much time viewing television "one may well lose opportunities to grow as a human being."⁴⁴

Kubey and Csikszentmihalyi find that a mistaken cultural assumption underlies much of the appeal of television. For them the mistaken cultural assumption is narrowly one of thinking that physical and mental exertion are bad, and that they are unrelated to human growth and living a worthwhile life.⁴⁵ In contrast, for us, the more comprehensive mistaken assumption is that technology generally can fulfill our aspirations for freedom and enrichment. Considered from the standpoint of the vision of technology, television is a paradigmatic example of and not an exception to the unimpeded development of technological culture. As for liberation, it is a commodity which does not make demands—in dress, transportation, or manners, or even having to be at home when a program is aired. Following the device's split between means and ends, people exert themselves in labor and expect to relax completely in leisure. They want amusement, not challenge or disturbance. In terms of prosperity, with video cassette players and hundreds of channels, the most, the most varied, and the latest programs can be watched. Advertisements, too, and the settings of the programs themselves celebrate this prosperity of technology. In short, the incredible attraction of television is that it is the homeplace of the vision people are still spellbound by. It confirms them in that vision and tells them what's what in the universe. Its glamour binds and soothes while simultaneously disappointing them with the flatness and shallowness of its nourishment, its ironic unfulfillment. Television as an exemplar of a final commodity represents the ultimate appeal of the promise of technology. It is, then, the success story of the technology. Television is the vision of technological culture.

Much of the remainder of our leisure time is spent, if not in front of a television, then with the commodities advertised on it. But these other forms of final consumption, technological alternatives to television, even

more interactive ones, will manifest the same basic pattern and similar ironic consequences. We have already seen how second homes can relieve people of a midlife crisis and the problems of daily life in the city, and it is now easy to see how the second homes conspicuous leisure agrees with technology. So let us look at how the wild Crazies can be packaged and sold as a commodity for final consumption without materially damaging them as clearcuts would do.

Llama trekking is becoming fashionable. Unlike packing with horses in traditional western style, llamas carry gear but not riders, and, so, trekkers must hike. Moreover, llamas are easier on the environment because they do not eat or weigh as much as a horse, and, being sure-footed with a cloven hoof, they are probably as easy on the trails as elk. What could possibly be inappropriate and insulating about their use?

They can be used appropriately, no doubt, but it pays to attend to the reasons given for using them. Consider an article entitled "Hiking with Llamas in Montana" from the travel section of the Sunday *New York Times*. Aply enough, it is focused on a trek in the Crazies from the Cottonwood Canyon trailhead. The article begins:

"What else can we take?" asked [the guide] . . . I need more weight to balance the llama's packs." "How about a few more liters of wine," suggested his wife . . . "How about some chairs? . . ." "Got camp chairs . . . got solar shower, got the stove and Dutch oven, got the fishing rods, got all the coolers. Got a lot of wine . . ." ⁴⁶

Here one stays clean and comfortable, has plenty to do, and has everything one could possibly wish for. "Released from the dread of freeze-dried lasagna and powdered lemonade," trekkers can look forward, we are told, to dinners like this one served the first night.

We began with wine spritzers, smoked salmon and cheese. The meal progressed through a saffron-infused paella loaded with shrimp, chicken and sausage, a bottle of stream-chilled champagne, a banana concoction flambéed with a self-mocking flourish, coffee and liqueurs. ⁴⁷

Such a meal is billed to enrich the camping experience. And what about the liberation from hardship and misery? The author gives us a measure of these forces one can expect to encounter. Rising early in the morning, "I washed my face in a stream, barely suppressing a howl as the icy water hit my skin. The thought of subjecting my teeth to the same torture was too horrible to contemplate." Finally, an appeal is made directly to the vision of technology, for, while it is acknowledged that such trips are obviously

costly, they are so literally disburdening that, "Llama trips are a dream come true." ⁴⁸

The problem with llama trekking in the above account is twofold. On the one hand, it insulates one from the conditions of the place, smoothing out even the forbidding ruggedness of the Crazies, narrowing ones contact with them, and making wilderness an easily consumable package. The blank spot one is freed up for is filled with such an array of goods that ones attention is directed from the place to the specialness of what is packed in. Our everyday problem-solving is never thrown radically into question—one always has something to do. In these ways, using llamas (as advertised above), impoverishes our encounter with a wild place. We experience a wilderness veneer, a wilderness completely on our own terms.

The second problem has to do with the standards to which such llama trekking appeals. They are technological through and through, valuing technological novelty and availability. Others, sharing these standards, will greet them by asking if such standards can be met better in other ways, perhaps more socially just, considering these are public lands. If disburdened hiking is all one wants and if it takes the importation of plenty of wine, French cuisine, inflatable rafts and solar showers to make the experience really attractive and palatable, then why not punch a road further into the core of the mountains, nearer to the lakes? Then such experience would be available to nearly everyone and not just yuppies.

The next step would be to put a climb up Crazy Peak in my basement or at a health club. The dream would be to mimic the sights, smells, wind, terrain and other features of the experience in uncrude ways. Where does the separation pattern's wounding split between means and ends occur here? Certainly I am not cashing in my prior exertion in labor for what could be called a passive affair such as when I drive a car. Challenge and enjoyment may occur at the same time. Certainly my mind is not reading a book while my body pedals an exercycle, splitting mind from body, myself from others. Body and mind do attend to the same pictured world of the Crazies, the same scenes and terrain. But virtual reality is at our disposal, on our turf and terms, and is disconnected from the larger, fourfold context of our lifeworld. ⁴⁹ As such, it is essentially merely an instance of domination at the expense of relationship. The correlational coexistent thing has dropped away entirely. In the setting of virtual reality we may see into a playful story of Chuang Tzu's in a new and disquieting way. After having had a vivid dream the night before that he was butterfly, Chuang Tzu said he found himself puzzled about whether he was now Chuang Tzu who had dreamed he was a butterfly or a butterfly who is now dreaming it is Chuang Tzu. The vertigo we feel here is not one from heights, although it is not as safe as it might be were we merely on a flight simulator.

The old saw is that humans are acquisitive animals, and the more we have the more we want. From an ethnocentric perspective, it makes some sense why we would think this is "just human nature." Now we can see that people have an insatiable need to have more because they have so superficially. What is it to consume something? "To consume is to use up an isolated entity without preparation, resonance, or consequence."⁵⁰ A commodity made available by a device is free of preparation, does not resonate with the natural, cultural, or local world, and, since it takes so little time and makes such a muted statement, it is unlikely to turn a person's life one way or another, or affirm forcefully its present direction. A person consumes it and is ready to follow the same basic pattern in consuming more. Seeing how much consuming people can cram into their day and their life becomes a norm.

So while technology is successful on its own terms when it makes goods available, its success is merely a pyrrhic victory. What makes good things rich and involving has been lost. We have been seduced by a shallow semblance. Thus, technology fails to deliver happiness, not because it fails to make goods available, but because such goods as it does make available turn out to be merely ironic goods. What seemed promising in the appearance is disappointing in the reality. Our aspirations for freedom and happiness go awry when we attempt to procure them with devices.

So does technology deliver the goods? Does technology help people live more rewarding and meaningful lives? As people make the things that count in their lives technologically available, they empty them of depth and they lose them. It is a lesson our culture has not learned yet when we let television, like Kirk's magic wand, turn us into something less than members of the animal kingdom.

Allegiance to the Vision

Later we will see how Thoreau provides us with a language to make the above point about the shallowness of consumption even more comprehensible when we consider what the villagers would be missing were they, following the vision of technology, to reduce Walden pond to a mere reservoir tapped with plumbing machinery to supply the commodity dishwasher. For now we need to see through this vision and its trivializing underlying ethic. Why does everything get reduced to resources, machinery and commodities? What is the appeal? So long as it seems plausible to us as a culture that we can procure happiness through technological means, we will continue to rearrange nature and culture according to this vision. The final commodities that constitute this goal will be numerous, varied and highly refined in terms of the other standards of ubiquity, instantaneity,

safety and ease. To procure final commodities in this way is the ideal and counts as the larger success story of technology by its own standards.

Accordingly, to be affluent is to possess and consume the most. Affluence has glamour. "It is the embodiment of the free, rich, and imperial life that technology has promised. So it appears from below whence it is seen by most people,"⁵¹ or, as it is seen from the perspective of less affluent countries. This, then, is the goal and blueprint of the technological society our culture is so busily building. This more specifically drawn picture constitutes the *basic vision and framework* of technology.

People's agreement with the framework of technology underlies their concern with the national economy, showing a common commitment to a high and rising standard of living. An increased standard of living cashes out into more numerous, more refined, and varied commodities. Since our culture's commitment to the technological good life is so firm, political campaigns can focus on the economy without ever having to confront deeply the issue of the kind of quality of life the economy provides.

If our culture tacitly grants this commitment to the promise of technology, certain hypothetical imperatives fall out. To produce the most commodities, much of the machinery of production must be centralized, forbiddingly complex, and, by way of productive devices, it must reduce contact of the laborers within the process to narrow points. Normally people do not take up the challenge of making room for intrinsically good work at the expense of technological progress. On the contrary, the degradation of work takes the next step with the elimination of work when more productive and reliable robots take over. A union official is quoted by Borgmann as saying, "We don't like the idea of losing jobs, but it's part of life."⁵² It is part of life, *granted* that technology is the goal.

Earlier we saw that these hypothetical imperatives were at work in the development of the Crazy Mountains. So much raw material would go to waste in the Crazy Mountains if the pinebark beetle infestation were to continue. The notion of waste was contingent upon the assumption that the volume of wood was needed by the technological society. Given that the technological society has consumption as its goal, that bonus pay and second homes are a standard package of a good life, and given that the raw material of Cottonwood Canyon is believed to forward this goal, then the timber of Cottonwood Canyon needs treatment, needs to be logged.

If this vision and framework has at least people's unspoken allegiance, then several consequences follow. First, consumer culture is not at bottom created by advertising as a more Marxist view might hold. To be sure, advertising does regulate, heighten, and make consumption palpable, and, it replaces art by providing us with our orientation in a technological setting where much of the traditional fabric and many of its landmarks have

been dissolved by the device paradigm. Still, the basic framework condition which people move *from* when addressed by advertisements is already in place: From a historical perspective, universal consumption for all is the outcome of the vision of technology and the hold it has upon our culture.⁵³ Then, too, from the perspective of the structure of the device, devices provide an isolated entity to be used up and disposed of. Hence, devices call for a life of consumption and nothing more. Consumer culture is, thus, a consequence of the vision of technology and of the technological device.

This prior cultural condition advertising addresses shows tellingly when people make genuine decisions for consumption. Where do these choice points occur? Most people do not have choices whether they can become potters, poets, wheelwrights, musicians, or follow out liberal arts vocations, whether they can buy from a shoemaker or shoestore, whether they can remain farmers in the face of low product prices and high machinery and energy costs, or whether a nuclear plant gets built near them. But the absence of these kinds of choices rests on a prior agreement to consumption as a way of life. More to the point here, there are occasions when these outside constraints are missing, times when people shape and choose, and at these moments they show how disposed they are to exchange things for commodities. Here we discover whether they are with or against the rule of the vision of technology.

It is to take a condescending view if one excuses families who surrender and betray their tradition by saying that advertisements told them to eat out more often and to refurbish their homes according to the dictates of the Sears catalogue.⁵⁴

When things and engagement are exchanged for commodities and consumption,

what moves one to take the step, so firsthand experience tends to show, is the persistent glamour of the promise of technology; the relief that one looks forward to in having the burden of preparing another meal lifted from one's shoulders; the hope of richer engagement with the world on the basis of greater affluence; the desire to provide one's child with the fullest and easiest means of development; the impatience with things that require constant care and frequent repair; and the wish to affirm one's existence through the acquisition of property that commands respect.⁵⁵

Statistically as well, people do continue to adhere to the promise of liberation and enrichment.⁵⁶ These, then, are those moments when we choose

technology over what Kubey and Csikszentmihalyi spoke of as those vivid signposts that show what it is that makes life worth living.

III

Technological Subversion

Technology subverts in two different ways. On the one hand, technology subverts things when it procures a mere aspect of a thing of depth. Koolwhip replaces whipped cream. This happens, importantly, when we speak abstractly of valuable things, for technologists are apt to present us with an alternative means to the same end. That is, technologists seek functional equivalents for the same value, e.g., a wilderness experience substitute. So value talk, raising the value question, does not cut deeply enough since it likely already assumes the means/end split of the technological framework. If we assume this split, we will be subverted by the device pattern. Technology's underlying ethic of freedom and prosperity has been subverted in this way. So too Charles Taylor's ideals of self-realization, regardless of how well-articulated, will meet a similar fate, becoming subverted to trivialized and deviant modes, unless such ideals of authenticity remain tethered to correlational coexistent things.

In this same fashion, Leopold's land ethic, abstractly conceived outside the context of his book, is subverted by today's Forest Service. Leopold writes:

The 'key-log' which must be moved to release the evolutionary process for an ethic is simply this: quit thinking about decent land-use as solely an economic problem. Examine each question in terms of what is ethically and aesthetically right, as well as what is economically expedient. A thing is right when it tends to preserve the integrity, stability, and beauty of the biotic community. It is wrong when it tends otherwise.⁵⁷

As we saw earlier, the Forest Service by law is required to manage the forest in terms other than economic expediency alone. The doctrine of multiple use, for instance, makes it consider other uses for the forest than those we typically associate with economic expediency, such as timber, grazing and mining. Recreation and big game habitat are the dominant goals for some areas. We saw that timber and hunting were supposed to be treated as equally important in the management plans for part of Cottonwood Canyon. Even if an area is selected to be managed primarily for timber, economic expediency is not the only factor taken into account.

Environmental Impact Statements and Assessments are undertaken so that factors of biological uniqueness, diversity and stability constrain the final management plans and their specifications such as the size, shape and placement of clearcuts. Finally, aesthetics and cultural factors get taken into account in the planning process and can constrain the final management of an area, such as Cottonwood Canyon. The timber sale proposal for Cottonwood Canyon leaves a strip of trees between the road and the clearcuts so that recreationists driving the road will not witness what, to the "average forest visitor," is ugly. And, of course, the clearcuts shaped like avalanche chutes are more costly than they would be if the agency were completely insensitive to beauty. So the Forest Service in a technologically subversive manner is employing the principles set forth by Leopold; yet most people who turn to Leopold's land ethic for guidance are not happy with the way the land is treated by the agency. The Forest Service may maintain some of the scenic factors of the canyon, but the real beauty and special character of the place will have been destroyed.

Technology subverts in a second way when heeding its pull overrides all other commitments and claims. If technology is the agreed upon cultural goal, then whatever comes into conflict with the steps thought to be necessary to achieve that goal will normally be defeated. We saw this defeat and subversion in our culture's willingness to destroy good work for the sake of production, and in choices of commodities over things. The same subversive force undermines the claims of social justice. In this area people do not need to develop an ethic from the ground up as Rolston does for the environment. Respect for the equality of persons in this country has been a force since the beginning, yet glaring inequality persists. Borgmann argues that our society believes that "pervasive relative deprivation fuels the motor of technological advancement."⁵⁸ They think that what the rich have today they will have tomorrow. Hence, so long as they continue to hold this belief, and so long as they are enthralled by the promise of technology, people will tolerate inequality. Global problems of misery and poverty claim people in affluent countries, too. Yet, with too few exceptions, they callously disregard these people in poverty so long the well-off are within the framework of the blessings of technology and continue to be fascinated with affluence.⁵⁹ Living with less affluence will seem to them a sacrifice of their prosperity as long as the vision of technology has their allegiance.

We see this subversion when people face the claims of excellence. People are still somewhat claimed by standards of traditional excellence. Often these standards are at the heart of the kind of life they would have liked to lead if less constrained or wish their children might.⁶⁰ However, our society does not act on behalf of this concern for excellence, e.g.,

people's command of science is weak, participation in politics is minimal, television programs which critics like least are being watched most and so on. Moreover, although people have room and means to shape their leisure, they do not spend their time at those kinds of activities which could suggest a dedication to the pursuit of excellence. Rather all activities suggesting a tie to excellence constitute less than a quarter of the time spent watching television.⁶¹ This means that over eighty-five percent of their leisure time is devoted to something other than traditional excellence, that is, mostly consumption. So here again technology subverts the claims people still feel.

The Underlying Pull to Dominate Nature

Importantly for the task at hand—although it is only one case among many such subversions—technology subverts all other claims of nature and the natural world. From the standpoint of the device paradigm, the environment and nature are areas to bring under control and secure. Nature can be made available in the form of final commodities through roads, trains, the tourist industry, game farms, Disney movies and so forth. It can be brought under control in a defensive and insulating way through securing health, shelter, warmth, clothing and food. It can be brought under control for production through the exploitation of timber, oil, coal, and mineral resources. In bringing nature under control, technology has very often been voracious and short-sighted, yet, when technology proceeds this way, it fails by its own standards of safety and ease. It may saddle us with greater hardships than it relieves us from. Hence, the environmental crisis resulting from the physical limits to growth is not simply a point of concern raised by those living at the margins of the ruling paradigm. A telling sign of this is the fact that it has come to the center of attention. Everyone understands the need for trade-offs, achieving balance between environmental concerns and future growth. We need to live between the rims. This means that the environmental crisis is not radically a challenge to the framework of technology from outside its own terms, but rather challenges the framework *from within* to become more comprehensive.

How will environmental stability be sought in terms of the vision of technology? For this safety and ease, the device paradigm must be extended to the whole globe and take an enlightened, long-term view. This may fail, but how is it likely to succeed were it to succeed? Apart from economizing through engineering ingenuity and the elimination of wasteful practices, citizens of advanced countries, Borgmann believes alarmingly enough, will make the poor, and poor from other countries, pay the price of coming to terms with our physical limits. If successful, the technology

society's global future will tend toward a physically homeostatic equilibrium, though this is not likely soon.⁵² (Even if scientists and engineers save the planet, the question for us, unlike other critics of technology, always is: will it be a planet worth living on?)

If this goal and its blueprint have people's tacit if not explicit assent, what else will matter about nature? Clearly, nature from this standpoint is viewed with merely sophisticated, enlightened self-interest. Clearly, just as in the cases cited earlier, whatever other claims are registered from nature and natural things in their own right, so long as people remain in agreement with and enthralled by the promise of technology, these claims will not be heard and made room for, especially when making room for them would entail some *sacrifice* of a high and rising standard of living and hence to the good life of consumption. We found this agreement with technology to be at the bottom of the imperatives to develop Cottonwood Canyon. Then, too, the initial road bulldozed into Cottonwood Canyon makes the wilderness more commodiously available for a summer home. As should be clear by now, the underlying ethic of technology competes with and defeats an environmental ethic, and so the claims made upon people by nature and natural beings in accordance with Rolston's environmental altruism are likely to fare worse than the claims of social justice. So long as the glamorous appeal of consumption has their ears, their heeding, the story of the domination of nature, in one variant or another, will subvert the story of a new respect for nature. When environmentalists argue people can have it both ways, both the environment they want and the technological good life, the contradiction, when spelled out by a theory of technology, becomes obvious. So, if the claims of nature and natural things in their own right are to be respected, honored, and acted upon with consequence, the ruling vision of technology must be protested and successfully challenged.

How can this challenge take place? How can we begin to initiate a reform of technology? We have taken a major step in this direction by locating the fundamental problem as technology and by evaluating critically the vision of technology in terms of what it does to things and its ironic consequences. Reflection enables us to step back and regard the appeal of consumption as glamorously attractive but disappointing. The hold that the vision of technology has upon us is thereby weakened, yet more needs to be and can be done. It is not enough merely to cut back or cut out entirely the amount of leisure time devoted to television, for otherwise we will easily fall back into its spell or that of another form of consumption. To counter consumption more generally we need to find attractive and substantial alternative things.

Things need more of a say in our lives. What we have yet to learn is the complete force of the appeal of things and appropriate ways to communicate that appeal. So that we may heed them, their voices must be given a hearing again and for the first time. Then, unlike those choice-points illustrated above, when we choose between technology and things of final importance, we will know the full weight of what we are agreeing to and protesting.

Technology will subvert us if we speak abstractly of correlational coexistent things, so how can we better awaken an understanding of the force of the appeal of these things? How best can we speak of *the things that matter* in the ways they matter? Over a cup of coffee in graduate school, I fell to reminiscing about the Crazies with a fellow student. As I stumbled about seeking the right words, I was surprised to learn that she understood well enough; she could see what mattered to me, that this came to life in my manner of speaking. Isn't that the way of it so often? To be intelligible the felt presence of the thing exacts a kind of showing. In chapter 11, "Learning to Speak Again," I will develop this idea. In the next chapter, I will exhibit it. My signpost will be an actual landmark, the Crazy Mountains themselves. When it comes to bringing home to us the full appeal, the real force of correlational coexistent things, philosophy and theory need to give way to narrative and poetry. Ultimately, it will be these powers of things and what they tell us that enable us to turn technology.

232

Notes

52. *Ibid.*, 123.

53. *Ibid.*, 52.

54. *Ibid.*, 105.

55. *Ibid.*

56. *Ibid.*, 107.

57. Aldo Leopold, "The Land Ethic," in *A Sand County Almanac and Sketches Here and There* (New York: Oxford University Press, 1968), 224.

58. Borgmann, *Technology*, 112.

59. *Ibid.*, 127.

60. *Ibid.*, 126.

61. *Ibid.*, 127-29.

62. *Ibid.*, 148. Yet, even if technology succeeds by its own standards, it exacts a price. "If that technological totalitarianism comes to pass, life will take on an essentially secure, trite, and predictable cast," 144.

Chapter 5

1. *The Ethics of Authenticity* (Cambridge: Harvard University Press, 1991), 100–01.
2. Marx Weber's term.
3. See Taylor, "The Iron Cage," *The Ethics of Authenticity*, 93–108.
4. Albert Borgmann, *Technology and the Character of Contemporary Life: A Philosophical Inquiry* (Chicago: University of Chicago Press, 1984), 41ff.
5. *Ibid.*, 36.
6. David Ehrenfeld, *The Arrogance of Humanism* (New York: Oxford University Press, 1981).
7. Borgmann, *Technology*, 42.
8. *Ibid.*, 196.
9. *Ibid.*, 196–210.
10. *Ibid.*, 42.
11. Witold Rybczynski, *Home: The Short History of an Idea* (New York: Penguin, 1986), 219–20.
12. Borgmann, *Technology*, 47.
13. Don Ihde, *Technics and Praxis* (Boston: D. Reidel, 1979), 21.
14. Borgmann, *Technology*, 53.
15. *Ibid.*, 140–41.
16. Robert Kubey and Mihaly Csikszentmihalyi, *Television and the Quality of Life: How Viewing Shapes Everyday Experience* (Hillsdale, NJ: Lawrence Erlbaum Associates, 1990), ix.
17. *Ibid.*, 173–74.
18. *Ibid.*, 102–03.
19. *Ibid.*, 108–17, 173.
20. *Ibid.*, 189.
21. *Ibid.*, 173–74.
22. *Ibid.* 79.
23. *Ibid.*, 171–74.
24. *Ibid.* 208–09, 213.

25. *Ibid.*, 183.
26. *Ibid.*, 101, 119–27, 173.
27. *Ibid.*, 83.
28. *Ibid.*, 172.
29. *Ibid.*, 108–17, 172.
30. *Ibid.*, 127–29, 172–73.
31. *Ibid.*, 133–35, 172–73.
32. *Ibid.*, 101.
33. *Ibid.*, 185–90.
34. *Ibid.*, 181–216.
35. *Ibid.*, 157–69, 172.
36. *Ibid.*, 117.
37. *Ibid.*, 168, 173, 184ff.
38. *Ibid.*, 157–169, 172–73.
39. *Ibid.*, 204.
40. *Ibid.*, 190.
41. *Ibid.*, 145.
42. *Ibid.*, 147.
43. *Ibid.*, 100–01 (emphasis mine), see also, 196.
44. *Ibid.*, 207.
45. *Ibid.*, 215–16.
46. Nancy Debevoise, "Hiking with Llamas in Montana," *New York Times* (April 20, 1986), Sec. 20, 10.
47. *Ibid.*
48. *Ibid.*
49. In his most recent work Borgmann sees virtual reality or hyperreality vs. focal realism as pivotal. Accordingly, he has modified the terms of his criticism, e.g., disposability and discontinuity vs. commanding presence and continuity. See his *Crossing the Postmodern Divide* (Chicago: University of Chicago Press, 1992).
50. Borgmann, *Technology*, 53.
51. *Ibid.*, 223.

Philosophical Ethics

by

Deborah Johnson

Chapter 2 from
Computer Ethics, Third Edition
Prentice Hall, 2001

Before embarking on analysis of the ethical issues surrounding computer and information technology, it will be helpful to discuss the nature of ethical analysis, and to become familiar with some traditional ethical concepts and theories. This chapter shows how ethical analysis can proceed so as to produce insight and better understanding. The chapter also explains concepts and theories that philosophers have found particularly useful in discussing ethical issues.

We often overhear or participate in discussions of ethical issues. Think, for example, of the heated discussions you have heard about government restrictions on individual freedom (e.g., censorship of the Internet, the right to assisted suicide). Or think of discussions about abortion, affirmative action, and the distribution of wealth in our society. Often when individuals are asked to explain why they think a behavior or policy is wrong, they have difficulty articulating their reasons. Sometimes it seems that individuals who are expressing moral opinions are simply reacting as they think most people in their society react or they espouse ideas they heard friends or relatives espouse. Many who have fairly strong moral beliefs have only a very vague sense of why the behavior or policy is unfair or irresponsible or harmful. These unexamined beliefs can be the starting place for ethical analysis, though it is important to understand that they are only starting places.

Discussions at this level may quickly end unresolved because the individuals involved are not able to provide good reasons for believing as they do. It is difficult or impossible to discuss the issues rationally, let alone resolve them. If discussion stays merely at the level of statements of belief, discussants will walk away thinking that everyone is entitled to his or her own opinion and there is no point talking about ethics, except perhaps to see where others stand. Discussants won't have learned anything or come to understand the ethical issues any better.

This book is an undertaking in philosophical analysis, and philosophical analysis proceeds on the premise that we must examine the reasons we have for our moral or ethical beliefs. In philosophical, ethical analysis the reasons for moral beliefs are articulated, and then critically evaluated. The reasons you give for holding an ethical belief or taking a position on an ethical issue can be thought of as an argument for a claim. The argument has to be "put on the table," and once there, it can be evaluated in terms of its plausibility, coherence, and consistency. Once stated, we can ascertain whether the argument does, indeed, support the claim being made or the position being taken.

This critical evaluation is often done in the context of trying to convince someone to reject a position, or to adopt another position, but it may also be done simply to explore a claim. When you critically evaluate the argument supporting a claim, you come to understand the claim more fully. A critical examination of the underpinnings of moral beliefs sometimes leads to a change in belief, but it may also simply lead to stronger and better understood beliefs.

In philosophical analysis, not only must you give reasons for your claims, you are also expected to be consistent from one argument or topic to the next. For example, instead of having separate, isolated views on abortion and capital punishment, philosophical analysis would lead you to recognize that both your views on abortion and your views on capital punishment rest on a claim about the value of human life and what abrogates it. Philosophical analysis would lead you to inquire whether the claim you made about the value of human life in the context of a discussion of capital punishment is consistent with the claim you made about the value of human life in the context of a discussion of abortion. If the claims appeared to be inconsistent from the one context to the next, then you would be expected to change one of your claims or provide an account of how the two positions can be understood as consistent. In other words, you would show that seemingly inconsistent views are in fact consistent.

Philosophical analysis is an ongoing process. It involves a variety of activities. It involves

expressing a claim and putting forward an argument or reasons for the claim, and it involves critical examination of the argument. If the argument does not hold up to critical examination, then it might be reformulated into a revised argument, perhaps rejecting aspects of the original argument but holding on to a core idea. The revised argument, then, has to be critically examined, and so on, with ongoing reformulation and critique. Philosophers often refer to this process as a *dialectic* (which is related to the word *dialogue*). We pursue an argument to see where it goes and to find out what you would have to know or assert to defend the argument and establish it on a firm footing.

In addition to moving from claims to reasons and arguments, and from one formulation of an argument to another, better formulation, the dialectic also moves back and forth from cases to principles or theory. To illustrate, take the issue of euthanasia. Suppose you start out by making the claim that euthanasia is wrong. You articulate a principle as the reason for this claim. Say, the principle is that human life has the highest value and, therefore, human life should never be intentionally ended. You might then test this principle by seeing how it applies in a variety of euthanasia cases. For example, is it wrong to use euthanasia when the person is conscious but in extreme pain? When the person is unconscious and severely brain damaged? When the person is terminally ill? When the person is young or elderly? Since your principle concerns the value of human life, it has implications beyond the issue of euthanasia. Hence, you might also test it by applying it to completely different types of cases. Is the intentional taking of human life wrong when it is done in a war situation? Is intentional killing wrong when it comes to capital punishment? Given your position on these cases, you may want to qualify the principle or you may hold to the principle and change your mind about the cases. For example, after seeing how the principle applies in various cases, you may want to qualify it so that you now assert that one should never intentionally take a human life except in self-defense or except when taking a life will save another life. Or you might reformulate the principle so that it specifies that the value of human life has to do with its quality. When the quality of life is significantly and permanently diminished, while it is still not permissible to intentionally kill, it is morally permissible to let a person die.

The dialogue continues as the dialectic leads to a more and more precise specification of the principle and the argument. The process clarifies what is at issue and what the possible positions are. It moves from somewhat inchoate ideas to better and better arguments, and more defensible and better articulated positions.

The dialectic (from an initial belief to an argument, from argument to better argument, and from theory to case, and back) does not always lead to definitive conclusions or unanimous agreement. Therefore, it is important to emphasize that understanding can be improved, progress can be made, even when one has not reached definitive conclusions. Through the dialectic we learn which arguments are weaker and stronger and why. We come to understand the ideas that underpin our moral beliefs. We develop deeper and more consistent beliefs and we come to understand how moral ideas are interrelated and interdependent.

As you will see in a moment, a familiarity with traditional ethical theories will help in articulating the reasons for many of your moral beliefs. Ethical theories provide frameworks in which arguments can be cast. Moreover, ethical theories provide some common ground for discussion. They establish a common vocabulary and frameworks within which, or against which, ideas can be articulated.

DISTINGUISHING DESCRIPTIVE AND NORMATIVE CLAIMS

In any discussion of ethics, it is important to recognize the distinction between descriptive and normative claims. In a sense and partly, this is the distinction between facts and values, but the matter of what counts as a fact is very contentious in philosophy. So, it will be better to stay with the terms descriptive and normative. *Descriptive* statements are statements that describe a state of affairs in the world. For example, “The car is in the driveway.” And “Georgia is south of Tennessee.” In addressing ethical issues and especially the ethical issues surrounding computer and information technology, it is quite common to hear seemingly factual statements about human beings. The following are descriptive statements: “Such and such percentage of the people surveyed admitted to having made at least one illegal copy of computer software.” “The majority of individuals who access pornographic Web sites are males between the ages of 14 and 35.” “Such and such percentage of U.S. citizens use the Internet to obtain information on political candidates.” “In all human societies, there are some areas of life that are considered private.” These statements describe what human beings think and do. They are *empirical* claims in the sense that they are

statements that can be verified or proven false by examining the state of affairs described. To be sure, it may not be easy to verify or disconfirm claims like these, but in principle it is possible. Observations can be made, surveys can be administered, people can be asked, and so on.

Social scientists gather empirical data and report their findings, both on moral and nonmoral matters. When it comes to morality, psychologists and sociologists might do such things as identify the processes by which children develop moral concepts and sensibilities. Or they may measure how individuals value and prioritize various goods such as friendship, privacy, and autonomy. When anthropologists go to other cultures, they may describe complex moral rules in that culture. They are describing lived and observed moral systems. Similarly, historians may trace the development of a particular moral notion in an historical period.

All of these social scientific studies are descriptive studies of morality; they examine morality as an empirical phenomenon. They don't, however, tell us what is right and wrong. They don't tell us what people *should* think or do, only what people, in fact, think and do.

In contrast, philosophical ethics is *normative*. The task of philosophical ethics is to explore what human beings ought to do, or more accurately, to evaluate the arguments, reasons, and theories that are proffered to justify accounts of morality. Ethical theories are prescriptive. They try to provide an account of why certain types of behavior are good or bad, right or wrong. Descriptive statements may come into play in the dialectic about philosophical ethics, but normative issues cannot be resolved just by pointing to the facts about what people do or say or believe. For example, the fact (if it were true) that many individuals viewed copying proprietary software as morally acceptable would not make it so. The fact that individuals hold such a belief is not an argument for the claim that it is morally permissible to copy proprietary software. You might wish to explore why individuals believe this to see if they have good reasons for the belief. Or you might wish to find out what experiences have led individuals to draw this conclusion. Still, in the end, empirical facts are not alone sufficient to justify normative claims. Figuring out what is right and wrong, what is good and what is bad, involves more than a descriptive account of states of affairs in the world.

The aim of this book is not to describe how people behave when they use computers. For this, the reader should consult social scientists—sociologists, anthropologists, political scientists, and psychologists. Rather the aim of this book is to help you understand how people ought to behave when they use computers and what rules or policies ought to be adopted with regard to computer and information technology.

ETHICAL RELATIVISM

We can begin our examination of ethical concepts and theories by examining a prevalent, often unexamined moral belief. Many believe that “ethics is relative.” This seems like a good starting place. This claim can be examined carefully and critically. We can begin by formulating the idea as a theory consisting of a set of claims backed by reasons.

The idea of ethical relativism seems to be something like this: “What is right for you may not be right for me,” or “I can decide what is right for me, but you have to decide for yourself.” When we take this idea and formulate it into a more systematic account, it seems to encompass a negative claim (something that it denies), and a positive claim, (something it asserts). The negative claim appears to be: “There are no universal moral norms.” According to this claim, there isn't a single standard for all human beings. One person may decide that it is right for him to tell a lie in certain circumstances, another person may decide that it is wrong for her to tell a lie in exactly the same circumstances, and both people could be right. So, the claim that “right and wrong are relative” means in part that there are no universal rights and wrongs.

The positive claim of ethical relativism is more difficult to formulate. Sometimes ethical relativists seem to be asserting that right and wrong are relative to the individual, and sometimes they seem to assert that right and wrong are relative to the society in which one lives. I am going to focus on the latter version, and on this version the relativist claims that what is morally right for me, an American living in the twenty-first century, could be different than what is right for a person living, say, in Asia in the fifth century. The positive claim of relativism is that right and wrong are relative to your society.

Ethical relativists often cite a number of descriptive facts to support these claims:

1. They point to the fact that cultures vary a good deal in what they consider to be right and wrong. For example, in some societies, infanticide is acceptable while in other societies it is considered wrong. In some societies,

- it is considered wrong for women to go out in public without their faces being covered. Polygamy is permissible in some cultures; in others it is not. Examples of this kind abound.
2. Relativists also point to the fact that the moral norms of a given society change over time so that what was considered wrong at one time, in a given society, may be considered right at another time. Slavery in America is a good example of this since slavery was considered morally permissible by many in the United States at one time, but is now illegal and almost universally considered impermissible.
 3. Relativists also point to what we know about how people develop their moral ideas. We are taught the difference between right and wrong as children, and what we come to believe is right or wrong is the result of our upbringing. It depends on when, where, how, and by whom we were raised. If I had been born in certain Middle Eastern countries, I might believe that it is wrong for a woman to appear in public without her face covered. Yet because I was raised in the United States in the twentieth century, by parents who had Western ideas about gender roles and public behavior, I do not believe this. Of course, parents are not the only determinant of morality. A person develops moral ideas from the experiences he or she has in school, at work, with peers, and so on.

It is useful to note that we have already made progress simply by clearly and systematically formulating the idea of ethical relativism, an idea you may have entertained or heard expressed, but never had a chance to examine carefully. Moreover, we have been able to identify and articulate some reasons thought to support ethical relativism. With the idea and supporting evidence now “on the table,” we can carefully and critically examine them.

The facts which ethical relativists point to cannot be denied. For example, I would not want to take issue with the claims that:

1. There is and always has been a good deal of diversity of belief about right and wrong.
2. Moral beliefs change over time within a given society.
3. Social environment plays an important role in shaping the moral ideas you have.

However, there does seem to be a problem with the connection between these facts and the claims of ethical relativism. Do these facts show that there are no universal moral rights or wrongs? Do they show that right and wrong are relative to your society?

On more careful examination, it appears that the facts cited by ethical relativists do not support their claims. To put this another way, we can, without contradiction, accept the facts and still deny ethical relativism. The facts do not necessitate that there are no universal moral standards or that ethics is relative. Lest there be no confusion, you should recognize that “ethics is relative” could be interpreted either as an empirical or a normative claim. As an empirical claim, it asserts that ethical beliefs vary; as a normative claim it asserts that right and wrong (not just beliefs about, but what is actually right and wrong) vary.

If we understand the claim “ethics is relative” to be a description of human behavior, then it *does* follow from the facts cited. Indeed, it is redundant of the facts cited, for as a description of human behavior, it merely repeats what the facts have said. Ethical beliefs vary. Individuals believe different things are right and wrong depending on how and by whom they have been raised and where and when they live.

On the other hand, if we understand “ethics is relative” to be a normative claim, a claim asserting the negative and/or positive parts of ethical relativism, then it is not redundant, and the facts do not support the claims. Here the leap from facts to conclusion is problematic for a number of reasons. For one, the argument goes from a set of “is” claims to an “ought” claim and the ought-claim just doesn’t follow (in a straightforward way) from the is-claims. The argument goes like this: “People do *a*; people do *b*; people do *c*; and therefore people ought to do *x*.”

Moreover, the facts are compatible with the opposite conclusion. That is, it is possible that a universal moral code applies to everyone even though some or all fail to recognize it. Centuries ago when some people believed the earth was flat and others claimed that it was round, the earth’s shape was not relative. The fact that there is diversity of opinion on right and wrong does not tell us anything about whether right and wrong are relative. The facts are compatible both with the claim that there is no universal right and wrong and with the claim that there is a universal right and wrong.

Taking this one step further, let’s consider the fact that our moral beliefs are shaped by our social environment. While it is true that our moral beliefs are shaped by our social environment, this says nothing about the rightness or wrongness of what we believe. Racism and sexism are good examples of moral attitudes we may acquire from our environment but which turn out on reflection to be unjustifiable (bad) ideas.

We must also be careful about what is inferred from the fact that there is diversity in moral beliefs. This diversity may be misleading; that is, it may be superficial rather than deep. Relativists seem to be focusing on specific practices and there is still the possibility that universal nouns underlie these. Moral principles such as “never intentionally harm another person” or “always respect human beings as ends in themselves” are of such generality that they could be operative in many or all cultures but expressed in different ways. What is meant by “harm,” “respect,” and “human being” may vary although there is some principle to which all people adhere. So, it is possible that there are some universal principles at work, but they are hidden from sight due to the diversity of expression or interpretation of the principle.

Social scientists have certainly tried to find patterns within the apparent diversity. Some have asserted, for example, that all cultures have prohibitions on incest or, more recently, that while there is a great deal of diversity about what is considered private, all cultures consider some aspect of the lives of individuals private.

Even so, while such patterns have important implications for the study of ethics, we have to remember that establishing patterns across cultures is descriptive, and it is another matter to determine what these claims imply about how people ought to behave. In a moment, when we examine utilitarianism, we will see an example of a very general normative principle that is compatible with a diversity of practices. Utilitarianism is a form of consequentialism, and such theories assert that individuals should always do what will maximize good consequences. Individuals in quite different situations may be doing very different things but all in accordance with this same principle.

In any case, the facts pointed to by relativists do not support their claim that there are no universal moral rights and wrongs. Nor do the facts cited support the ethical relativist’s claim that right and wrong are relative to one’s society. Pointing to what people believe to be right and wrong tells us nothing about what is right or wrong. The fact that people behave in accordance with the norms of their society is not evidence for the claim that they ought to.

It is important to keep in mind that the criticism I have just made of the ethical relativist’s argument does not establish that there are universal rights and wrongs. The criticisms show only that the arguments ethical relativists might put forward to support their position do not work. You may be able to come up with a different argument on behalf of ethical relativism, and then your argument would have to be carefully and critically examined.

Before you try to defend ethical relativism, however, there are some serious problems with the theory and you ought to be aware of these. Ethical relativism, as I have formulated it, appears to be self-contradictory. The negative and positive claims appear to contradict each other. In saying that right and wrong are relative to one’s society, ethical relativists seem to be asserting that one is bound by the rules of their society. The relativist seems to be saying that what is right for me is defined by my society, and what is right for a member of an African tribe is what is set by the standards of her or his tribe. It would seem, then, that I ought to do what is considered right in my society, and everyone else ought to do what is considered right in their society. Notice, however, that if this is what ethical relativists mean, they are affirming a universal moral principle. On the one hand, they deny that there are universal rights and wrongs, and, on the other hand, they assert one. If I have accurately depicted ethical relativism, then it appears to be an utterly incoherent (self-contradictory) theory.

If this were a book about ethical relativism alone, I would try to resurrect the theory by reformulating its claims and bringing in other arguments to support it. All I will do instead is to point to what I think is an important moral motive buried in relativism. Often what ethical relativists are trying to do is make the point that no one should denigrate, ridicule, and disrespect people who have beliefs that are different from their own. In other words, you shouldn’t judge people from other times or places by the standards of your own morality. It is arrogant, relativists might say, to believe that you as an individual or a member of a particular society have the correct moral views and that anyone who doesn’t agree with you is wrong. Such relativists would argue that we ought to respect people with moral beliefs different from our own.

This seems an important and worthy point that some relativists want to make. Still, it should be noted that to take this position is, again, to take a universal position. You are claiming that “everyone ought” to adopt a position which might be characterized as tolerance or respect for others.

So, it would seem that we cannot assert both that everyone ought to respect the views of others and at the same time hold that ethics is relative. If toleration is the motive behind relativism, this motive has an implicit universal character and that conflicts with relativism’s claim that there is no universal right and wrong. To see the contradiction, consider the case of someone who lives in a society that does not believe

in toleration. According to relativism, this person need not be tolerant of others. Relativism says right and wrong is relative to your society and in this person's society there is nothing wrong with being intolerant. Thus, it would seem that if underlying one's belief in relativism is the belief that everyone should be tolerant of the beliefs of others, relativism is not going to be an acceptable theory, at least not if it is formulated as I have formulated it.

Case Illustration

To see these and other problems with ethical relativism, consider a hypothetical case. Suppose, by a distortion of history, that computers were developed to their present sophistication in the late 1930s and early 1940s. World War II is in progress. You are a German citizen working for a large computer company. You are in charge of the sales division and you personally handle all large orders. You are contacted by representatives of the German government. The German government has not yet fully automated its operations (computers are still relatively new) and it wants now to purchase several large computers and several hundred smaller computers to be networked.

You read the newspapers and know how the war is proceeding so you have a pretty good idea of how the German government will use the computers. It is quite likely they will use the computers to help keep track of their troops and equipment, to identify Jews and monitor their activities, to build more efficient gas chambers, and so on. The question is, if you were an ethical relativist would it be permissible for you to sell the computers to Hitler and his government?

The question reveals some practical problems with relativism. Relativism specifies that what is right for you is what is considered right in your society. But, how do you figure out what the standards of your society are? Are the standards of your society what the political leaders say and do or what the majority in the society believe? If these are different, what should you do? To put this in another way, is Hitler necessarily abiding by the standards of his society or is he going against these? If he is going against these standards, then perhaps he is doing wrong and you would be doing wrong to support him. It may not be easy to tell whether Hitler is adhering to or rejecting the standards of his society. Hence, it may not be so easy to use relativism to guide your actions.

This leads to another problem with relativism. Suppose Hitler and most German citizens agree that Hitler's agenda is right. Nevertheless, you disagree. Relativism seems to rule out the possibility of resistance or rebellion in such a situation. If someone rebels against the standards of her society, it would seem she is doing wrong for she is acting against relativism's claim that what is right for you is what is considered right in your society. Many of our greatest heroes, Socrates, Martin Luther King, Ghandi, even Jesus, would, on this account, be considered wrong or bad. They acted against the standards of their societies.

So, if Hitler and most Germans agreed that the German agenda was right, it would seem that you, as a relativist, would have to conclude that it is right for you to sell the computers to the German government (even if you personally objected to Hitler's agenda).

Now suppose that one of your friends from the United States or somewhere else finds out about the sale and asks you why you did this. What do you say? You answer: It was the right thing to do because it was consistent with the standards and beliefs in my society. From your friend's perspective, this may seem a very feeble answer. The fact that some type of behavior is the standard in your society seems an inadequate moral reason for adopting the standard as your own. It doesn't seem a very good reason for acting in a certain way, especially when the act has significant negative consequences.

Summarizing what has been said so far about the problems with relativism, it suffers from three types of problems. First, the evidence that is used to support it, does not support it. Second, proponents cannot assert both the negative and the positive claims of relativism without inconsistency. By claiming that everyone is bound by the rules of his or her society, the ethical relativist makes a universal claim and yet the relativist claims there are no universal rights and wrongs. And, third, the theory, as the Hitler case illustrates, does not seem to help in making moral decisions. Relativism, at least as I have formulated it, does not help us figure out what to do in tough situations. It recommends that we adhere to the standards in our society and yet it doesn't help us figure out what these standards are. Moreover, doing something because it is the standard in your society does not seem a good reason for doing something.

Where do we stand now? It is important to note that we have made progress even though we have not formulated a moral theory that is defensible. Partly our progress is negative. That is, we have identified

some arguments that don't work. At the same time, we have learned about some of the difficulties in taking a relativist position and are therefore in a better position to reformulate the theory. Perhaps, most important of all, we have seen the challenge of developing and defending ethical claims.

Our exploration of ethical relativism has hardly scratched the surface. You may want to reformulate ethical relativism so as to avoid some of the arguments given against it. You may want, for the time being, to take what might be called "an agnostic position." As an agnostic, you claim that you don't yet know whether there are universal rights and wrongs but you would also claim that you do not have sufficient reasons for ruling out the possibility either. You will wait and see, keeping an open mind, and being on the alert for implausible and inconsistent claims.

UTILITARIANISM

Utilitarianism is an ethical theory claiming that what makes behavior right or wrong depends wholly on the consequences. In putting the emphasis on consequences, utilitarianism affirms that what is important about human behavior is the outcome or results of the behavior and not the intention a person has when he or she acts. On one version of utilitarianism, what is all important is happiness-producing consequences (Becker and Becker, 1992). Crudely put, actions are good when they produce happiness and bad when they produce the opposite, unhappiness. The term *utilitarianism* derives from the word *utility*. According to utilitarianism actions, rules, or policies are good because of their usefulness (their utility) in bringing about happiness.

Lest there be no confusion, philosophers are not always consistent in the way they use the terms utilitarianism and consequentialism. Sometimes, consequentialism is seen as the broadest term referring to ethical theories that claim that what makes an action right or wrong is the consequences and not the internal character of action. Utilitarianism is, then, a particular version of this type of theory with the emphasis specifically on happiness-producing consequences. That is the way I shall use these terms, though I warn readers that the distinction sometimes is made in just the opposite way, that is, with utilitarianism seen as the broadest theory and consequentialism as a particular form of utilitarianism.

In any case, in the version on which I will focus, the claim is that in order to determine what they should do, individuals should follow a basic principle. The basic principle is this: *Everyone ought to act so as to bring about the greatest amount of happiness for the greatest number of people.*

But, what, you may ask, is the "proof" of this theory? Why should each of us act to bring about the greatest amount of happiness? Why shouldn't we each seek our own interest?

Intrinsic and Instrumental Value

Utilitarians begin by focusing on values and asking what is so important, so valuable to human beings, that we could use it to ground an ethical theory. They note that among all the things in the world that are valued, we can distinguish things that are valued because they lead to something else from things that are valued for their own sake. The former are called *instrumental* goods and the latter *intrinsic* goods. Money is a classic example of something that is instrumentally good. It is not valuable for its own sake, but rather has value as a means for acquiring other things. On the other hand, intrinsic goods are not valued because they are a means to something else. They have qualities or characteristics that are valuable in themselves. Knowledge is sometimes said to be intrinsically valuable. So, is art because of its beauty. You might also think about environmental debates in which the value of nature or animal or plant species or ecosystems are said to be valuable independent of their value to human beings. The claim is that these things have value independent of their utility to human beings.

Having drawn this distinction between instrumental and intrinsic goods, utilitarians ask what is so valuable that it could ground a theory of right and wrong? It has to be something intrinsically valuable, for something which is instrumentally valuable is dependent for its goodness on whether it leads to another good. If you want *x* because it is a means to *y*, then *y* is what is truly valuable and *x* has only secondary or derivative value. Utilitarianism, as I am using the term, claims that happiness is the ultimate intrinsic good, because it is valuable for its own sake. Happiness cannot be understood as simply a means to something else. Indeed, some utilitarians claim that everything else is desired as a means to happiness and that, as a result, everything else has only secondary or derivative (instrumental) value.

To see this, take any activity that people engage in and ask why they do it. Each time you will find

that the sequence of questions ends with happiness. Take, for example, your career choice. Suppose that you have chosen to study computer science so as to become a computer professional. Why do you want to be a computer professional? Perhaps you believe that you have a talent for computing, and you believe you will be able to get a well-paying job in computer science—one in which you can be creative and somewhat autonomous. Then we must ask, why are these things important to you? That is, why is it important to you to have a career doing something for which you have a talent? Why do you care about being well paid? Why do you desire a job in which you can be creative and autonomous? Suppose that you reply by saying that being well paid is important to you because you want security or because you like to buy things or because there are people who are financially dependent on you. In turn, we can ask about each of these. Why is it important to be secure? Why do you want security or material possessions? Why do you want to support your dependents? The questions will continue until you point to something that is valuable in itself and not for the sake of something else. It seems that the questions can only stop when you say you want whatever it is because you believe it will make you happy. The questioning stops here because it doesn't seem to make sense to ask why someone wants to be happy.

A discussion of this kind could go off in the direction of questioning whether your belief is right. Will a career as a computer professional make you happy? Will it really bring security? Will security or material possessions, in fact, make you happy? Such discussions center on whether or not you have chosen the correct means to your happiness. However, the point that utilitarians want to make is that any discussion of what you should seek in life, and what is valuable, will not stop until we get to happiness.

It makes no sense, utilitarians argue, to ask why people value happiness. Happiness is the ultimate good. All our actions are directly or indirectly aimed at happiness. It is happiness for which we all strive. Utilitarians seem to believe that this is simply part of our human nature. Human beings are creatures who seek happiness. And, since happiness is the ultimate good, utilitarians believe that morality must be based on creating as much of this good as possible. Thus, all actions should be evaluated in terms of their "utility" for bringing about happiness.

According to utilitarianism, when an individual is faced with a decision about what to do, the individual should consider his or her alternatives, predict the consequences of each alternative, and choose that action which brings about the most good consequences, that is, the most happiness. So, the utilitarian principle provides a decision procedure. When you have to decide what to do, consider the happiness-unhappiness consequences that will result from your various alternatives. The alternative that produces the most overall net happiness (good minus bad) is the right action. To be sure, the right action may be one that brings about some unhappiness, but that is justified if the action also brings about so much happiness that the unhappiness is outweighed, or as long as the action has the least net unhappiness of all the alternatives.

Be careful not to confuse utilitarianism with *egoism*. Egoism is a theory that specifies that one should act so as to bring about the greatest number of good consequences for yourself. What is good is what makes "me" happy or gets me what I want. Utilitarianism does not say that you should maximize your own good. Rather, total happiness is what is at issue. Thus, when you evaluate your alternatives, you have to ask about their effects on the happiness of everyone. This includes effects on you, but your happiness counts the same as the happiness of others. It may turn out to be right for you to do something that will diminish your own happiness because it will bring about a marked increase in overall happiness.

The decision-making process proposed in utilitarianism seems to be at the heart of a good deal of social decision making. That is, legislators and public policy makers seem to seek policies that will produce good consequences, and they often opt for policies that may have some negative consequences but on balance, bring about more good than harm. Cost-benefit or risk-benefit analysis aims at quantifying net good consequences. This involves weighing the potential benefits of a project, such as construction of a new waste disposal plant, against the risks of harm in undertaking the project. It involves calculating and weighing the negative and positive effects of a project in deciding whether to go forward with it. In the case of a waste disposal plant, for example, we look at alternative ways to handle the waste, the various costs and benefits of each alternative, the good and bad effects of locating the plant here or there, and so on. We balance the benefits of the plant against the risk of harm and other negative consequences to all those who will be affected.

Acts versus Rules

As mentioned earlier, there are several formulations of utilitarianism and proponents of various versions disagree on important details. One important and controversial issue of interpretation has to do with whether the focus should be on rules of behavior or individual acts. Utilitarians have recognized that it would be counter to overall happiness if each one of us had to calculate at every moment what all the consequences of every one of our actions would be. Not only is this impractical, because it is time consuming and because sometimes we must act quickly, but often the consequences are impossible to foresee. Thus, there is a need for general rules to guide our actions in ordinary situations.

Rule-utilitarians argue that we ought to adopt rules that, if followed by everyone, would, in the long run, maximize happiness. Take, for example, telling the truth. If individuals regularly told lies, it would be very disruptive. You would never know when to believe what you were told. In the long run, a rule obligating people to tell the truth has enormous beneficial consequences. Thus, “tell the truth” becomes a utilitarian moral rule. “Keep your promises,” and “Don’t reward behavior that causes pain to others,” are also rules that can be justified on utilitarian grounds. According to rule-utilitarianism, if the rule can be justified in terms of the consequences that are brought about from people following it, then individuals ought to follow the rule.

Act-utilitarians put the emphasis on individual actions rather than rules. They believe that even though it may be difficult for us to anticipate the consequences of our actions, that is what we should be trying to do. Take, for example, a case where lying may bring about more happiness than telling the truth. Say you are told by a doctor that tentative test results indicate that your spouse may be terminally ill. You know your spouse well enough to know that this knowledge, at this time, will cause your spouse enormous stress. He or she is already under a good deal of stress because of pressures at work and because someone else in the family is very ill. To tell your spouse the truth about the test results will cause more stress and anxiety, and this stress and anxiety may turn out to be unnecessary if further tests prove that the spouse is not terminally ill. Your spouse asks you what you and the doctor talked about. Should you lie or tell the truth? An act-utilitarian might say that the right thing to do in such a situation is to lie, for little good would come from telling the truth and a good deal of suffering (perhaps unnecessary suffering) will be avoided from lying. A rule-utilitarian would agree that good might result from lying in this one case, but in the long run, if we cannot count on people telling the truth (especially our spouses), more bad than good will come. Think of the anxiety that might arise if spouses routinely lied to one another. Thus, according to rule-utilitarians, we must uphold the rule against lying; it would be wrong to lie.

Act-utilitarianism treats rules simply as “rules of thumb,” general guidelines to be abandoned in situations where it is clear that more happiness will result from breaking them. Rule-utilitarians, on the other hand, take rules to be strict. They justify moral rules in terms of the happiness consequences that result from people following them. If a rule is justified, then an act that violates the rule is wrong.

In either case, it should be clear that the utilitarian principle can be used to formulate a decision procedure for figuring out what you should do in a situation. In fact, many utilitarians propose that the utilitarian principle be used to determine the laws of a society. Laws against stealing, killing, breaking contracts, fraud, and so on can be justified on utilitarian grounds. Utilitarianism is also often used as a principle for evaluating the laws that we have. If a law is not producing good consequences or is producing a mixture of good and bad effects, and we know of another approach that will produce better net effects, then that information provides the grounds for changing the law. Punishment is a good example of a social practice that can be evaluated in terms of its utility. According to utilitarianism, since punishment involves the imposition of pain, if it does not produce sonic good consequences, then it is not justified. Typically utilitarians focus on the deterrent effect of punishment as the good consequence counterbalancing the pain involved.

Earlier I mentioned that utilitarianism might be said to capture part of the idea in relativism. According to utilitarianism, the morally right thing to do in a given situation will depend entirely on the situation. In one situation, it may be right to lie, in another situation in which the circumstances are different, it may be wrong to lie. Even rule-utilitarians must admit that the rule that will produce the most happiness will vary from situation to situation. A simple example would be to suppose a natural environment in which water is scarce. In such a situation, a rule prohibiting individuals from putting water in swimming pools and watering lawns would be justified. The rule would be justified because the alternative would lead to bad consequences. On the other hand, in a natural environment in which water is abundant, such a rule would not be justified.

So, even though utilitarians assert a universal principle, the universal principle has varying implications depending on the situation. This means that utilitarianism is consistent with varying laws and practices at different times or in different places depending on the specific circumstances.

Now that the fundamentals of utilitarianism have been explained, it is worth remembering, once again, that we are engaged in a dialectic. We have developed the idea of utilitarianism; we have made the case for the theory. The theory has been “put on the table,” so to speak. Even though it has been developed only in its most rudimentary form, the theory now needs to be critically scrutinized.

Critique of Utilitarianism

One of the important criticisms of utilitarianism is that when it is applied to certain cases, it seems to go against some of our most strongly held moral intuitions. In particular, it seems to justify imposing enormous burdens on some individuals for the sake of others. According to utilitarianism, every person is to be counted equally. No one person’s unhappiness or happiness is more important than another’s. However, since utilitarians are concerned with the total amount of happiness, we can imagine situations where great overall happiness might result from sacrificing the happiness of a few. Suppose, for example, that having a small number of slaves would create great happiness for a large number of individuals. The individuals who were made slaves would be unhappy, but this would be counterbalanced by significant increases in the happiness of many others. This seems justifiable (if not obligatory) according to utilitarianism. Another more contemporary example would have us imagine a situation in which by killing one person and using all their organs for transplantation, we would be able to save ten lives. Killing one to save ten would seem to maximize good consequences. Critics of utilitarianism argue that since utilitarianism justifies such practices as slavery and killing of the innocent, it has to be wrong. It is, therefore, unacceptable as an account of morality.

In defending the theory from this criticism, some utilitarians argue that utilitarianism does not justify such unsavory practices. Critics, they argue, are forgetting the difference between short-term and long-term consequences. Utilitarianism is concerned with all the consequences and when long-term consequences are taken into account, it becomes clear that such practices as slavery and killing innocent people to use their organs could never be justified. In the long run, such practices have the effect of creating so much fear in people that net happiness is diminished rather than increased. Imagine the fear and anxiety that would prevail in a society in which anyone might at any time be taken as a slave. Or imagine the reluctance of anyone to go to a hospital if there was even a remote possibility that they might be killed if by chance they were there when multiple organs were needed to save lives. The good effects of such practices could never counterbalance these bad effects.

Other utilitarians boldly concede that there are going to be some circumstances in which what seem to be repugnant practices should be accepted because they bring about consequences having a greater net good than would be brought about by other practices, that is, because they are consistent with the principle of utility. So, for example, according to these utilitarians, if there are ever circumstances in which slavery would produce more good than ill, then slavery would be morally acceptable. These utilitarians acknowledge that there may be circumstances in which some people should be sacrificed for the sake of total happiness.

In our dialogue about ethics, it is important to pick up on our strongly held moral intuitions for they are often connected to a moral principle or theory. In the case of utilitarianism, the intuition that slavery is always wrong (or that it is wrong to kill the innocent for the sake of some greater good) points to an alternative moral theory. A concrete case will help us further understand utilitarianism and introduce a different theory, one that captures the moral intuition about the wrongness of slavery and killing the innocent.

Case Illustration

Not long ago, when medical researchers had just succeeded in developing the kidney dialysis machine, a few hospitals acquired a limited number of these expensive machines. Hospitals soon found that the number of patients needing treatment on the machines far exceeded the number of machines they had available or could afford. Decisions had to be made as to who would get access to the machines, and these

were often life-death decisions. In response, some hospitals set up internal review boards composed of medical staff and community representatives. These boards were charged with the task of deciding which patients should get access to the dialysis machines. The medical condition of each patient was taken into account, but the decisions were additionally made on the basis of the personal and social characteristics of each patient: age, job, number of dependents, social usefulness of job, whether the person had a criminal record, and so on. The review committees appeared to be using utilitarian criteria. The resource—kidney dialysis machines—was scarce and they wanted to maximize the benefit (the good consequences) of the use of the machines. Thus, those who were most likely to benefit and to contribute to society in the future would get access. Individuals were given a high ranking for access to the machines if they were doctors (with the potential to save other lives), if they had dependents, if they were young, and so on. Those who were given lower priority or no priority for access to the machines were those who were so ill that they were likely to die even with treatment, those who were older, those who were criminals, those without dependents, and so on.

As the activities of the hospital review boards became known to the public, they were criticized. Critics argued that your value as a person cannot be measured by your value to the community. The review boards were valuing individuals on the basis of their social value and this seemed dangerous. Everyone, it was argued, has value in and of themselves.

The critique of this method for deciding who should live and who should die suggested a principle that is antithetical to utilitarianism. It suggested that each and every person, no matter what their social role or lot in life, has value and should be respected. To treat individuals as if they are a means to some social end seems the utmost in disrespect. And, that is exactly what a policy of allocating scarce resources according to social value does. It says, in effect, that people have value only as means to the betterment of society, and by that criteria some individuals are much more valuable than others.

The critics of distribution of kidney dialysis on the basis of social utility proposed as an alternative that scarce medical resources should be distributed by a lottery. In a lottery, everyone has an equal chance. Everyone counts the same. This, they argued, was the only fair method of distribution.

The kidney dialysis issue is just a microcosm of all medical resources. Doctors, medical equipment, and medical research are expensive and we have a finite amount of money to spend. Hence, lines have to be drawn—on what level of care goes to who, at what stage in their life, and so on. Distributive decisions have to be made.

The important point for our purposes is that the formulation of utilitarianism we have been considering leads to methods of distribution that seem to be unfair or unjust. So while the core idea in utilitarianism seems plausible (i.e., that everyone's happiness or well-being should be counted), utilitarianism does not seem to adequately handle the distribution of benefits and burdens. The criticism of the hospital review boards for distributing access to kidney machines according to social value goes to the heart of this criticism. Critics argue that people are valuable in themselves, not for their contribution to society. They argue that utilitarian programs are often unfair because in maximizing overall good, they impose an unfair burden on some individuals, and as such treat those individuals merely as means to social good.

I will now turn to an ethical theory that articulates the reasoning underlying the critique of utilitarianism. Before doing so, however, it is important to note that the dialectic could go off in a different direction. The debate about utilitarianism is rich and there are many moves that could be made in reformulating the theory and defending it against its critics. It is also important to note that whatever its weaknesses, utilitarianism goes a long way in providing a systematic account of many of our moral notions.

DEONTOLOGICAL THEORIES

In utilitarianism, what makes an action or a rule right or wrong is outside the action; it is the consequences of the action or rule that make it right or wrong. By contrast, deontological theories put the emphasis on the internal character of the act itself.¹ What makes an action right or wrong for deontologists is the principle

¹ The term *deontology* is derived from the Greek words *deon* (duty) and *logos* (science). Etymologically, then, deontology means the science of duty. According to the *Encyclopedia of Philosophy*, its current usage is more specific, referring to an ethical theory which holds that "at least some acts are morally obligatory regardless of their consequences for human weal or woe." (Edwards, 1967)

inherent in the action. If an action is done from a sense of duty, if the principle of the action can be universalized, then the action is right. For example, if I tell the truth (not just because it is convenient for me to do so, but) because I recognize that I must respect the other person, then I act from duty and my action is right. If I tell the truth because I fear getting caught or because I believe I will be rewarded for doing so, then my act is not morally worthy.

I am going to focus here on the theory of Immanuel Kant. If we go back for a moment to the allocation of dialysis machines, Kant's moral theory is applicable because it proposes what is called a categorical imperative specifying that we should never treat human beings merely as means to an end. We should always treat human beings as ends in themselves. Although Kant is not the only deontologist, I will continue to refer to him as I discuss deontology.

The difference between deontological theories and consequentialist theories was illustrated in the discussion of allocation of dialysis machines. Deontologists say that individuals are valuable in themselves, not because of their social value. Utilitarianism is criticized because it appears to tolerate sacrificing some people for the sake of others. In utilitarianism, right and wrong are dependent on the consequences and therefore vary with the circumstances. By contrast, deontological theories assert that there are some actions that are always wrong, no matter what the consequences. A good example of this is killing. Even though we can imagine situations in which intentionally killing one person may save the lives of many others, deontologists insist that intentional killing is always wrong. Killing is wrong even in extreme situations because it means using the person merely as a means and does not treat the human being as valuable in and of himself. Deontologists do often recognize self-defense and other special circumstances as excusing killing, but these are cases when, it is argued, the killing is not exactly intentional. (The person attacks me. I would not, otherwise, aim at harm to the person, but I have no other choice but to defend myself.)

At the heart of deontological theory is an idea about what it means to be a person, and this is connected to the idea of moral agency, Charles Fried (1978) put the point as follows:

[T]he substantive contents of the norms of right and wrong express the value of persons, of respect for personality. What we may not do to each other, the things which are wrong, are precisely those forms of personal interaction which deny to our victim the status of a freely choosing, rationally valuing, specially efficacious person, the special status of moral personality. (pp. 28–29)

According to deontologists, the utilitarians go wrong when they fix on happiness as the highest good. Deontologists point out that happiness cannot be the highest good for humans. The fact that we are rational beings, capable of reasoning about what we want to do and then deciding and acting, suggests that our end (our highest good) is something other than happiness. Humans differ from all other things in the world insofar as we have the capacity for rationality. The behavior of other things is determined simply by laws of nature. Plants turn toward the sun because of photosynthesis. They don't think and decide which way they will turn. Physical objects fall by the law of gravity. Water boils when it reaches a certain temperature. In contrast, human beings are not entirely determined by laws of nature. We have the capacity to legislate for ourselves. We decide how we will behave. As Kant describes this, it is the difference between acting in accordance with law (plants and stones do) and acting in accordance with the *conception* of law.

The capacity for rational decision making is the most important feature of human beings. Each of us has this capacity; each of us can make choices, choices about what we will do, and what kind of persons we will become. No one else can or should make these choices for us. Moreover, we should recognize this capacity in others.

Notice that it makes good sense that our rationality is connected with morality, for we could not be moral beings at all unless we had this rational capacity. We do not think of plants or fish or dogs and cats as moral beings precisely because they do not have the capacity to reason about their actions. We are moral beings because we are rational beings, that is, because we have the capacity to give ourselves rules (laws) and follow them.

Where utilitarians note that all humans seek happiness, deontologists emphasize that humans are creatures with goals who engage in activities directed toward achieving these goals (ends), and that they use their rationality to formulate their goals and figure out what kind of life to live. In a sense, deontologists pull back from fixing on any particular value as structuring morality and instead ground morality in the capacity of each individual to organize his or her own life, make choices, and engage in activities to realize

their self-chosen life plans. What morality requires is that we respect each of these beings as valuable in themselves and refrain from valuing them only insofar as they fit into our own life plans.

As mentioned before, Kant put forward what he called the *categorical imperative*. While there are several versions of it, I will focus on the second version which goes as follows: *Never treat another human being merely as a means but always as an end*. This general rule is derived from the idea that persons are moral beings because they are rational, efficacious beings. Because we each have the capacity to think and decide and act for ourselves, we should each be treated with respect, that is with recognition of this capacity.

Note the “merely” in the categorical imperative. Deontologists do not insist that we never use another person as a means to an end, only that we never “merely” use them in this way. For example, if I own a company and hire employees to work in my company, I might be thought of as using those employees as a means to my end (i.e., the success of my business). This, however, is not wrong if I promise to pay a fair wage in exchange for work and the employees agree to work for me. I thereby respect their ability to choose for themselves. What would be wrong would be to take them as slaves and make them work for me. It would also be wrong to pay them so little that they must borrow from me and remain always in on, debt. This would be exploitation. This would show disregard for the value of each person as a “freely choosing, rationally valuing, specially efficacious person.” Similarly, it would be wrong for me to lie to employees about the conditions of their work. Suppose, for example, that while working in my plant, employees will be exposed to dangerous, cancer-causing chemicals. I know this but don’t tell the employees because I am afraid they will quit. In not being forthcoming with this information, I am, in effect, manipulating the employees to serve my ends. I am not recognizing them as beings of value with their own life-plans and the capacity to choose how they will live their lives.

Case Illustration

Though utilitarianism and Kantian theory were contrasted in the case illustration about allocation of scarce medical resources, another case will clarify even more. Consider a case involving computers. Suppose a professor of sociology undertakes research on attitudes toward sex and sexual behavior among high school students. Among other things, she interviews hundreds of high school students concerning their attitudes and behavior. She knows that the students will never give her information unless she guarantees them confidentiality, so before doing the interviews, she promises each student that she alone will have access to the raw interview data, and that all publishable results will be reported in statistical form. Thus, it would be impossible to identify information from individual students.

Suppose, however, that it is now time to analyze the interview data and she realizes that it will be much easier to put the data into a computer and use the computer to do the analysis. To assure the confidentiality she promised, the professor will have to code the data so that names do not appear in the database and will have to make an effort to secure the data. She has hired graduate students to assist her and she wonders whether she should let the graduate students handle the raw data. Should she allow the graduate assistants to code and process the data?

At first glance it would seem that from a consequentialist point of view, the professor should weigh the good that will come from the research, and from doing it quickly on a computer, against the possible harm to herself and her subjects if information is leaked. The research may provide important information to people working with high school students and may help her career prosper. Still, the advantage of doing it quickly may be slight. She must worry about the effect of a leak of information on the students. Also, since she has explicitly promised confidentiality to the student-subjects, she has to worry about the effects on her credibility as a social researcher and on social science research in general if she breaks her promise. That is, her subjects and many others may be reluctant in the future to trust her and other social scientists if she breaks the promise and they find out.

Thus, there seem good reasons to say that from a consequentialist point of view the professor should not violate her promise of confidentiality. Fortunately, there are ways to code data before putting it into the computer or turning it over to her graduate students. She must do the coding herself and keep the key to individual names confidential.

This is how a consequentialist might analyze the situation. Interestingly, a deontologist might well come to the same conclusion though the reasoning would be quite different. The sociologist is doing a study that will advance human knowledge and, no doubt, further her career. There is nothing wrong with

this as long as it does not violate the categorical imperative. The question here is whether she is treating her subjects merely as means to knowledge and her own advancement, or whether she is truly recognizing those subjects as ends in themselves. Were the sociologist to ignore her promise of confidentiality to the students, she would not be treating each subject as an end. Each student made a choice based on her pledge of confidentiality. She would be treating them merely as means if she were to break her promise when it suited her. Thus, out of respect for the subjects, the sociologist must code the data herself so as to maintain the promised confidentiality.

The two theories do not, then, come to very different conclusions in this case. However, the analysis is very different in that the reasons given for coming to the conclusion are very different. In other cases, these theories lead to dramatically different conclusions.

Our dialogue on utilitarianism and Kantian theory could continue. I have presented only the bare bones of each theory. However, in the interest of getting to the issues surrounding computers, we must move on and put a few more important concepts and theories “on the table.”

RIGHTS

So far, very little has been said about rights though we often use the language of rights when discussing moral issues. “You have no right to tell me what to do.” “I have a right to do that.” Ethicists often associate rights with deontological theories. The categorical imperative requires that each person be treated as an end in himself or herself, and it is possible to express this idea by saying that individuals have “a right to” the kind of treatment that is implied in being treated as an end. The idea that each individual must be respected as valuable in himself or herself implies that we each have rights not to be interfered with in certain ways, for example, not to be killed or enslaved, to be given freedom to make decisions about our own lives, and so on.

An important distinction that philosophers often make here is between negative rights and positive rights. Negative rights are rights that require restraint by others. For example, my right not to be killed requires that others refrain from killing me. It does not, however, require that others take positive action to keep me alive. Positive rights, on the other hand, imply that others have a duty to do something to or for the right holder. So, if we say that I have a positive right to life, this implies not just that others must refrain from killing me, but that they must do such things as feed me if I am starving, give me medical treatment if I am sick, swim out and save me if I am drowning, and so on. As you can see, the difference between negative and positive rights is quite significant.

Positive rights are more controversial than negative rights because they have implications that are counter-intuitive. If every person has a positive right to life, this seems to imply that each and every one of us has a duty to do whatever is necessary to keep all people alive. This would seem to suggest that, among other things, it is our duty to give away any excess wealth that we have to feed and care for those who are starving or suffering from malnutrition. It also seems to imply that we have a duty to supply extraordinary lifesaving treatment for all those who are dying. In response to these implications, some philosophers have argued that individuals have only negative rights.

While, as I said earlier, rights are often associated with deontological theories, it is important to note that rights can be derived from other theories as well. For example, we can argue for the recognition of a right to property on utilitarian grounds. Suppose we ask why individuals should be allowed to have private property in general and, in particular, why they should be allowed to own computer software. Utilitarians would argue for private ownership of software on grounds that much more and better software will be created if individuals are allowed to own (and then license or sell) it. Thus, they argue that individuals should have a legal right to ownership in software because of the beneficial consequences of acknowledging such a right.

Another important thing to remember about rights is the distinction between legal and moral (or natural or human) rights. Legal rights are rights that are created by law. Moral, natural, or human rights are claims independent of law. Such claims are usually embedded in a moral theory or a theory of human nature.

The utilitarian argument is an argument for creating or recognizing a legal right; it is not an argument to the effect that human beings have a natural right, for example, to own what they create. In Chapter 6 we will focus on property rights in computer software and there we will explore both natural and utilitarian property rights.

Rights and Social Contract Theories

Rights are deeply rooted in the tradition of social contract theories. In this tradition the idea of a social contract (between individuals, or between individuals and government) is hypothesized to explain and justify the obligations that human beings have to one another. Many of these theories imagine human beings in a state of nature and then show that reason would lead individuals in such a state to agree to live according to certain rules, or to give power to a government to enforce certain rules. The depiction of a state of nature in which human beings are in a state of insecurity and uncertainty is used to suggest what human nature is like and to show that human nature necessitates government. That is, in such a state any rational human beings would agree (make a contract) to join forces with others even though this involves giving up some of their natural freedom. The agreement (the social contract) creates obligations and these are the basis of more obligation.

An argument of this kind is made by several social contract theorists and each specifies the nature and limits of our obligations differently. One important difference, for example, is in whether morality exists prior to the social contract. Hobbes argues that there is no justice or injustice in a state of nature; humans are at war with out another and each individual must do what they must to preserve themselves. Locke, on the other hand, specifies a natural form of justice in the state of nature. Human beings have rights in the state of nature and others can treat individuals unjustly. Government is necessary to insure that natural justice is implemented properly because without government, there is no certainty that punishments will be distributed justly.

Rawlsian justice

In 1971, John Rawls, a professor at Harvard University, introduced a new version of social contract theory (though some argue it is not a social contract theory in the traditional sense). Rawls introduced the theory in a book entitled simply *A Theory of Justice*. The theory may well be one of the most influential moral theories of the twentieth century, for not only did it generate an enormous amount of attention, it influenced discussion among economists, social scientists, and public policy makers.

Rawls was primarily interested in questions of distributive justice. In the tradition of a social contract theorist, he tries to understand what sort of contract between individuals would be just. Rawls recognizes that we can't arrive at an account of justice and the fairness of social arrangements by reasoning about what rules particular individuals would agree to. He understands that individuals are self-interested and therefore will be influenced by their own experiences and their own situation when they think about fair arrangements. Thus, if some group of us were to get together in something like a state of nature (suppose a group is stranded on an island or a nuclear war occurs and only a few survive), the rules we would agree to would not necessarily be a just system. It would not necessarily exemplify justice.

The problem is that we would each want rules that would favor us. Smart people would want rules that favored intelligence. Strong people would want a system that rewarded strength. Women would not want rules that were biased against women, and so on. The point is that there is no reason to believe that the outcome of a negotiation in which people expressed their preferences would result in rules of justice and just institutions. In this sense, Rawls believes that justice has to be blind in a certain way.

Rawls specifies, therefore, that in order to get at justice, we have to imagine that the individuals who get together to decide on the rules for society are behind a veil of ignorance. The veil of ignorance is such that individuals do not know what characteristics they will have. They do not know whether they will be male or female, black or white, highly intelligent or moderately intelligent or retarded, physically strong or in ill-health, musically talented, successful at business, indigent and so on.

At the same time, these individuals would be rational and self-interested and would know something about human nature and human psychology. In a sense, what Rawls is suggesting here is that we have to imagine *generic* human beings. They have abstract features that human beings generally have (i.e., they are rational and self-interested). And, they have background knowledge (i.e., general knowledge of how humans behave and interact and how they are affected in various ways).

According to Rawls, justice is what individuals would choose in such a situation. Notice that what he has done, in a certain sense, is eliminate bias in the original position. Once a society gets started, once particular individuals have characteristics, their views on what is fair are tainted. They cannot be objective. So, justice, according to Rawls is what people would choose in the original position where they are rational

and self-interested, informed about human nature and psychology but behind a veil of ignorance with regard to their own characteristics. Rawls argues that individuals in the original position would agree to two rules. These are the rules of justice and they are “rules of rules” in the sense that they are general principles constraining the formulation of specific rules. The rules of justice are:

1. Each person should have an equal right to the most extensive basic liberty compatible with a similar liberty for others.
2. Social and economic inequalities should be arranged so that they are both (a) reasonably expected to be to everyone’s advantage and (b) attached to positions and offices open to all.

These general principles assure that no matter where an individual ends up in the lottery of life (in which characteristics of intelligence, talents, physical abilities, and so on, are distributed), he or she would have liberty and opportunity. He or she would have a fair shot at a decent life.

While Rawls’ account of justice has met with criticism, it goes a long way in providing a framework for envisioning and critiquing just institutions. This discussion of Rawls is extremely abbreviated as were the accounts of Kant and utilitarianism. Perhaps the most important thing to keep in mind as we proceed to the issues surrounding computer and information technology is that rights-claims and claims about justice and fairness generally presume a much more complicated set of claims. Such claims should never be accepted as primitive truths. The underlying argument and embedded assumptions should be uncovered and critically examined.

VIRTUE ETHICS

Before moving on to the ethical issues surrounding computer and information technology, one other tradition in ethical theory should be mentioned. In recent years, interest has arisen in resurrecting the tradition of virtue ethics, a tradition going all the way back to Plato and Aristotle. These ancient Greek philosophers pursued the question: What is a good person? What are the virtues associated with being a good person? For the Greeks *virtue* meant excellence, and ethics was concerned with excellences of human character. A person possessing such qualities exhibited the excellences of human good. To have these qualities is to function well as a human being.

The list of possible virtues is long and there is no general agreement on which are most important, but the possibilities include courage, benevolence, generosity, honesty, tolerance, and self-control. Virtue theorists try to identify the list of virtues and to give an account of each—What is courage? What is honesty? They also give an account of why the virtues are important.

Virtue theory seems to fill a gap left by other theories we considered, because it addresses the question of moral character, while the other theories focused primarily on action and decision making. What sort of character should we be trying to develop in ourselves and in our children. We look to moral heroes, for example, as exemplars of moral virtue. Why do we admire such people? What is it about their character and their motivation that are worthy of our admiration?

Virtue theory might be brought into the discussion of computer technology and ethics at any number of points. The most obvious is, perhaps, the discussion of professional ethics, where we want to think about the characteristics of a good computer professional. Good computer professionals will, perhaps, exhibit honesty in dealing with clients and the public. They should exhibit courage when faced with situations in which they are being pressured to do something illegal or act counter to public safety. A virtue approach would focus on these characteristics and more, emphasizing the virtues of a good computer professional.

INDIVIDUAL AND SOCIAL POLICY ETHICS

One final distinction will be helpful. In examining problems or issues, it is important to distinguish levels of analysis, in particular that between macro and micro level issues or approaches. One can approach a problem from the point of view of social practices and public policy, or from the point of view of individual choice. Macro level problems are problems that arise for groups of people, a community, a state, a country. At this level of analysis, what is sought is a solution in the form of a law or policy that specifies

how people in that group or society ought to behave, what the rules of that group ought to be. When we ask the following questions, we are asking macro level questions: Should the United States grant software creators a legal right to own software? Should software engineers be held liable for errors in the software they design? Should companies be allowed to electronically monitor their employees?

On the other hand, micro level questions focus on individuals (in the presence or absence of law or policy). Should I make a copy of this piece of software? Should I lie to my friend? Should I work on a project making military weapons? Sometimes these types of questions can be answered simply by referring to a rule established at the macro level. For example, legally I can make a back-up copy of software that I buy, but I shouldn't make a copy and give it to my friend. Other times, there may be no macro level rule or the macro level rule may be vague or an individual may think the macro level rule is unfair. In these cases, individuals must make decisions for themselves about what they ought to do.

The theories just discussed inform both approaches, but in somewhat different ways, so it is important to be clear on which type of question you are asking or answering.

CONCLUSION

While the focus of our attention will now shift to the ethical issues surrounding computer and information technology, the deep questions and general concerns of ethical theories will continue to haunt us. The dialogue is ongoing. Remember that science is never done. In both science and ethics, we look for reasons supporting the claims that we make, and we tell stories (develop arguments and theories) to answer our questions. We tell stories about why the physical world is the way it is, why human beings behave the way they do, why lying and killing are wrong, and so on. The stories we tell often get better and better over time. They get broader (more encompassing) and richer, sometimes more elegant, sometimes allowing us to see new things we never noticed before. The stories generally lead to new questions. So it is with ethics as well as science.

Computer ethics should be undertaken with this in mind, for the task of computer ethics involves working with traditional moral concepts and theories, and extending them to situations with somewhat new features. The activity brings insight into the situations arising from use of computer and information technology, and it may also bring new insights into ethical concepts and theories.

The Altered Nature of Human Action

by
Hans Jonas

Translated by Hans Jonas with the collaboration of David Herr

Chapter 1 from
The Imperative of Responsibility: In Search of an Ethics for the Technological Age
The University of Chicago Press, 1984

All previous ethics—whether in the form of issuing direct enjoinders to do and not to do certain things, or in the form of defining principles for such enjoinders, or in the form of establishing the ground of obligation for obeying such principles—had these interconnected tacit premises in common: that the human condition, determined by the nature of man and the nature of things, was given once for all; that the human good on that basis was readily determinable; and that the range of human action and therefore responsibility was narrowly circumscribed. It will be the burden of the present argument to show that these premises no longer hold, and to reflect on the meaning of this fact for our moral condition. More specifically, it will be my contention that with certain developments of our powers the nature of human action has changed, and, since ethics is concerned with action, it should follow that the changed nature of human action calls for a change in ethics as well: this not merely in the sense that new objects of action have added to the case material on which received rules of conduct are to be applied, but in the more radical sense that the qualitatively novel nature of certain of our actions has opened up a whole new dimension of ethical relevance for which there is no precedent in the standards and canons of traditional ethics.

The novel powers I have in mind are, of course, those of modern technology. My first point, accordingly, is to ask how this technology affects the nature of our acting, in what ways it makes acting under its dominion different from what it has been through the ages. Since throughout those ages man was never without technology, the question involves the human difference of modern from previous technology.

I. The Example of Antiquity

Let us start with an ancient voice on man's powers and deeds which in an archetypal sense itself strikes, as it were, a technological note—the famous Chorus from Sophocles' *Antigone*.

Many the wonders but nothing more wondrous than man.
This thing crosses the sea in the winter's storm,
making his path through the roaring waves.
And she, the greatest of gods, the Earth—
deathless she is, and unwearied—he wears her away as the ploughs go up and down from year to
year
and his mules turn up the soil.

The tribes of the lighthearted birds he ensnares, and the races
of all the wild beasts and the salty brood of the sea,
with the twisted mesh of his nets, he leads captive, this clever man.
He controls with craft the beasts of the open air,
who roam the hills. The horse with his shaggy mane
he holds and harnesses, yoked about the neck,
and the strong bull of the mountain.

Speech and thought like the wind and the feelings that make the town,
he has taught himself, and shelter against the cold,

refuge from rain. Ever resourceful is he.
 He faces no future helpless. Only against death
 shall he call for aid in vain. But from baffling maladies
 has he contrived escape.

Clever beyond all dreams the inventive craft that he has
 which may drive him one time or another to well or ill.
 When he honors the laws of the land and the gods' sworn right
 high indeed is his city; but stateless the man
 who dares to do what is shameful.

[Lines 334—370]

1. *Man and Nature*

This awestruck homage to man's powers tells of his violent and violating irruption into the cosmic order, the self-assertive invasion of nature's various domains by his restless cleverness; but also of his building—through the self-taught powers of speech and thought and social sentiment—the home for his very humanity, the artifact of the city. The raping of nature and the civilizing of man go hand in hand. Both are in defiance of the elements, the one by venturing into them and overpowering their creatures, the other by securing an enclave against them in the shelter of the city and its laws. Man is the maker of his life *qua* human, bending circumstances to his will and needs, and except against death he is never helpless.

Yet there is a subdued and even anxious quality about this appraisal of the marvel that is man, and nobody can mistake it for immodest bragging. Unspoken, but self-evident for those times, is the pervading knowledge behind it all that, for all his boundless resourcefulness, man is still small by the measure of the elements: precisely this makes his sallies into them so daring and allows those elements to tolerate his forwardness. Making free with the denizens of land and sea and air, he yet leaves the encompassing nature of those elements unchanged, and their generative powers undiminished. He cannot harm them by carving out his little kingdom from theirs. They last, while his schemes have their short-lived way. Much as he harries Earth, the greatest of gods, year after year with his plough—she is ageless and unwearied; her enduring patience he must and can trust, and to her cycle he must conform. And just as ageless is the sea. With all his netting of the salty brood, the spawning ocean is inexhaustible. Nor is it hurt by the plying of ships, nor sullied by what is jettisoned into its deeps. And no matter how many illnesses he contrives to cure, mortality does not bow to his cunning.

All this holds because before our time man's inroads into nature, as seen by himself, were essentially superficial and powerless to upset its appointed balance. (Hindsight reveals that they were not always so harmless in reality.) Nor is there a hint, in the *Antigone* chorus or anywhere else, that this is only a beginning and that greater things of artifice and power are yet to come—that man is embarked on an endless course of conquest. He had gone thus far in reducing necessity, had learned by his wits to wrest that much from it for the humanity of his life, and reflecting upon this, he was overcome by awe at his own boldness.

2. *The Man-Made Island of the "City"*

The room he has thus made was filled by the city of men—meant to enclose, and not to expand—and thereby a new balance was struck within the larger balance of the whole. All the good or ill to which man's inventive craft may drive him one time or another is inside the human enclave and does not touch the nature of things.

The immunity of the whole, untroubled in its depth by the importunities of man, that is, the essential immutability of Nature as the cosmic order, was indeed the backdrop to all of mortal man's enterprises, including his intrusions into that order itself. Man's life was played out between the abiding and the changing: the abiding was Nature, the changing his own works. The greatest of these works was the city, and on it he could confer some measure of abiding by the laws he made for it and undertook to honor. But no long-range certainty pertained to this contrived continuity. As a vulnerable artifact, the cultural construct can grow slack or go astray. Not even within its artificial space, with all the freedom it gives to man's determination of self, can the arbitrary ever supersede the basic terms of his being. The very inconstancy of human fortunes assures the constancy of the human condition. Chance and luck and folly, the great equalizers in human affairs, act like an entropy of sorts and make all definite designs in the long run revert to the perennial norm. Cities rise and fall, rules come and go, families prosper and decline; no change is there to stay, and in the end, with all the temporary deflections balancing each other out, the state of man is

as it always was. So here, too, in his very own artifact, the social world, man's control is small and his abiding nature prevails.

Still, this citadel of his own making, clearly set off from the rest of things and entrusted to him, was the whole and sole domain of man's responsible action. Nature was not an object of human responsibility—she taking care of herself and, with some coaxing and worrying, also of man: not ethics, only cleverness applied to her. But in the city, the social work of art, where men deal with men, cleverness must be wedded to morality, for this is the soul of its being. It is in this intrahuman frame, then, that all traditional ethics dwells, and it matches the size of action delimited by this frame.

II. Characteristics of Previous Ethics

Let us extract from the above those characteristics of human action which are relevant for a comparison with the state of things today.

1. All dealing with the nonhuman world, that is, the whole realm of *techne* (with the exception of medicine), was ethically neutral—in respect both of the object and the subject of such action: in respect of the object, because it impinged but little on the self-sustaining nature of things and thus raised no question of permanent injury to the integrity of its object, the natural order as a whole; and in respect of the agent subject it was ethically neutral because *techne* as an activity conceived itself as a determinate tribute to necessity and not as an indefinite, self-validating advance to mankind's major goal, claiming in its pursuit man's ultimate effort and concern. The real vocation of man lay elsewhere. In brief, action on nonhuman things did not constitute a sphere of authentic ethical significance.

2. Ethical significance belonged to the direct dealing of man with man, including the dealing with himself: all traditional ethics is *anthropocentric*.

3. For action in this domain, the entity "man" and his basic condition was considered constant in essence and not itself an object of reshaping *techne*.

4. The good and evil about which action had to care lay close to the act, either in the praxis itself or in its immediate reach, and were not matters for remote planning. This proximity of ends pertained to time as well as space. The effective range of action was small, the time span of foresight, goal-setting, and accountability was short, control of circumstances limited. Proper conduct had its immediate criteria and almost immediate consummation. The long run of consequences beyond was left to chance, fate, or providence. Ethics accordingly was of the here and now, of occasions as they arise between men, of the recurrent, typical situations of private and public life. The good man was the one who met these contingencies with virtue and wisdom, cultivating these powers in himself, and for the rest resigning himself to the unknown.

All enjoinders and maxims of traditional ethics, materially different as they may be, show this confinement to the immediate setting of the action. "Love thy neighbor as thyself"; "Do unto others as you would wish them to do unto you"; "Instruct your child in the way of truth"; "Strive for excellence by developing and actualizing the best potentialities of your being *qua* man"; "Subordinate your individual good to the common good"; "Never treat your fellow man as a means only but always also as an end in himself"—and so on. Note that in all these maxims the agent and the "other" of his action are sharers of a common present. It is those who are alive now and in some relationship with me who have a claim on my conduct as it affects them by deed or omission. The ethical universe is composed of contemporaries, and its horizon to the future is confined by the foreseeable span of their lives. Similarly confined is its horizon of place, within which the agent and the other meet as neighbor, friend, or foe, as superior and subordinate, weaker and stronger, and in all the other roles in which humans interact with one another. To this proximate range of action all morality was geared.

It follows that the *knowledge* that is required—besides the moral will to assure the morality of action fitted these limited terms: it was not the knowledge of the scientist or the expert, but knowledge of a kind readily available to all men of good will. Kant went so far as to say that "human reason can, in matters of morality, be easily brought to a high degree of accuracy and completeness even in the most ordinary intelligence";¹ that "there is no need of science or philosophy for knowing what man has to do in order to be honest and good, and indeed to be wise and virtuous. . . . [Ordinary intelligence] can have as good hope of hitting the mark as any philosopher can promise himself";² and again: "I need no elaborate acuteness to find out what I have to do so that my willing be morally good. Inexperienced regarding the course of the world, unable to anticipate all the contingencies that happen in it," I can yet know how to act in accordance

with the moral law.³

Not every thinker in ethics, it is true, went so far in discounting the cognitive side of moral action. But even when it received much greater emphasis, as in Aristotle, where the discernment of the situation and what is fitting for it makes considerable demands on experience and judgment, such knowledge has nothing to do with the science of things. It implies, of course, a general conception of the human good as such, a conception predicated on the presumed invariables of man's nature and condition, which may or may not find expression in a theory of its own. But its translation, into practice requires a knowledge of the here and now, and this is entirely nontheoretical. This "knowledge" proper to virtue (of the "where, when, to whom, and how") stays with the immediate issue, in whose defined context the action as the agent's own takes its course and within which it terminates. The good or bad of the action is wholly decided within that short-term context. Its authorship is unquestioned, and its moral quality shines forth from it, visible to its witnesses. No one was held responsible for the unintended later effects of his well-intentioned, well-considered, and well-performed act. The short arm of human power did not call for a long arm of predictive knowledge; the shortness of the one is as little culpable as that of the other. Precisely because the human good, known in its generality, is the same for all time, its realization or violation takes place at each time, and its complete locus is always the present.

III. New Dimensions of Responsibility

All this has decisively changed. Modern technology has introduced actions of such novel scale, objects, and consequences that the framework of former ethics can no longer contain them. The *Antigone* chorus on the *deinotes*, the wondrous power, of man would have to read differently now; and its admonition to the individual to honor the laws of the land would no longer be enough. The gods, too, whose venerable right could check the headlong rush of human action, are long gone. To be sure, the old prescriptions of the "neighbor" ethics—of justice, charity, honesty, and so on—still hold in their intimate immediacy for the nearest, day-by-day sphere of human interaction. But this sphere is overshadowed by a growing realm of collective action where doer, deed, and effect are no longer the same as they were in the proximate sphere, and which by the enormity of its powers forces upon ethics a new dimension of responsibility never dreamed of before.

1. *The Vulnerability of Nature*

Take, for instance, as the first major change in the inherited picture, the critical vulnerability of nature to man's technological intervention—unsuspected before it began to show itself in damage already done. This discovery, whose shock led to the concept and nascent science of ecology, alters the very concept of ourselves as a causal agency in the larger scheme of things. It brings to light, through the effects, that the nature of human action has *de facto* changed, and that an object of an entirely new order—no less than the whole biosphere of the planet—has been added to what we must be responsible for because of our power over it. And of what surpassing importance an object, dwarfing all previous objects of active man! Nature as a human responsibility is surely a novum to be pondered in ethical theory. What kind of obligation is operative in it? Is it more than a utilitarian concern? Is it just prudence that bids us not to kill the goose that lays the golden eggs, or saw off the branch on which we sit? But the "we" who sit here and who may fall into the abyss—who is it? And what is *my* interest in its sitting or falling?

Insofar as it is the fate of *man*, as affected by the condition of nature, which makes our concern about the preservation of nature a *moral* concern, such concern admittedly still retains the anthropocentric focus of all classical ethics. Even so, the difference is great. The containment of nearness and contemporaneity is gone, swept away by the spatial spread and time span of the cause-effect trains which technological practice sets afoot, even when undertaken for proximate ends. Their irreversibility conjoined to their aggregate magnitude injects another novel factor into the moral equation. Add to this their cumulative character: their effects keep adding themselves to one another, with the result that the situation for later subjects and their choices of action will be progressively different from that of the initial agent and ever more the fated product of what was done before. All traditional ethics reckoned only with noncumulative behavior.⁴ The basic situation between persons, where virtue must prove and vice expose itself, remains always the same, and every deed begins afresh from this basis. The recurring occasions which pose their appropriate alternatives for human conduct—courage or cowardice, moderation or excess, truth or mendacity, and so on—each time reinstate the primordial conditions from which action takes off.

These were never superseded, and thus moral actions were largely “typical,” that is, conforming to precedent. In contrast with this, the cumulative self-propagation of the technological change of the world constantly overtakes the conditions of its contributing acts and moves through none but unprecedented situations, for which the lessons of experience are powerless. And not even content with changing its beginning to the point of unrecognizability, the cumulation as such may consume the basis of the whole series, the very condition of, itself. All this would have to be cointended in the will of the single action if this is to be a morally responsible one.

2. *The New Role of Knowledge in Morality*

Knowledge, under these circumstances, becomes a prime duty beyond anything claimed for it heretofore, and the knowledge must be commensurate with the causal scale of our action. The fact that it cannot really be thus commensurate, that is, that the predictive knowledge falls behind the technical knowledge that nourishes our power to act, itself assumes ethical importance. The gap between the ability to foretell and the power to act creates a novel moral problem. With the latter so superior to the former, recognition of ignorance becomes the obverse of the duty to know and thus part of the ethics that must govern the evermore necessary self-policing of our outsized might. No previous ethics had to consider the global condition of human life and the far-off future, even existence, of the race. These now being an issue demands, in brief, a new conception of duties and rights, for which previous ethics and metaphysics provide not even the principles, let alone a ready doctrine.

3. *Has Nature “Rights” Also?*

And what if the new kind of human action would mean that more than the interest of man alone is to be considered—that our duty extends farther, and the anthropocentric confinement of former ethics no longer holds? It is at least not senseless anymore to ask whether the condition of extrahuman nature, the biosphere as a whole and in its parts, now subject to our power, has become a human trust and has something of a moral claim on us not only for our ulterior sake but for its own and in its own right. If this were the case it would require quite some rethinking in basic principles of ethics. It would mean to seek not only the human good but also the good of things extrahuman, that is, to extend the recognition of “ends in themselves” beyond the sphere of man and make the human good include the care for them. No previous ethics (outside of religion) has prepared us for such a role of stewardship—and the dominant, scientific view of *Nature* has prepared us even less. Indeed, that view emphatically denies us all conceptual means to think of Nature as something to be honored, having reduced it to the indifference of necessity and accident, and divested it of any dignity of ends. But still, a silent plea for sparing its integrity seems to issue from the threatened plenitude of the living world. Should we heed this plea, should we recognize its claim as morally binding because sanctioned by the nature of things, or dismiss it as a mere sentiment on our part, which we may indulge as far as we wish and can afford to do? If the former, it would (if taken seriously in its theoretical implications) push the necessary rethinking beyond the doctrine of action, that is, ethics, into the doctrine of being, that is, metaphysics, in which all ethics must ultimately be grounded. On this speculative subject I will say no more here than that we should keep ourselves open to the thought that natural science may not tell the whole story about Nature.

IV. Technology as the “Calling” of Mankind

1. *Homo Faber over Homo Sapiens*

Returning to strictly intrahuman considerations, there is another ethical aspect to the growth of *techne* as a pursuit beyond the pragmatically limited terms of former times. Then, so we found, *techne* was a measured tribute to necessity, not the road to mankind’s chosen goal—a means with a finite measure of adequacy to well-defined proximate ends. Now, *techne* in the form of modern technology has turned into an infinite forward-thrust of the race, its most significant enterprise, in whose permanent, self-transcending advance to ever greater things the vocation of man tends to be seen, and whose success of maximal control over things and himself appears as the consummation of his destiny. Thus the triumph of *homo faber* over his external object means also his triumph in the internal constitution of *homo sapiens*, of whom he used to be a subsidiary part. In other words, technology, apart from its objective works, assumes ethical significance by the central place it now occupies in human purpose. Its cumulative creation, the expanding artificial environment, continuously reinforces the particular powers in man that created it, by compelling

their unceasing inventive employment in its management and further advance, and by rewarding them with additional success—which only adds to the relentless claim. This positive feedback of functional necessity and reward—in whose dynamics pride of achievement must not be forgotten—assures the growing ascendancy of one side of man's nature over all the others, and inevitably at their expense. If nothing succeeds like success, nothing also entraps like success. Outshining in prestige and starving in resources whatever else belongs to the fullness of man, the expansion of his power is accompanied by a contraction of his self-conception and being. In the image he entertains of himself—the programmatic idea which determines his actual being as much as it reflects it—man now is evermore the maker of what he has made and the doer of what he can do, and most of all the preparer of what he will be able to do next. But who is “he”? Not you or I: it is the aggregate, not the individual doer or deed that matters here; and the indefinite future, rather than the contemporary context of the action, constitutes the relevant horizon of responsibility. This requires imperatives of a new sort. If the realm of making has invaded the space of essential action, then morality must invade the realm of making, from which it has formerly stayed aloof, and must do so in the form of public policy. Public policy has never had to deal before with issues of such inclusiveness and such lengths of anticipation. In fact, the changed nature of human action changes the very nature of politics.

2. *The Universal City as a Second Nature*

For the boundary between “city” and “nature” has been obliterated: the city of men, once an enclave in the nonhuman world, spreads over the whole of terrestrial nature and usurps its place. The difference between the artificial and the natural has vanished, the natural is swallowed up in the sphere of the artificial, and at the same time the total artifact (the works of man that have become “the world” and as such envelop their makers) generates a “nature” of its own, that is, a necessity with which human freedom has to cope in an entirely new sense.

Once it could be said *Fiat justitia, pereat mundus*, “Let justice be done, and may the world perish”—where “world,” of course, meant the renewable enclave in the imperishable whole. Not even rhetorically can the like be said anymore when the perishing of the whole through the doings of man—be they just or unjust—has become a real possibility. Issues never legislated come into the purview of the laws which the total city must give itself so that there will be a world for the generations of man

3. *Man's Presence in the World as an Imperative*

That there *ought* to be through all future time such a world fit for human habitation, and that it ought in all future time to be inhabited by a mankind worthy of the human name, will be readily affirmed as a general axiom or a persuasive desirability of speculative imagination (as persuasive and undemonstrable as the proposition that there being a world at all is “better” than there being none): but as a *moral* proposition, namely, a practical *obligation* toward the posterity of a distant future, and a principle of decision in present action, it is quite different from the imperatives of the previous ethics of contemporaneity; and it has entered the moral scene only with our novel powers and range of prescience.

The *presence of man* in the world had been a first and unquestionable given, from which all idea of obligation in human conduct started out. Now it has itself become an object of obligation: the obligation namely to ensure the very premise of all obligation, that is, the *foothold* for a moral universe in the physical world—the existence of mere *candidates* for a moral order. This entails, among other things, the duty to preserve this physical world in such a state that the conditions for that presence remain intact; which in turn means protecting the world's vulnerability from what could imperil those very conditions. The difference this makes for ethics may be illustrated in one example.

V. Old and New Imperatives

1. Kant's categorical imperative said: “Act so that you *can* will that the maxim of your action be made the principle of a universal law.” The “can” here invoked is that of reason and its consistency with itself: *Given* the existence of a community of human agents (acting rational beings), the action must be such that it can without self-contradiction be imagined as a general practice of that community. Mark that the basic reflection of morals here is not itself a moral but a logical one: The “I *can* will” or “I *cannot* will” expresses logical compatibility or incompatibility, not moral approbation or revulsion. But there is no self-contradiction in the thought that humanity would once come to an end, therefore also none in the thought that the happiness of present and proximate generations would be bought with the unhappiness or

even nonexistence of later ones—as little as, after all, in the inverse thought that the existence or happiness of later generations would be bought with the unhappiness or even partial extinction of present ones. The sacrifice of the future for the present is logically no more open to attack than the sacrifice of the present for the future. The difference is only that in the one case the series goes on, and in the other it does not (or: its future ending is contemplated). But that it *ought to go on*, regardless of the distribution of happiness or unhappiness, even with a persistent preponderance of unhappiness over happiness, nay, of immorality over morality⁵—this cannot be derived from the rule of self-consistency *within* the series, long or short as it happens to be: it is a commandment of a very different kind, lying outside and “prior” to the series as a whole, and its ultimate grounding can only be metaphysical.

2. An imperative responding to the new type of human action and addressed to the new type of agency that operates it might run thus: “Act so that the effects of your action are compatible with the permanence of genuine human life”; or expressed negatively: “Act so that the effects of your action are not destructive of the future possibility of such life”; or simply: “Do not compromise the conditions for an indefinite continuation of humanity on earth”; or, again turned positive: “In your present choices, include the future wholeness of Man among the objects of your will.”

3. It is immediately obvious that no rational contradiction is involved in the violation of this kind of imperative. I can will the present good with sacrifice of the future good. Just as I can will my own end, I can will that of humanity. Without falling into contradiction with myself, I can prefer a short fireworks display of the most extreme “self-fulfillment,” for myself or for the world, to the boredom of an endless continuation in mediocrity.

However, the new imperative says precisely that we may risk our own life—but not that of humanity; and that Achilles indeed had the right to choose for himself a short life of glorious deeds over a long life of inglorious security (with the tacit premise that a posterity would be there to know and tell of his deeds), but that we do not have the right to choose, or even risk, nonexistence for future generations on account of a better life for the present one. Why we do not have this right, why on the contrary we have an obligation toward that which does not yet exist and never need exist at all—an obligation not only toward its fortunes in case it happens to exist, but toward its coming to exist in the first place, to which as nonexistent “it” surely has no claim: to underpin this proposition theoretically is by no means easy and without religion perhaps impossible. At present, our imperative simply posits it without proof, as an axiom.

4. It is also evident that the new imperative addresses itself to public policy rather than private conduct, which is not in the causal dimension to which that imperative applies. Kant's categorical imperative was addressed to the individual, and its criterion was instantaneous. It enjoined each of us to consider what would happen *if* the *maxim* of my present action were made, or at this moment already were, the principle of a universal legislation; the self-consistency or inconsistency of such a *hypothetical* universalization is made the test for my *private* choice. But it was no part of the reasoning that there is any probability of my private choice in fact becoming universal law, or that it might contribute to its becoming that. Indeed, *real* consequences are not considered at all, and the principle is one not of objective responsibility but of the subjective quality of my self-determination. The new imperative invokes a different consistency: not that of the act with itself, but that of its eventual *effects* with the continuance of human agency in times to come. And the “universalization” it contemplates is by no means hypothetical—that is, a purely logical transference from the individual “me” to an imaginary, causally unrelated “all” (“if everybody acted like that”); on the contrary, the actions subject to the new imperative—actions of the collective whole—have their universal reference in their actual scope of efficacy: they “totalize” themselves in the progress of their momentum and thus are bound to terminate in shaping the universal dispensation of things. This adds a *time* horizon to the moral calculus which is entirely absent from the instantaneous logical operation of the Kantian imperative: whereas the latter extrapolates into an ever-present order of abstract compatibility, our imperative extrapolates into a predictable real *future* as the open-ended dimension of our responsibility.

VI. Earlier Forms of “Future-oriented Ethics”

Now it may be objected that with Kant we have chosen an extreme example of the ethics of subjective intention (*Gesinnungsethik*), and that our assertion of the present-oriented character of all former ethics, as holding among contemporaries, is contradicted by several ethical forms of the past. The following three examples come to mind: the conduct of earthly life (to the point of sacrificing its entire happiness)

with a view to the eternal salvation of the soul; the long-range concern of the legislator and statesman for the future common weal; and the politics of utopia, with its readiness to use those living now as a mere means to a goal that lies in a future after their time, or to exterminate them as obstacles in its way—of which revolutionary Marxism is the prime example.

1. *The Ethics of Fulfillment in the Life Hereafter*

Of these three cases the first and third share the trait of placing the future above the present as the possible locus of absolute value, thus demoting the present to a mere preparation for the future. An important difference is that in the religious case the acting down here is not credited with bringing on the future bliss by its own causality (as revolutionary action is supposed to do), but is merely supposed to *qualify* the agent for it, namely, in the eyes of God, to whom faith must entrust its realization. That qualification, however, consists in a life pleasing to God, of which in general it may be assumed that it is the best, most worthwhile life in itself anyway, thus worthy to be chosen for its own sake and not merely for that of eventual future bliss. Indeed, when chosen mainly from that reward motive, the life in question would lose in worth and therewith even in its qualifying strength. That is to say, the latter is the greater, the less intended it is. When we then ask what human qualities are held to procure the qualification, that is, to constitute a life pleasing to God, we must look at the life prescriptions of the particular creeds—and these we may often find to be just those prescriptions of justice, charity, purity of heart, etc., which would, or could, be prescribed by an innerworldly ethic of the classical sort as well. Thus in the “moderate” version of the belief in the soul’s salvation (of which, if I am not mistaken, Judaism is an example) we still deal, after all, with an ethics of contemporaneity and immediacy, notwithstanding the transcendent goal; and what ethics it might concretely be in this or that historical case—that is not deducible from the transcendent goal as-such (of whose content no idea can be formed anyway), but is told by the way in which the “life pleasing to God,” said to be the precondition for it, was in each instance given material content.

It may happen, however, that the content is such—and this is the case in the “extreme” forms of the soul salvation doctrine—that its practice, that is, the fulfillment of the “precondition,” can in no way be regarded as of value in itself but is merely the stake in a wager, with whose loss, that is, the failure to attain the eternal reward, all would be lost. For in this case of the dreadful metaphysical bet as elaborated by Pascal, the stake is one’s entire earthly existence with all its possibilities of enjoyment and fulfillment, whose very renunciation is made the price of eternal salvation. In this category belong all those forms of radical mortification of the flesh, of life-denying asceticism, whose practitioners would have cheated themselves out of everything if their expectations were disappointed. This otherworldly wager differs from the calculus of ordinary, this-worldly hedonism, with its considered risks of sometime-renunciations and deferments, merely by the totality of its *quid pro quo* and the surpassing nature of the chance for which the stakes are risked. But just this surpassing expectation moves the whole undertaking out of the realm of ethics. Between the finite and the infinite, the temporal and the eternal, there is no commensurability and thus no meaningful comparison; that is, there is neither a qualitative nor a quantitative sense in which one is *preferable* to the other. Concerning the *value* of the goal, whose informed appraisal ought to form an essential element of *ethical* decision, there is nothing but the empty assertion that it is the ultimate value. Also lacking is the *causal* relation—which at least *ethical* thinking requires—between the action and its (hoped-for) result; that “result,” so we saw, is conceived not as being effected by present renunciation but merely as promised from elsewhere in compensation for it.

If one inquires *why* the this-worldly renunciation is considered so meritorious that it may dare to expect this kind of indemnification or reward, one answer might be that the flesh is sinful, desire is evil, and the world is impure. In this case (as in the somewhat different case where individuation as such is regarded as bad) asceticism does represent, after all, a genuine instrumentality of action and a path to internal goal-achievement through one’s own performance: the path, namely, from impurity to purity, from sinfulness to sanctity, from bondage to freedom, from selfhood to self-transcendence. Insofar as it is such a “path,” asceticism is already in itself the *best* sort of life by the metaphysical criteria assumed. But in this case we are dealing again with an ethic of the here and now: a form—albeit a supremely egotistic and individualistic form—of the ethic of self-perfection, whose inward exertions may indeed attain to those peak moments of spiritual illumination, which are a present foretaste of the future reward: a mystical experience of the Absolute.

In sum, we can say that, insofar as this whole complex of otherworldly striving falls within ethics at all (as do, for instance, the aforementioned “moderate” forms in which a life good in itself forms the condition for eternal reward), it too fits our thesis concerning the orientation of all previous ethics to the

present.

2. *The Statesman's Responsibility for the Future*

What about the examples of *innerworldly* future-oriented ethics, which alone do really belong to rational ethics in that they reckon with a known cause-effect pattern? We mentioned in the second place the long-range care of the legislator and statesman for the future good of the commonwealth. Greek political theory is on the whole silent about the *time* aspect which interests us here; but this silence itself is revealing. Something can be gathered from the praise of great lawgivers like Solon and Lycurgus or from the censure of a statesman like Pericles. The praise of the lawgiver includes, it is true, the durability of his creation, but not his planning ahead of something that is to come about only in aftertimes and not attainable already to his contemporaries. His endeavor is to create a viable political structure, and the test of viability is in the enduring of his creation—a changeless enduring if possible. The best state, so it was thought, is also the best for the future, precisely because the stable equilibrium of its present ensures its future as such; and it will then, of course, be the best state in that future as well, since the criteria of a good order (of which durability is one) do not change. They do not change because human nature does not change, which with its imperfections is included in the conception which the wise lawgiver must have of a viable political order. This conception thus aims not at the ideally perfect state but rather at the realistically best, that is, the best possible state—and this is now just as possible, and just as imperiled, as it will always be. But this very peril, which threatens all order with the disorder of the human passions, makes necessary, in addition to the singular, founding wisdom of the lawgiver, the continuous, governing wisdom of the statesman. The reproach of Socrates against the politics of Pericles, be it noted, is not that, in the end after his death, his grandiose schemes came to nought, but rather that with such grandiose schemes (including their initial successes) he had already in his own time turned the Athenians' heads and corrupted their civic virtues. Athens' current misfortune thus was blamed not on the eventual failure of those policies but on the blemish at their roots, which even "success" in their own terms would not have made better in retrospect. What would have been good at that time would be that still today and would most probably have survived into the present.

The foresight of the statesman thus consists in the wisdom and moderation he devotes to the present. This present is not here for the sake of a future different from (and superior to) it in type, but rather proves itself—luck permitting—in a future still like itself, and so must be as justified already in itself as its succession is hoped to be. Duration, in short, results as a concomitant of what is good now and at all times. Certainly, political action has a wider time span of effect and responsibility than private action, but its ethics, according to the premodern view, is still none other than the present-oriented one, applied to a life form of longer duration.

3. *The Modern Utopia*

a) This changes only with what, in my third example, I called the politics of utopia, which is a thoroughly modern phenomenon and presupposes a previously unknown, dynamic eschatology of history. The religious eschatologies of earlier times do not yet represent this case, although they prepare for it. Messianism, for example, does not ordain a messianic politics, but leaves the coming of the Messiah to divine dispensation. Human behavior is implicated in it only in the sense that it can make itself worthy of the event through fulfilling those very norms to which it is subject even without such a prospect. Here we find to hold on the collective scale what we previously found to hold on the personal scale with regard to otherworldly hopes: the here and now is certainly overarched by them, but is not entrusted with their active realization. It serves them the better, the more faithful it remains to its own God-given law, whose fulfillment lies entirely within itself.

b) Here, too, there did occur the extreme form, where the "urgers of the end" took matters into their own hands and with one last thrust of earthly action tried to bring about the messianic kingdom or millennium, for which they considered the time ripe. In fact, some of the chiliastic movements, especially at the beginning of the modern era, lead into the neighborhood of utopian politics, particularly when they are not content with merely having made a start and clearing the path, but when they make a positive beginning with the Kingdom of God, of whose *contents* they have a definite conception. Insofar as ideas of social equality and justice play a role in this conception, the characteristic motivation of modern utopian ethics is already there: but not yet the yawning gulf, stretching across generations, between now and later, means and end, action and goal, which marks the modern, secularized eschatology, that is, modern political utopianism. It is still an ethic of the self-vindicating present, not of the retroactively vindicating future: the

true man is already there, and even, in the “community of the saints,” the kingdom of God from the moment they realize it in their own midst, as ordained and held to be possible in the dawning fulness of time. The assault, however, against the establishments of the world that still oppose its spreading, is made in the expectation of a Jericho-like miracle, not as a mediated process of historical causation. The last step to the innerworldly utopian ethic of history is yet to be taken.

c) Only with the advent of modern *progress*, both as a fact and as an idea, did the possibility emerge of conceiving everything past as a stepping-stone to the present and of everything present as a stepping-stone to the future. When this notion (which in itself, as unlimited, distinguishes no stage as final and leaves to each the immediacy of its own present) is wed with a secularized eschatology which assigns to the absolute, defined in terms of this world, a finite place in time, and when to this is added a conception of a teleological dynamism which leads to the final state of affairs—then we have the conceptual prerequisites for a utopian politics. “To found the kingdom of heaven already upon earth” (Heinrich Heine) presupposes some idea of what such an earthly kingdom of heaven would look like (or so one would think—but on this point the theory displays a remarkable blank). In any case, even lacking such an idea, the resolute secular eschatology entails a conception of human events that radically demotes to provisional status all that goes before, stripping it of its independent validity and at best making it the vehicle for reaching the promised state of things that is yet to come—a means to the future end which alone is worthy in itself.

Here in fact is a break with the past, and what we have said concerning the present-oriented character of all previous ethics and their common premise of the persistence of human nature is no longer true of the teaching which represents this break most clearly, the Marxist philosophy of history and its corresponding ethic of action. Action takes place for the sake of a future which neither the agent nor the victim nor their contemporaries will live to enjoy. The obligations upon the now issue from that goal, not from the good and ill of the contemporary world; and the norms of action are just as provisional, indeed just as “inauthentic,” as the conditions which it will transmute into the higher state. The ethic of revolutionary eschatology considers itself an ethic of transition, while the consummate, true ethic (essentially still unknown) will only come into its own after the harsh interim morality (which can last a long time) has created the conditions for it and thereby abrogated itself.

Thus there already exists, in Marxism, a future-oriented ethic, with a distance of vision, a time span of affirmed responsibility, a scope of object (= all of future humanity), and a depth of concern (the whole future nature of man)—and, as we might already add, with a sense for the powers of technology—which in all these respects stands comparison with the ethic for which we want to plead here. All the more important it is to determine the relation between these two ethical positions which, as answers to the unprecedented modern situation and especially to its technology, have so much in common over against premodern ethics and yet are so different from one another. This must wait until we have heard more about the problems and tasks which the ethic here envisaged has to deal with, and which are posed by the colossal progress of technology. For technology’s power over human destiny has overtaken even that of communism, which no less than capitalism thought merely to make use of it. We say this much in advance: while both positions concern themselves with the utopian possibilities of this technology, the ethic we are looking for is not eschatological and, in a sense yet to be specified, is anti-utopian.

VII. Man as an Object of Technology

Our comparison dealt with the historical forms of the ethics of contemporaneity and immediacy, for which the Kantian case served only as an example. What stands in question is not their validity within their own frame of reference but their sufficiency for those new dimensions of human action which transcend that frame. Our thesis is that the new kinds and dimensions of action require a commensurate ethic of foresight and responsibility which is as novel as the eventualities which it must meet. We have seen that these are the eventualities that arise out of the works of *homo faber* in the era of technology. But among those novel works we have not mentioned yet the potentially most ominous class. We have considered *techne* only as applied to the nonhuman realm. But man himself has been added to the objects of technology. *Homo faber* is turning upon himself and gets ready to make over the maker of all the rest. This consummation of his power, which may well portend the overpowering of man, this final imposition of art on nature, calls upon the utter resources of ethical thought, which never before has been faced with elective alternatives to what were considered the definite terms of the human condition.

1. *Extension of Life Span*

Take, for instance, the most basic of these “givens,” man’s *mortality*. Who ever before had to make up his mind on its desirable and *eligible* measure? There was nothing to choose about the upper limit, the “threescore years and ten, or by reason of strength fourscore.” Its inexorable rule was the subject of lament, submission, or vain (not to say foolish) wish-dreams about possible exceptions—strangely enough, almost never of affirmation. The intellectual imagination of a George Bernard Shaw and a Jonathan Swift speculated on the privilege of not having to die, or the curse of not being able to die. (Swift with the latter was the more perspicacious of the two.) Myth and legend toyed with such themes against the acknowledged background of the unalterable, which made the earnest man rather pray “teach us to number our days that we may get a heart of wisdom” (Psalm 90). Nothing of this was in the realm of doing and effective decision. The question was only how to relate to the stubborn fact.

But lately the dark cloud of inevitability seems to lift. A practical hope is held out by certain advances in cell biology to prolong, perhaps indefinitely extend, the span of life by counteracting biochemical processes of aging. Death no longer appears as a necessity belonging to the nature of life, but as an avoidable, at least in principle tractable and long-delayable, organic malfunction. A perennial yearning of mortal man seems to come nearer fulfillment. And for the first time we have in earnest to ask the questions “How desirable is this? How desirable for the individual, and how for the species?” These questions involve the very meaning of our finitude, the attitude toward death, and the general biological significance of the balance of death and procreation. Even prior to such ultimate questions are the more pragmatic ones of who should be eligible for the boon: Persons of particular quality and merit? Of social eminence? Those who can pay for it? Everybody? The last would seem the only just course. But it would have to be paid for at the opposite end, at the source. For clearly, on a population-wide scale, the price of extended age must be a proportional slowing of replacement, that is, a diminished access of new life. The result would be a decreasing proportion of youth in an increasingly aged population. How good or bad would that be for the general condition of man? Would the species gain or lose? And how right would it be to preempt the place of youth? Having to die is bound up with having been born: mortality is but the other side of the perennial spring of “natality” (to use Hannah Arendt’s term). This had always been ordained; now its meaning has to be pondered in the sphere of decision.

To take the extreme (not that it will ever be obtained): if we abolish death, we must abolish procreation as well, for the latter is life’s answer to the former, and so we would have a world of old age with no youth, and of known individuals with no surprises of such that had never been before. But this perhaps is precisely the wisdom in the harsh dispensation of our mortality: that it grants us the eternally renewed promise of the freshness, immediacy, and eagerness of youth, together with the supply of otherness as such. There is no substitute for this in the greater accumulation of prolonged experience: it can never recapture the unique privilege of seeing the world for the first time and with new eyes; never relive the wonder which, according to Plato, is the beginning of philosophy; never the curiosity of the child, which rarely enough lives on as thirst for knowledge in the adult, until it wanes there too. This ever renewed beginning, which is only to be had at the price of ever repeated ending, may well be mankind’s hope, its safeguard against lapsing into boredom and routine, its chance of retaining the spontaneity of life. Also, the role of the *memento mori* in the individual’s life must be considered, and what its attenuation to indefiniteness may do to it. Perhaps a nonnegotiable limit to our expected time is necessary for each of us as the incentive to number our days and make them count.

So it could be that what by intent is a philanthropic gift of science to man, the partial granting of his oldest wish—to escape the curse of mortality—turns out to be to the detriment of man. I am not indulging in prediction and, in spite of my noticeable bias, not even in valuation. My point is that already the promised gift raises questions that had never to be asked before in terms of practical choice, and that no principle of former ethics, which took the human constants for granted, is competent to deal with them. And yet they must be dealt with ethically and by principle and not merely by the pressure of interests.

2. *Behavior Control*

It is similar with all the other, quasi-utopian possibilities which progress in the biomedical sciences has partly already placed at our disposal and partly holds in prospect for eventual translation into technological know-how. Of these, *behavior control* is much nearer to practical readiness than the still hypothetical prospect I have just been discussing, and the ethical questions it raises are less profound but

have a more direct bearing on the moral conception of man. Here again, the new kind of intervention exceeds the old ethical categories. They have not equipped us to rule, for example, on mental control by chemical means or by direct electrical action on the brain via implanted electrodes—undertaken, let us assume, for defensible and even laudable ends. The mixture of beneficial and dangerous potentials is obvious, but the lines are not easy to draw. Relief of mental patients from distressing and disabling symptoms seems unequivocally beneficial. But from the relief of the *patient*, a goal entirely in the tradition of the medical art, there is an easy passage to the relief of *society* from the inconvenience of difficult individual behavior among its members: that is, the passage from medical to social application: and this opens up an indefinite field with grave potentials. The troublesome problems of rule and unruliness in modern mass society make the extension of such control methods to nonmedical categories extremely tempting for social management. Numerous questions of human rights and dignity arise. The difficult question of preempting versus enabling care insists on concrete answers. Shall we induce learning attitudes in schoolchildren by the mass administration of drugs, circumventing the appeal to autonomous motivation? Shall we overcome aggression by electronic pacification of brain areas? Shall we generate sensations of happiness or pleasure or at least contentment through independent stimulation (or tranquilizing) of the appropriate centers—*independent*, that is, of the objects of happiness, pleasure, or content and their attainment in personal living and achieving? Candidacies could be multiplied. Business firms might become interested in some of these techniques for performance increase among their employees.

Regardless of the question of compulsion or consent, and regardless also of the question of undesirable side-effects, each time we thus bypass the human way of dealing with human problems, short-circuiting it by an impersonal mechanism, we have taken away something from the dignity of personal selfhood and advanced a further step on the road from responsible subjects to programmed behavior systems. Social functionalism, important as it is, is only one side of the question. Decisive is the question of what kind of individuals the society is composed of—to make its existence valuable as a whole. Somewhere along the line of increasing social manageability at the price of individual autonomy, the question of the worthwhileness of the whole human enterprise must pose itself. Answering it involves the image of man we entertain. We must think it anew in light of the things we can do with it or to it now and could never do before.

3. Genetic Manipulation

This holds even more with respect to the last object of a technology applied on man himself—the *genetic* control of future men. This is too wide a subject for the cursory treatment of these prefatory remarks, and it will have its own chapter in a later “applied part” to succeed this volume. Here I merely point to this most ambitious dream of *homo faber*, summed up in the phrase that man will take his own evolution in hand, with the aim of not just preserving the integrity of the species but of modifying it by improvements of his own design. Whether we have the right to do it, whether we are qualified for that creative role, is the most serious question that can be posed to man finding himself suddenly in possession of such fateful powers. Who will be the image-makers, by what standards, and on the basis of what knowledge? Also, the question of the moral right to experiment on future human beings must be asked. These and similar questions, which demand an answer before we embark on a journey into the unknown, show most vividly how far our powers to act are pushing us beyond the terms of all former ethics.

VIII. The “Utopian” Dynamics of Technical Progress and the Excessive Magnitude of Responsibility

The ethically relevant common feature in all the examples adduced is what I like to call the inherently “utopian” drift of our actions under the conditions of modern technology, whether it works on nonhuman or on human nature, and whether the “utopia” at the end of the road be planned or unplanned. By the kind and size of its snowballing effects, technological power propels us into goals of a type that was formerly the preserve of Utopias. To put it differently, technological power has turned what used and ought to be tentative, perhaps enlightening plays of speculative reason into competing blueprints for projects, and in choosing between them we have to choose between extremes of remote effects. The one thing we can really know of them is their extremism as such—that they concern the total condition of nature on our globe and the very kind of creatures that shall, or shall not, populate it. In consequence of the inevitably

“utopian” scale of modern technology, the salutary gap between everyday and ultimate issues, between occasions for common prudence and occasions for illuminated wisdom, is steadily closing. Living now constantly in the shadow of unwanted, built-in, automatic utopianism, we are constantly confronted with issues whose positive choice requires supreme wisdom—an impossible situation for man in general, because he does not possess that wisdom, and in particular for contemporary man, because he denies the very existence of its object, namely, objective value and truth. We need wisdom most when we believe in it least.

If the new nature of our acting then calls for a new ethics of long-range responsibility, coextensive with the range of our power, it calls in the name of that very responsibility also for a new kind of humility—a humility owed, not like former humility to the smallness of our power, but to the excessive magnitude of it, which is the excess of our power to act over our power to foresee and our power to evaluate and to judge. In the face of the quasi-eschatological potentials of our technological processes, ignorance of the ultimate implications becomes itself a reason for responsible restraint—as the second best to the possession of wisdom itself.

One other aspect of the required new ethics of responsibility for and to a distant future is worth mentioning: the doubt it casts on the capacity of representative government, operating by its normal principles and procedures, to meet the new demands. For according to those principles and procedures, only *present* interests make themselves heard and felt and enforce their consideration. It is to them that public agencies are accountable, and this is the way in which concretely the respecting of rights comes about (as distinct from their abstract acknowledgment). But the *future* is not represented, it is not a force that can throw its weight into the scales. The nonexistent has no lobby, and the unborn are powerless. Thus accountability to them has no political reality behind it in present decision-making, and when they can make their complaint, then we, the culprits, will no longer be there.

This raises to an ultimate pitch the old question of the power of the wise, or the force of ideas not allied to self-interest, in the body politic. What force shall represent the future in the present? That is a question for political philosophy, and one on which I dare not voice my woefully uncertain ideas. They would be premature here anyway. For before that question of enforcement can become practical, the new ethics must find its theory, on which do's and don'ts can be based. That is: before the question of what *force*, comes the question of what *insight* or value-knowledge will represent the future in the present.

IX. The Ethical Vacuum

And here is where I come to a standstill, where we all come to a standstill. For the very same movement which put us in possession of the powers that have now to be regulated by norms—the movement of modern knowledge called science—has by a necessary complementarity eroded the foundations from which norms could be derived; it has destroyed the very idea of norm as such. Not, fortunately, the feeling for norm and even for particular norms. But this feeling becomes uncertain of itself when contradicted by alleged knowledge or at least denied all support by it. It always has a difficult time against the loud clamors of greed and fear. Now it must in addition blush before the frown or smirk of superior knowledge which has certified it as unfounded and incapable of foundation. First it was nature that was “neutralized” with respect to value, then man himself. Now we shiver in the nakedness of a nihilism in which near-omnipotence is paired with near-emptiness, greatest capacity with knowing least for what ends to use it.

It is moot whether, without restoring the category of the sacred, the category most thoroughly destroyed by the scientific enlightenment, we can have an ethics able to cope with the extreme powers which we possess today and constantly increase and are almost compelled to wield. Regarding those consequences that are imminent enough still to hit ourselves, fear can do the job—fear which is so often the best substitute for genuine virtue or wisdom. But this means fails us toward the more distant prospects, which here matter the most, especially as the beginnings seem mostly innocent in their smallness. Only awe of the sacred with its unqualified veto is independent of the computations of mundane fear and the solace of uncertainty about distant consequences. However, religion in eclipse cannot relieve ethics of its task; and while of faith it can be said that as a moving force it either is there or is not, of ethics it is true to say that it must be there.

It must be there because men act, and ethics is for the ordering of actions and for regulating the power to act. It must be there all the more, then, the greater the powers of acting that are to be regulated;

and as it must fit their size, the ordering principle must also fit their kind. Thus, novel powers to act require novel ethical rules and perhaps even a new ethics.

“Thou shalt not kill” was enunciated because man has the power to kill and often the occasion and even the inclination for it—in short, because killing is actually done. It is only under the *pressure* of real habits of action, and generally of the fact that always action already takes place, without *this* having to be commanded first, that ethics as the ruling of such acting under the standard of the good or the permitted enters the stage. Such a *pressure* emanates from the novel technological powers of man, whose exercise is given with their existence. *If* they really are as novel in kind as here contended, and if by the kind of their potential consequences they really have abolished the moral neutrality which the technical commerce with matter hitherto enjoyed—then their pressure bids us to seek for new prescriptions in ethics which are competent to assume their guidance, but which first of all can hold their own theoretically against that very pressure.

In this chapter we have developed our *premises*, namely, first, that our collective technological practice constitutes a new kind of human action, and this not just because of the novelty of its methods but more so because of the unprecedented nature of some of its objects, because of the sheer magnitude of most of its enterprises, and because of the indefinitely cumulative propagation of its effects. From all three of these traits, our second premise follows: that what we are doing in this manner is, regardless of the particulars of any of its immediate purposes, no longer ethically neutral as a whole. With this exposition of the ethical question, the task of seeking an answer, and first of all a rational principle for it, only begins.

1. Immanuel Kant, *Groundwork of the Metaphysics of Morals*, preface.

2. *Ibid.*, chap. 1.

3. *Ibid.* (I have followed H. J. Paton’s translation with some changes.)

4. Except in self-cultivation and in education. E.g., the practice of virtue is also a “learning” of its discipline and as such progressive; it strengthens the moral powers and makes their exercise habitual (as the converse is true of bad habits). But naked primal nature can always break through again. The most virtuous can be caught in the destructive tempest of passion, and the most wicked may experience conversion. Is the same still possible with the cumulative changes in the conditions of existence which technology deposits on its path?

5. On this last point, the biblical God changed his mind to an all-encompassing “yes” after the Flood.

Chapter 7 from *A Social History of American Technology*, by Ruth Schwartz Cowan. Oxford University Press, 1997.

Industrial Society and Technological Systems

by

Ruth Schwartz Cowan

BETWEEN 1870 and 1920, the United States changed in ways that its founders could never have dreamed possible. Although American industrialization began in the 1780s, the nation did not become an industrialized society until after the Civil War had ended. The armistice agreed to at Appomattox signaled, although the participants probably did not realize it, the beginning of the take-off phase of American industrialization. Having begun as a nation of farmers, the United States became a nation of industrial workers. Having begun as a financial weakling among the nations, by 1920 the United States had become the world's largest industrial economy.

What did this transformation mean to the people who lived through it? When a society passes from preindustrial to industrial conditions, which is what happened in the United States in the years between 1870 and 1920, people become less dependent on nature and more dependent on each other. This is one of history's little ironies. In a preindustrial society, when life is unstable, the whims of the weather and the perils of natural cycles are most often to blame. In an industrial society, when life is unstable, the individuals become more dependent on one another because they are linked together in large, complex networks that are, at one and the same time, both physical and social: technological systems.

Industrialization, Dependency, and Technological Systems

Many Americans learned what it means to become embedded in a set of technological systems in the years between 1870 and 1920. Today we have become so accustomed to these systems that we hardly ever stop to think about them; although they sustain our lives, they nonetheless remain mysterious. In the late twentieth century, people have tended to think that, if anything, industrialization has liberated them from dependency, not encased them in it, but that is not the case. We can see this clearly by imagining how a woman might provide food for a two-year-old child in a non-industrialized society.

In a hunter-gatherer economy, she might simply go into the woods and collect nuts or walk to the waterside and dig for shellfish. In a premodern agricultural community (such as the one that some of the native peoples of the eastern seaboard had created), she might work with a small group of other people to plant corn, tend it, harvest it, and shuck it. Then she herself might dry it, grind it into meal, mix it with water, and bake it into a bread for the child to eat. In such a community, a woman would be dependent on the cooperation of several other people in order to provide enough food for her child, but all of those people would be known to her and none of them would be involved in an activity in which she could not have participated if necessity had demanded.

In an industrialized economy (our own, for example), an average woman's situation is wholly different. In order to get bread for a child, an average American woman is dependent on thousands of other people, virtually all of them totally unknown to her, many of them living and working at a considerable distance, employing equipment that she could not begin to operate, even if her life (quite literally) depended on it and even if she had the money (which isn't likely) to purchase it. A farmer grew the wheat using internal combustion engines and petroleum-derivative fertilizers. Then the wheat was harvested and transported to an organization that stored it under stable conditions, perhaps for several years. Then a milling company may have purchased it and transported it (over thousands of miles of roads or even ocean) to a mill, where it was ground by huge rollers powered by electricity (which itself may have been generated thousands of miles away). Then more transportation (all of this transportation required petroleum, which itself had to be processed and transported) was required: to a baking factory, where dozens of people (and millions of dollars of machinery) were used to turn the flour into bread. Then transportation again: to a market, where the woman could purchase it (having gotten herself there in an automobile, which itself had to be manufactured somewhere else, purchased at considerable expense, and supplied with fuel)—all of this before a slice of it could be spread with peanut butter to the delight of a two year old.

The point should, by now, be clear. People who live in agricultural societies are dependent on natural processes: they worry, with good reason, about whether and when there will be a drought or a flood, a plague of insects or of fungi, good weather or bad. People who live in industrial societies are not completely independent of such natural processes, but are more so than their predecessors (many floodplains have been controlled; some droughts can be offset by irrigation). At the same time, they are much more dependent on other people and on the technological systems that other people have designed and constructed. The physical parts of these systems are networks of connected objects: tractors, freight cars, pipelines, automobiles, display cases. The social parts are networks of people and organizations that make the connections between objects possible: farmers, bakers, and truck drivers; grain elevators, refineries, and supermarkets.

Preindustrialized societies had such networks of course (some of them are described in Chapter 2), but in industrialized societies, the networks are more complex and much denser—all of which makes it much harder for individuals to extricate themselves. A small change very far away can have enormous effects very quickly. Daily life can be easily disrupted for reasons that ordinary people can find hard to understand, and even experts can have difficulty comprehending.

People live longer and at a higher standard of living in industrial societies than in preindustrial ones, but they are not thereby rendered more independent (although advertising writers and politicians would like them to think they are) because, in the process of industrialization, one kind of dependency is traded for another: nature for technology. Americans learned what it meant to make that trade in the years between 1870 and 1920. We can begin understanding what they experienced if we look at some of the technological systems that were created or enlarged during those years.

The Telegraph System

The very first network that Americans experienced really looked like a network: the elongated spider's web of electric wires that carried telegraph signals. The fact that electricity could be transmitted long distances through wires had been discovered in the middle of the eighteenth century. Once a simple way to generate electric currents had been developed (a battery, or voltaic pile, named after the man who invented it, Alessandro Volta) many people began experimenting with various ways to send messages along the wires. An American portrait painter, Samuel F. B. Morse, came up with a practicable solution (see Chapter 6). Morse developed a transmitter that emitted a burst of electric current of either short or long duration (dots and dashes). His receiver, at the other end of the wire, was an electromagnet, which, when it moved, pushed a pencil against a moving paper tape (thus recording the pattern of dots and dashes). The most creative aspect of Morse's invention was his code, which enabled trained operators to make sense out of the patterns of dots and dashes.

In 1843, after Morse had obtained a government subvention, he and his partners built the nation's first telegraph line between Baltimore and Washington. By 1845, Morse had organized his own company to build additional lines and to licence other telegraph companies so that they could build even more lines, using the instruments he had patented. In a very short time, dozens of competing companies had entered the telegraph business, and Morse had all he could do to try to collect the licensing fees to which he was entitled. By 1849, almost every state east of the Mississippi had telegraph service, much of it provided by companies that were exploiting Morse's patents without compensating him.

Beginning around 1850, one of these companies, the New York and Mississippi Valley Printing Telegraph Company, began buying up or merging with all the others; in 1866, it changed its name to the Western Union Telegraph Company. In the decades after the Civil War, Western Union had an almost complete monopoly on telegraph service in the United States; a message brought to one of its offices could be transmitted to any of its other offices in almost all fairly large communities in the United States. Once the message was delivered, recipients could pick it up at a Western Union office. During these decades, only one company of any note succeeded in challenging Western Union's almost complete monopoly on telegraph service. The Postal Telegraph Company specialized in providing pick-up and delivery services for telegrams; yet even at the height of its success, it never managed to corner more than 25 percent of the country's telegraph business.

In 1866, when Western Union was incorporated, it already controlled almost 22,000 telegraph offices around the country. These were connected by 827,000 miles of wire (all of it strung from a virtual forest of telegraph poles, many of them running along railroad rights of way), and its operators were

handling something on the order of 58 million messages annually. By 1920, the two companies (Western and Postal) between them were managing more than a million miles of wire and 155 million messages. Yet other companies (many of the railroads, for example, several investment banking houses, several wire news services) were using Western Union and Postal Telegraph lines on a contractual basis to provide in-house communication services (the famous Wall Street stock ticker was one of them).

As a result, as early as 1860, and certainly by 1880, the telegraph had become crucial to the political and economic life of the nation. Newspapers had become dependent on the telegraph for quick transmission of important information. The 1847 war with Mexico was the first war to have rapid news coverage, and the Civil War was the first in which military strategy depended on the quick flow of battle information over telegraph lines. During the Gilded Age (1880-1900), the nation's burgeoning financial markets were dependent on the telegraph for quick transmission of prices and orders. Railroad companies used the telegraph for scheduling and signaling purposes since information about deviations in train times could be quickly transmitted along the lines. The central offices of the railroads utilized telegraph communication to control the financial affairs of their widely dispersed branches. When the Atlantic cable was completed in 1866, the speed and frequency of communication between nations increased, thereby permanently changing the character of diplomatic negotiations. The cable also laid the groundwork for the growth of international trade (particularly the growth of multinational corporations) in the later decades of the century.

In short, by 1880, if by some weird accident all the batteries that generated electricity for telegraph lines had suddenly run out, the economic and social life of the nation would have faltered. Trains would have stopped running; businesses with branch offices would have stopped functioning; newspapers could not have covered distant events; the president could not have communicated with his European ambassadors; the stock market would have had to close; family members separated by long distances could not have relayed important news—births, deaths, illnesses—to each other. By the turn of the century, the telegraph system was both literally and figuratively a network, linking together various aspects of national life—making people increasingly dependent on it and on one another.

The Railroad System

Another system that linked geographic regions, diverse businesses, and millions of individuals was the railroad. We have already learned (in Chapter 5) about the technical developments (the high-pressure steam engine, the swivel truck, the T-rail) that were crucial to the development of the first operating rail lines in the United States in the 1830s. Once the technical feasibility of the railroad became obvious, its commercial potential also became clear. The railroad, unlike canals and steamboats, was not dependent on proximity to waterways and was not (as boats were) disabled when rivers flooded or canals froze.

During the 1840s, American entrepreneurs had begun to realize the financial benefits that railroading might produce and railroad-building schemes were being concocted in parlors and banks, state houses, and farm houses all across the country. By the 1850s, a good many of those schemes had come to fruition. With 9,000 miles of railroad track in operation, the United States had more railroad mileage than all other western nations combined; by 1860, mileage had more than trebled, to 30,000 miles.

The pre-Civil War railroad system was not yet quite a technological system because, large as it was, it still was not integrated as a network. Most of the existing roads were short-haul lines, connecting such major cities as New York, Chicago, and Baltimore with their immediate hinterlands. Each road was owned by a different company, each company owned its own cars, and each built its tracks at the gauge (width) that seemed best for the cars it was going to attempt to run and the terrain over which the running had to be done. This lack of integration created numerous delays and additional expenses. In 1849, it took nine transshipments between nine unconnected railroads (and nine weeks of travel) to get freight from Philadelphia to Chicago. In 1861, the trip between Charleston and Philadelphia required eight car changes because of different gauges. During and immediately after the Civil War, not a single rail line entering either Philadelphia or Richmond made a direct connection with any other, much to the delight of the local teamsters, porters, and tavern keepers.

The multifaceted processes summed up under the word “integration” began in the years just after the Civil War and accelerated in the decades that followed. The rail system grew ever larger, stretching from coast to coast (with the completion of the Union Pacific Railroad in 1869), penetrating into parts of the country where settlement did not yet even exist. There were roughly 53,000 miles of track in 1870, but

there were 93,000 miles by the time the next decade turned, and 254,000—the all-time high—by 1920. In that half century, the nation's population tripled, but its rail system grew sevenfold; the forty-eight states of the mainland United States became physically integrated, one with the other.

The form of the rail system was just as significant as its size. By 1920, what had once been a disjointed collection of short (usually north-south) lines had been transformed into a network of much longer trunk lines (running from coast to coast, east-west), each served by a network of shorter roads that connected localities (the limbs) with the trunks. Passengers could now travel from New York to San Francisco with only an occasional change of train and freight traveled without the necessity of transshipments. What had remade this kind of integration possible was not a technological change, but a change in the pattern of railroad ownership and management.

From the very beginning of railroading, railroad companies had been joint-stock ventures (see Chapter 5). Huge amounts of capital had been required to build a railroad: rights of way had to be purchased, land cleared, bridges built, locomotives ordered, passenger cars constructed, freight cars bought. Once built, railroads were very expensive to run and to maintain: engines had to be repaired, passengers serviced, freight loaded, tickets sold, stations cleaned. Such a venture could not be financed by individuals, or even by partnerships. Money had to be raised both by selling shares of ownership in the company to large numbers of people and by borrowing large sums of money by issuing bonds.

As a result, both American stockbroking and American investment banking were twin products of the railroad age. Some of America's largest nineteenth-century fortunes were made by people who knew not how to build railroads, but how to finance them: J. P. Morgan, Leland Stanford, Jay Gould, Cornelius Vanderbilt, and George Crocker. These businessmen consolidated the railroads. They bought up competing feeder lines; they sought control of the boards of directors of trunk lines; they invested heavily in the stock of feeder roads until the feeders were forced to merge with the trunks. When they were finished, the railroads had become an integrated network, a technological system. In 1870, there had been several hundred railroads, many of which were in direct competition with each other. By 1900, virtually all the railroad mileage in the United States was either owned or controlled by just seven (often mutually cooperative) railroad combinations, all of which owed their existence to the machinations of a few very wealthy investment bankers.

As railroad ownership became consolidated, the railroad system became physically integrated. The most obvious indicator of this integration was the adoption of a standard gauge, which made it unnecessary to run different cars on different sets of tracks. By the end of the 1880s, virtually every railroad in the country had voluntarily converted to a gauge of 4 feet, 8 ½ inches in order to minimize both the expense and the delays of long distance travel. On this new integrated system, the need for freight and passengers to make repeated transfers was eliminated; as a result, costs fell while transportation speed increased.

The railroad system had a profound impact on the way in which Americans lived. By 1900, the sound of the train whistle could be heard in almost every corner of the land. Virtually everything Americans needed to maintain and sustain their lives was being transported by train. As much as they may have grumbled about freight rates on the railroads (and there was much injustice, particularly to farmers, to grumble about) and as much as they may have abhorred the techniques that the railroad barons had used to achieve integration, most Americans benefited from the increased operational efficiency that resulted.

In the years in which population tripled and rail mileage increased seven times, freight tonnage on the railroads went up elevenfold. Cattle were going by train from the ranches of Texas to the slaughterhouses of Chicago; butchered beef was leaving Chicago in refrigerated railroad cars destined for urban and suburban kitchens. Lumber traveled from forests to sawmills by train; two-by-four beams to build houses on the treeless plains left the sawmills of the Pacific Northwest on flatcars. Some petroleum went from the well to the refinery by train; most kerosene and gasoline went from the refinery to the retailer by train. Virtually all the country's mail traveled by train, including cotton cloth and saddles, frying pans and furniture ordered from the mail-order companies that had begun to flourish in the 1880s.

Even as fundamental and apparently untransportable a commodity as time was affected by the integration of the rail system, for scheduling was an important facet of integration. People who were going to travel by train had to know what time their trains would leave, and if connection had to be made, trains had to be scheduled so as to make the connections possible. Schedules also had to be constructed, especially on heavily trafficked lines, to ensure that trains did not collide. But scheduling was exceedingly difficult across the long distances of the United States because communities each established their own time on the basis of the position of the sun. When it was noon in Chicago, it was 12:30 in Pittsburgh (which is to the east of Chicago) and 11:30 in Omaha (to the west). The train schedules printed in Pittsburgh in the

early 1880s listed six different times for the arrival and departure of each train. The station in Buffalo had three different clocks.

Sometime in the early 1880s, some professional railroad managers and the editors of several railroad publications agreed to the idea, first proposed by some astronomers, that the nation should be divided into four uniform time zones. By common agreement among the managers of the country's railroads, at noon (in New York) on Sunday, November 18, 1883, railroad signalmen across the country reset their watches. The zones were demarcated by the 75th, 90th, 105th, and 120th meridians. People living in the eastern sections of each zone experienced, on that otherwise uneventful Sunday, two noons, and people living in the western sections, skipped time. Virtually everyone in the country accepted the new time that had been established by the railroads, although Congress did not actually confirm the arrangement by legislation for another thirty-five years. Such was the pervasive impact of the integrated rail network.

The Petroleum System

In 1859, a group of prospectors dug a well in a farmyard in Titusville, Pennsylvania. Although they appeared to be looking for water, the prospectors were in fact searching for an underground reservoir of a peculiar oily substance that had been bubbling to the surface of nearby land and streams. Native Americans had used this combustible substance as a lubricant for centuries. The prospectors were hoping that if they could find a way to tap into an underground reservoir of this material, they could go into the business of selling it to machine shops and factories (as a machine lubricant, an alternative to animal fat) and to households and businesses (as an illuminant, an alternative to whale oil and candles).

The prospectors struck oil—and the American petroleum industry was born. Within weeks the news had spread, and hundreds of eager profiteers rushed into western Pennsylvania, hoping to purchase land, drill for oil, or find work around the wells. The Pennsylvania oil rush was as massive a phenomenon as the California gold rush a decade earlier.

The drillers soon discovered that crude petroleum is a mixture of oils of varying weights and characteristics. These oils, they learned, could be easily separated from one another by distillation, an ancient and fairly well-known craft. All that was need was a fairly large closed vat with it long outlet tube (called a still) and a fire. The oil was heated in the still and the volatile gases produced would condense in the outlet tube. A clever distiller (later called a refiner) could distinguish different portions (fractions) of the distillate from each other, and then only the economically useful ones needed to be bottled and sent to market.

The market for petroleum products boomed during the Civil War: northern factories were expanding to meet government contracts; the whaling industry was seriously hampered by naval operations; railroads were working overtime to transport men and materiel to battlefronts. By 1862, some 3 million barrels of crude oil were being processed every year. Under peacetime conditions the industry continued to expand; by 1872, the number of processed barrels had trebled.

Transportation of petroleum remained a problem, however. The wells were located in the rural, underpopulated Appalachian highlands of Pennsylvania, not only many miles away from the cities in which the ultimate consumers lived, but also many miles away from railroad lines that served those cities. Initially crude oil had been collected in barrels and had been moved (by horse and cart or by river barges) to railroad-loading points. There the barrels were loaded into freight cars for the trip to the cities (such as Cleveland and Pittsburgh) in which the crude was being refined and sold. The transportation process was cumbersome, time-consuming, and wasteful; the barrels leaked, the barges sometimes capsized, the wagons—operating on dirt roads—sometimes sank to their axles in mud.

Pipelines were an obvious solution, but a difficult one to put into practice given that no one had ever before contemplated building and then maintaining a continuous pipeline over the mountainous terrain and the long distances that had to be traversed. The first pipeline to operate successfully was built in 1865. Made of lap-welded cast-iron pipes, two inches in diameter, it ran for six miles from an oil field to a railroad loading point and had three pumping stations along the way. This first pipeline carried eighty barrels of oil an hour and had demonstrated its economic benefits within a year. Pipeline mileage continued to increase during the 1870s and 1880s (putting thousands of teamsters out of business), but virtually all of the lines were relatively short hauls, taking oil from the fields to the railroads. Throughout the nineteenth century and well into the twentieth, the railroads were still the principal long-distance transporters of both crude and refined oil. After the 1870s, the drillers, refiners, and railroads gradually dispensed with barrels

(thus putting thousands of coopers out of business) and replaced them with specially built tank cars, which could be emptied into and loaded from specially built holding tanks. As it was being constructed, the network of petroleum pipelines was thus integrated into the network of railroad lines. It was also integrated into the telegraph network. Oil refineries used the telegraph system partly to keep tabs on prices for oil in various localities and partly to report on the flow of oil through the lines.

The most successful petroleum entrepreneurs were the ones who realized that control of petroleum transportation was the key ingredient in control of the entire industry. The major actor in this particular economic drama was John D. Rockefeller. Rockefeller had been born in upstate New York, the son of a Talented patent medicine salesman, but he had grown up in Cleveland, Ohio, a growing commercial center (it was a Great Lake port and both a canal and railroad terminus), and had learned accountancy in a local commercial college. His first job was as a bookkeeper for what was then called a commission agent, a business that collected commissions for arranging the shipment of bulk orders of farm products. A commission agent's success depended on getting preferential treatment from railroads and shipping companies. Rockefeller carried this insight with him, first when he went into a partnership as his own commission agent and then, in 1865, when he became the co-owner of an oil refinery in Cleveland.

Rockefeller and his associates were determined to control the then chaotic business of oil refining. They began by arranging for a secret rebate on oil shipments from one of the two railroads then serving Cleveland. Then in the space of less than a month, using the rebate as an incentive, they managed to coerce other Cleveland refiners into selling out and obtained control of the city's refining. Within a year or two, Rockefeller was buying up refineries in other cities as well. He had also convinced the railroads that he was using that they should stop carrying oil to refineries owned by others, so that he was in almost complete control of the price offered to drillers. In the early 1870s, a group of drillers banded together to build pipelines that would take their oil to railroads with which Rockefeller wasn't allied. Rockefeller responded to this challenge by assembling a monopoly on the ownership of tank cars (since the pipelines did not go all the way to the refineries and railroad tank cars were still necessary), and by 1879, he had been so successful in squeezing the finances of the pipeline in companies that their stockholders were forced to sell out to him. In that year, as a result of their control both of refineries and pipelines, Rockefeller and his associates controlled 90 percent of the refined oil in the United States.

Having bought up the competing pipelines (having let other people take the risks involved in developing new technologies for building and maintaining those lines), Rockefeller was quick to see their economic value. In 1881, one of his companies completed a six-inch line from the Pennsylvania oil fields to his refinery in Bayonne, New Jersey—the first pipeline that functioned independently of the railroads. By 1900, Rockefeller had built pipelines to Cleveland, Philadelphia, and Baltimore, and Standard Oil (Rockefeller's firm) was moving 24,000 barrels of crude a day (he still used the railroads to move the oil after it had been refined).

By that point, hundreds of civil and mechanical engineers were working for Rockefeller's pipeline companies (which held several patents on pipeline improvements), and several dozen chemists and chemical engineers were working in his refineries (and developing new techniques, such as the Frasch process for taking excess sulfur out of petroleum). In addition, Standard Oil was pioneering financial, management, and legal techniques for operating a business that had to control a huge physical network, spread out over several states. Since the laws dealing with corporations differed in each state and since some of them prevented a corporation in one state from owning property in another, one of Rockefeller's attorneys worked out a corporate arrangement so that Standard Oil had a different corporation in each state in which it operated (Standard Oil of New Jersey, Standard Oil of Ohio, and so forth). The stock holders in each corporation turned their stock over to a group of trustees, who managed the whole enterprise from New York—the famous Standard Oil Trust, of which Rockefeller himself was the single largest stockholder and therefore the major trustee. (The trust, as a way to organize a complex business, was soon picked up in tobacco and sugar refining and other industries involved in large-scale chemical processing, leading Congress, worried about the monopolistic possibilities, to pass the Sherman Anti-Trust Act in 1890.)

By 1900, the Standard Oil Trust (which had successfully battled antitrust proceedings in court) controlled most of the oil produced in Pennsylvania, and it owned most of the new oil fields that had been discovered in Ohio and Indiana. Rockefeller's almost complete stranglehold on the industry wasn't broken until oil was discovered early in the twentieth century in Texas, Oklahoma, Louisiana, and California, outside the reach of the pipelines he controlled and the railroads with which he was associated. Increased competition was accompanied by the continued growth not only of the pipeline network, but also of the

industry as a whole: 26 million barrels of petroleum were processed in 1880, 45 million in 1890, 63 million in 1900, 209 million in 1910 (as gasoline was just beginning to edge out kerosene as the most important petroleum product), and 442 million in 1920 (when the Model T had been in production for almost eight years).

Like the telegraph and the railroad (and in combination with the telegraph and the railroad), the oil pipeline network had become a pervasive influence on the American economy and on the daily life of Americans. In the last decades of the nineteenth century, a very large number of Americans, especially those living outside of the major cities, used one of its products, kerosene, for heating and lighting their homes and for cooking. During the same decades, American industry became dependent on other fractions of petroleum to lubricate the machinery with which it was producing everything from luxurious cloth to common nails. Finally, in the early decades of the twentieth century, with the advent first of the internal combustion engine fueled by gasoline and then of automobiles and trucks powered by that engine, Americans discovered that access to petroleum was becoming a necessary condition not only of their working lives but also of their leisure time.

The Telephone System

Technologically the telephone was similar to the telegraph, but socially it was very different. The device patented by Alexander Graham Bell in 1876 was rather like a telegraph line: voices rather than signals could be transmitted by electric current because the transmitter lever and the receiving pencil had been replaced by very sensitive diaphragms. Aware of the difficulties that Morse had encountered in reaping profits from his patents—and aware that he had no head for business—Bell decided to turn over the financial and administrative details of creating a telephone network to someone else.

The businessmen and the attorneys who managed the Bell Telephone Company did their work well. While the railroad, telegraph, and petroleum networks had been integrated by corporate takeovers, the telephone system was integrated, from the very beginning, by corporate design. A crucial decision had been made early on: Bell Telephone would manufacture all the telephone instruments, then lease the instruments to local companies, which would operate telephone exchanges under license to Bell. This meant that for the first sixteen years of telephone network development (sixteen years was then the length of monopoly rights under a patent), the Bell Telephone Company could dictate, under the licensing agreements, common technologies for all the local telephone systems. Bell could also control the costs of telephone services to local consumers.

Because of this close supervision by one company, the telephone system was integrated from the very beginning. Between 1877 and 1893, the Bell Telephone Company, through its affiliated local operating companies, controlled and standardized virtually every telephone, every telephone line, and every telephone exchange in the nation. Indeed in the 1880s, the officers of Bell were confident that they could profitably begin long-distance service (that is, service that would connect one local operating company with another) precisely because all of the operating companies were using its standardized technology. Bell needed to hire physicists and electrical engineers to solve the technical problems involved in maintaining voice clarity over very long wires, but the organizational problems involved in connecting New York with Chicago and Chicago with Cleveland turned out to be minimal.

On the assumption that the telephone system would end up being used very similarly to the telegraph network, the officers of Bell had decided that their most important customers would be other businesses, particularly those in urban areas. They decided, as a marketing strategy, to keep rates fairly high, in return for which they would work to provide the clearest and most reliable service possible. By the end of the company's first year of operation, 3,000 telephones had been leased, 1 for every 10,000 people. By 1880, there were 60,000 (1 per 1,000), and when the Bell patents expired in 1893, there were 260,000 (1 per 250). About two thirds of these phones were located in businesses. Most of the country's business information was still traveling by mail and by telegraph (because businessmen wanted a written record of their transactions), but certain kinds of businesses were starting to find the telephone very handy: in 1891, the New York and New Jersey Telephone Company served 937 physicians and hospitals, 401 pharmacies, 363 liquor stores, 315 stables, 162 metalworking plants, 146 lawyers, 126 contractors, and 100 printing shops. After the Bell patents expired, independent telephone companies entered the business despite Bell's concerted effort to keep them out. By 1902, there were almost 9,000 such independent companies, companies not part of the Bell system. When the organizers of the Bell system had analogized the

telephone to the telegraph, they had made a crucial sociological mistake. They understood that in technological terms the telephone was similar to the telegraph, but they failed to understand that in social terms it was quite different. The telephone provided user-to-user communication (with the telegraph there were always intermediaries). In addition, the telephone was a form of voice communication; it facilitated emotional communication, something that was impossible with a telegraph. In short, what the organizers of the Bell system had failed to understand was that people would use the telephone to socialize with each other. The independent companies took advantage of Bell's mistake. Some of them offered services that Bell hadn't thought to provide. Dial telephones were one such service, allowing customers to contact each other without having to rely on an operator (who sat at a switchboard, manually connecting telephone lines, one to another, with plugs). Operators were notorious for relieving the boredom of their jobs by listening in on conversations, something many customers wanted to avoid. Party lines were another such service. Anywhere from two to ten residences could share the same telephone line and telephone number, which drastically lowered the costs of residential services. Many lower-income people turned out to be willing to put up with the inconvenience of having to endure the ringing of telephones on calls meant for other parties in exchange for having telephone service at affordable rates.

Yet other independent companies served geographic locales that the Bell companies had ignored. This was particularly the case in rural areas where there were farm households. Bell managers apparently hadn't thought that farmers would want telephones, but it turned out that they were wrong. Farm managers used telephones to get prompt reports on prices and weathers. Farm households used telephones to summon doctors in emergencies and to alleviate the loneliness of lives lived far from neighbors and relatives. In 1902, relatively few farm households had telephones, but as the independent companies grew, so did the number of farm-based customers; by 1920, just under 39 percent of all farm households in the United States had telephone service (while only 34 percent of nonfarm households did).

All this competition in telephone service had the net effect that any economist could have predicted: prices for telephone service fell, even in the Bell system. In order to keep the system companies competitive, the central Bell company had to cut the rates that it charged its affiliates for the rental of phones, and these savings were passed on to consumers. In New York City, as just one example, rates fell from \$150 for 1,000 calls in 1880 to \$51 in 1915 (figures adjusted for inflation).

As a result, in the period between 1894 and 1920, the telephone network expanded profoundly. Middle-class people began to pay for telephone service to their homes. Farm households became part of the telephone network (in record numbers). Retail businesses began to rely on telephones in their relations with their customers. By 1920, there were 13 million telephones in use in the country, 123 for every 1,000 people. Eight million of those 13 million phones belonged to Bell and 4 million to independent companies that connected to Bell lines. In just forty years, the telephone network, which provided point-to-point voice communication, had joined the telegraph, railroad, and petroleum networks as part of the economic and social foundation of industrial society.

The Electric System

Like the telegraph and telephone systems, the electric system was (and still is) quite literally a network of wires. Physicists, who had been experimenting with electricity since the middle of the eighteenth century, knew that under certain conditions electricity could produce light. Unfortunately, the first devices invented for generating a continuous flow of electricity—batteries—did not create a current strong enough for illumination. However, in 1831 the British experimenter Michael Faraday perfected a device that was based on a set of observations that scientists had made a decade earlier: an electric current will make a magnet move and a moving magnet will create an electric current. Faraday built an electric generator (a rotating magnet with a conducting wire round around it)—a device that could, unlike the battery, create a continuous flow of current strong enough to be used for lighting.

Within a short time, the generator was being used to power arc lamps in which the light (and a lot of heat) was produced by sparking across a gap in the conducting wires. Arc lamps were first used in British and French lighthouses in the 1860s; the generator that created the electricity was powered by a steam engine. A few years later, arc lamps were also being used for street lighting in some American cities. Unfortunately, arc lamps were dangerous; they had to be placed very far away from people and from anything that might be ignited by the sparks. By the mid-1870s, several people in several different countries were racing with each other to find a safer form of electrical lighting, the incandescent lamp. In

such a lamp, light would be derived from a glowing, highly resistant filament and not a spark; but the filament had to be kept in a vacuum so that it wouldn't oxidize (and disappear) too fast.

Thomas Alva Edison won the race. In 1878, when Edison started working on electrical lighting, he already had amassed a considerable reputation (and a moderate fortune) as an inventor. His first profitable invention had been the quadruplex telegraph, which could carry four messages at once, and he had also made successful modifications to the stock ticker, the telegraph system for relaying stock prices from the floor of the stock exchange to the offices of investors and brokers. These inventions had enhanced his reputation with Wall Street financiers and attorneys. In 1876, when he decided to become an independent inventor, building and staffing his own laboratory in Menlo Park, New Jersey, and again in 1878, when he decided that he wanted his laboratory to crack the riddle of electric lighting, he had no trouble borrowing money to invest in the enterprise.

Actually, they were enterprises. From the beginning, Edison understood that he wanted to build a technological system and a series of businesses to manage that system. The first of these businesses was the Edison Electric Light Company, incorporated for the purpose of financing research and development of electric lighting. Most of the stock was purchased by a group of New York financiers; Edison received stock in return for the rights to whatever lighting patents he might develop. Once Edison had actually invented a workable lightbulb (it had a carbonized thread as its filament), he proceeded to design other devices, and create other companies, that would all be parts of the system. The Edison Electric Illuminating Company of New York, founded in 1880, was created to build and maintain the very first central generating station providing electric service to customers. When this station opened its doors in 1882 (as its site Edison chose the part of Manhattan with the highest concentration of office buildings), it contained several steam-driven generators (built to Edison's design by the Edison Machine Company) and special cables to carry the electricity underground (made by the Edison Electric Tube Company). Customers who signed up for electric service had their usage measured by meters that Edison had invented; their offices were outfitted with lamp sockets that Edison had designed into which they were to place lightbulbs that another Edison company manufactured.

Information about this new system spread very fast (thanks to publicity generated by the Edison Electric Light Company), and within a few months (not even years), entrepreneurs were applying to Edison for licenses to build electric generating plants all over the country, indeed all over the world. Having been designed as a system, the electrical network grew very fast. There was only one generating plant in the country in 1882, but by 1902, there were 2,250, and by 1920, almost 4,000. These plants had a total generating capacity of 19 million kilowatts. Just over a third of the nation's homes were wired for electricity by 1920, by which time electricity was being used not only for lighting but also for cooling (electric fans), ironing (the electric iron replaced the so-called sad iron quickly), and vacuuming (the vacuum cleaner was being mass-produced by 1915).

The Edison companies (some of which eventually merged with other companies to become the General Electric Company) were not, however, able to remain in control of the electric system for as long (or as completely) as the Bell companies were able to dominate the telephone business or Standard Oil the petroleum business. Part of the reason for this lay in the principles of electromagnetic induction, which can be used to create electric motors as well as electric generators. The same experimenters who were developing electric generators in the middle years of the nineteenth century were also developing electric motors, and one of the first applications of those motors was in a business very different from the lighting business: electric traction for electric intraurban streetcars, often known as trolley cars. The first of these transportation systems was installed in Richmond, Virginia, in 1888 by a company owned by Frank Sprague, an electrical engineer who had briefly worked for Edison.

Sprague had invented an electric motor that, he thought, would be rugged enough to power carriages running day in and day out on city streets. As it turned out, the motor had to be redesigned, and redesigned again, before it worked very well, and Sprague also had to design trolley poles (for conducting the electricity from the overhead wires to the carriage) and a controlling system (so that the speed of the motor could be varied by the person driving the carriage). In the end, however, the electric streetcar was successful, and the days of the horse-pulled carriage were clearly numbered. Fourteen years after Sprague's first system began operating, the nation had 22,576 miles of track devoted to street railways.

Electric motors were also being used in industry. The earliest motors, like the streetcar motors, had been direct current (d.c.) motors, which needed a special and often fragile device (called a commutator) to transform the alternating current (a.c.) produced by generators. In 1888, an a.c. motor was invented by Nikola Tesla, a Serbian physicist who had emigrated to the United States. Tesla's patents were assigned to

the Westinghouse Company, which began both to manufacture and to market them. At that point, the use of electric motors in industry accelerated. The very first factory to be completely electrified was a cotton mill, built in 1894. As electric motors replaced steam engines, factory design and location changed; it was no longer necessary to build factories that were several stories high (to facilitate power transmission from a central engine) or to locate them near water sources (to feed the steam boilers). The first decade of the twentieth century was a turning point in the use of electric power in industry as more and more factories converted; by 1901, almost 400,000 motors had been installed in factories, with a total capacity of almost 5 million horsepower.

In short, the electrical system was more complex than the telephone and petroleum systems because it consisted of several different subsystems (lighting, traction, industrial power) with very different social goals and economic strategies; because of its complexity, no single company could dominate it. By 1895, when the first generating plant intended to transmit electricity over a long distance became operational (it was a hydroelectric plant built to take advantage of Niagara Falls, transmitting electricity twenty miles to the city of Buffalo), there were several hundred companies involved in the electric industry: enormous companies such as Westinghouse and General Electric that made everything from generators to lightbulbs; medium-sized companies, such as the ones that ran streetcar systems or that provided electric service to relatively small geographic areas; and small companies, which made specialized electric motors or parts for electric motors. Despite this diversity, the electric system was unified by the fact that its product, electric energy, had been standardized. By 1910, virtually all the generating companies (which, by now, had come to be called utility companies) were generating alternating current at sixty cycles per second. This meant that all electric appliances were made to uniform specifications and all transmission facilities could potentially be connected to one another. By 1920, electricity had supplanted gas, kerosene, and oils for lighting. In addition, it was being used to power sewing machines in ready-made clothing factories, to separate aluminum from the contaminants in its ores, to run projectors through which motion pictures could be viewed, to carry many thousands of commuters back and forth, and to do dozens of other chores in workplaces and residences. As transmission towers marched across the countryside and yet another set of wire-carrying poles were constructed on every city street, few Americans demonstrated any inclination to decline the conveniences that the youngest technical system—electricity—was carrying in its wake.

The Character of Industrialized Society

As inventors, entrepreneurs, and engineers were building all these multifarious technological systems, Americans were becoming increasingly dependent on them. Each time a person made a choice—to buy a kerosene lamp or continue to use candles, to take a job in an electric lamp factory or continue to be a farmer, to send a telegraph message instead of relying on the mails, to put a telephone in a shop so that customers could order without visiting—that person, whether knowingly or not, was becoming increasingly enmeshed in a technological system. The net effect of all that construction activity and all those choices was that a wholly new social order, and wholly different set of social and economic relationships between people, emerged: industrial society.

In industrial societies, manufactured products play a more important economic role than agricultural products. More money is invested in factories than in farms; more bolts of cloth are produced than bales of hay; more people work on assembly lines than as farm laborers. Just over half (53 percent) of what was produced in the United States was agricultural in 1869 and only a third (33 percent) was manufactured. In 1899 (just thirty years later), those figures were reversed: half the nation's output was in manufactured goods and only a third was agricultural, despite the fact that the nation's total farm acreage had increased rapidly as a result of westward migration. Manufacturing facilities were turning out products that were becoming increasingly important aspects of everyday life: canned corn and lightbulbs, cigarettes and underwear.

In a preindustrial society, the countryside is the base for economic and political power. In such societies, most people live in rural districts. Most goods that are traded are agricultural products; the price of fertile land is relatively high; and wealth is accumulated by those who are able to control that land. Industrialized societies are dominated by their cities. More people live and work in cities than on farms; most goods are manufactured in cities; most trade is accomplished there; wealth is measured in money and not in land. Furthermore, the institutions that control money—banks—are urban institutions.

As the nineteenth century progressed, more and more Americans began living either in the rural

towns in which factories were located (which, as a result, started to become small cities) or in the older cities that had traditionally been the center of artisanal production and of commerce. Native-born Americans began moving from the countryside to the city; many newly arrived Americans (and there were millions of newcomers to America in the nineteenth century) settled in cities. Just over half of all Americans (54 percent) were farmers or farm laborers in 1870, but only one in three was by 1910. Some American families underwent the rural-urban transition slowly: a daughter might move off the farm to a rural town when she married, and then a granddaughter might make her fortune in a big city. Others had less time: a man might be tending olive groves in Italy one day and working in a shoe factory in Philadelphia two months later.

During the 1840s, the population of the eastern cities nearly doubled, and several midwestern cities (St. Louis, Chicago, Pittsburgh, Cincinnati) began to grow. In 1860, there were nine port cities that had populations over 100,000 (Boston, New York, Brooklyn, Philadelphia, Baltimore, New Orleans, Chicago, Cincinnati, and St. Louis)—by 1910, there were fifty. Just as significantly, the country's largest cities were no longer confined to the eastern seaboard or to the Midwest. There were several large cities in the plains states, and half the population of the far west was living not in its fertile valleys or at the feet of its glorious mountains, but in its cities: Los Angeles, Denver, San Francisco, Portland, and Seattle. By 1920, for the first time in the nation's history, just slightly over half of all Americans lived in communities that had more than 10,000 residents.

Money was flowing in the same direction that people were; by 1900, the nation's wealth was located in its cities, not in its countryside. The nation's largest businesses and its wealthiest individuals were in its cities. J. P. Morgan and Cornelius Vanderbilt controlled their railroad empires from New York; Leland Stanford and Charles Crocker ran theirs from San Francisco; John D. Rockefeller operated from Cleveland and New York; Andrew Carnegie, at least initially, from Pittsburgh. Probably by 1880, and certainly by 1890, stock exchanges and investment bankers had become more important to the nation's economic health than cotton wharves and landed gentry.

This transition to an urban society had political consequences because political power tends to follow the trail marked out by wealth (and, in a democracy, to some extent by population). In the early years of the nineteenth century, when the independent political character of the nation was being formed, most Americans still lived on farms and American politics was largely controlled by people who earned their living directly from the land. After the Civil War, city residents (being both more numerous and more wealthy) began to flex their political muscles and to express their political interests more successfully. The first twelve presidents of the United States had all been born into farming communities, but from 1865 until 1912, the Republican party, then the party that most clearly represented the interests of big business and of cities, controlled the White House for all but eight years, and those eight years were the two terms served by Grover Cleveland, who before becoming president had been the mayor of Buffalo, New York.

The transition to an urban society also had economic and technological consequences. In a kind of historical feedback loop, industrialization caused cities to grow and the growth of cities stimulated more industrialization. Nineteenth-century cities were, to use the term favored by urban historians, walking cities. Since most residents could not afford either the cost or the space required to keep a horse and carriage, they had to be able to walk to work or to work in their own homes. Since businesses also had to be within walking distance of each other, this meant that as cities grew they became congested; more and more people had both to live and to work within the same relatively limited space. With congestion came disease; all nineteenth-century American cities were periodically struck by devastating epidemics: cholera, dysentery, typhoid fever.

Even before they understood the causes of these epidemics, city governments became convinced that they had to do something both to relieve the congestion and to control the diseases. Streets had to be paved, running water provided, sewers constructed, new housing encouraged. This meant that reservoirs had to be built, aqueducts and pumping stations constructed, trenches dug, pipes purchased, brickwork laid, new construction techniques explored. All of this municipal activity not only stimulated American industry but also served as a spur to the growth of civil engineering.

In addition, in the years between 1870 and 1920, many American cities actively stimulated industrialization by seeking out manufacturing interests and offering operating incentives to them. Many of the nation's older cities found themselves in economic trouble as railroad depots became more important than ports as nodes in the country's transportation system. In their distress, these cities decided that their futures lay not in commerce but in manufacturing, and they began to seek out manufacturing entrepreneurs to encourage industrial growth. By that time, the steam engine having been perfected and its manufacture

made relatively inexpensive, manufacturers had ceased to depend on waterwheels as a power source, which meant that they could easily (and profitably) establish their enterprises in cities rather than in the countryside; the development of the electric motor only served to increase this potential.

Minneapolis became a center of flour milling, Kansas City of meatpacking, Memphis of cotton seed oil production, Rochester of shoe manufacture, Schenectady of electric equipment, New York of ready-made clothing, Pittsburgh of steel and glass manufacture. Local banks helped manufacturers start up in business and local politicians helped recruit a docile labor force, all in the interests of stabilizing or augmenting a city's economy. Nationwide the net result was a positive impetus to the growth of industry; the processes of industrialization and urbanization are mutually reinforcing.

If American cities grew prodigiously during the second half of the nineteenth century, so, too, did the American population as a whole: between 1860 and 1920, the population of the United States more than tripled (from 31 million to 106 million). Some of the increase was the result of a high natural birthrate; in general, American families were larger than what is needed to keep a population at a stable size from one generation to the next. In addition, as the result of improvements in public health and improvements in the food supply, the death rate was declining and life expectancy was rising. People were living longer and that meant that in any given year a declining proportion of the total population was dying. On top of this, immigrants were arriving in record numbers. The figures are astounding; the total, between the end of the Civil War and the passage of the Immigration Restriction Acts (1924), came to over 30 million people. Like their native-born contemporaries, immigrants had a high birthrate and a declining death rate and more of their children lived past infancy and then enjoyed a longer life expectancy, all of which further contributed to the mushrooming size of the American population. This startling population increase—almost 20 percent per decade—reflects another crucial difference between societies that have become industrialized and those that have not. In a preindustrialized society the size of the population changes in a more or less cyclical fashion. If the weather cooperates and the crops are bounteous and peace prevails, people remain reasonably healthy and many children live past infancy; over the course of time the population will grow. But eventually the population will grow too large to be supported by the available land or the land itself will become infertile. Droughts may come or heavy rains; locusts may infest the fields or diseases may strike the cattle. Men will be drawn off to battle just when it is time to plow the fields or soldiers engaged in battles will trample the wheat and burn the barns. Then starvation will ensue. People will succumb to disease; fewer children will be born, and more of them will die in infancy. The population will shrink.

Under preindustrial conditions, such population cycles have been inexorable. Sometimes the cycle will take two generations to recur, sometimes two centuries, but it has recurred as long as there have been agricultural peoples who have been keeping records of themselves. Industrialization breaks this cyclical population pattern. Once a country has industrialized, natural disasters and wars do not seem to have a long-term effect on the size of its population; the rate of increase may slow for a few years or so, but there is still an increase. And the standard of living keeps rising as well. People stay relatively healthy; they live longer lives. Generally speaking, they can have as many (or as few) children as they want, knowing that, also generally speaking, most of their children will live past infancy. This is the salient characteristic that makes underdeveloped countries long for development: industrialized countries seem able to support extraordinarily large populations without any long-term collapse either in the size of the population or in the standard of living.

Industrialized countries can do this because agriculture industrializes at the same time that manufacturing does. In the transition to industrialization, what is happening on the farm is just as important as what is happening in the factories since, to put it bluntly, people cannot work if they cannot eat. These social processes—sustained growth of the population and the industrialization of agriculture—are interlocked. Both were proceeding rapidly in the United States between the years 1870 and 1920 as American farmers simultaneously pushed west and industrialized, settling new territory and developing more productive farming techniques. As the frontier moved westward roughly 400 million new acres were put under cultivation: virgin prairie became farms, fertile mountain valleys were planted in orchards, grassy hills became grazing land for sheep and cattle. The total quantity of improved acreage (meaning land that had been cleared or fenced or otherwise made suitable for agricultural use) in the United States multiplied two and a half times between 1860 and 1900.

This alone would have considerably expanded the nation's agricultural output, but newly introduced agricultural implements profoundly altered the work process of farming (particularly grain growing) and increased its productivity. The first of these was the reaper (patented by Cyrus McCormick in

1834 and in limited use even before the Civil War). The reaper, which was pulled by horses, replaced hand labor. Once a reaper had been purchased, a farm owner could quadruple the amount of acreage cut in one day or fire three day laborers who had previously been employed for the harvest or greatly increase the acreage put to plow (since the number of acres planted had always been limited by what could be reaped in the two prime weeks of harvest). The reaper was followed by the harvester (which made binding the grain easier), followed by the self-binder (which automatically bound the grain into shocks), and—in the far west—followed by the combine, a steam-driven tractor (which cut a swath of over forty feet, then threshed and bagged the grain automatically, sometimes at the rate of three 150-pound bags a minute). In those same years, haymaking was altered by the introduction of automatic cutting and baling machinery, and plowing was made considerably easier by the invention of the steel plow (John Deere, 1837) and the chilled-iron plow (James Oliver, 1868), both of which had the advantage of being nonstick surfaces for the heavy, wet soils of the prairies.

The net result, by 1900, was that American farmers were vastly more productive than they had been in 1860. Productivity has two facets: it is a measure both of the commodities being produced and of the labor being used to produce them. Statistics on wheat production indicate how radically American agriculture was changing in the second half of the nineteenth century. In 1866, there were roughly 15.5 million acres devoted to wheat production in the United States; farmers achieved average yields 9.9 bushels per acre, resulting in a total national production of about 152 million bushels. By 1898, acreage had roughly trebled (to 44 million), yields had almost doubled (to 15.3 bushels per acre), and the total production was 675 million bushels.

All this was accomplished with a marked saving of labor. By the hand method, 400 people and 200 oxen had to work ten hours a day to produce 20,000 bushels of wheat; by the machine method, only 6 people (and 36 horses) were required. Farms were getting larger, ownership was being restricted to a smaller and smaller number of people and more machinery was required for profitable farming (between 1860 and 1900, the annual value of farm implements manufactured in the United States went from \$21 million to \$101 million)—at the same time, the farms were becoming more productive.

What this means, put another way, was that a smaller proportion of the nation's people were needed to produce the food required by its ever larger population. Some people left their farms because they hated the farming life, some because they could not afford to buy land as prices began to rise, some because they were forced off the land by the declining profitability of small farms. The farming population (this includes both owners and laborers) began to shrink in relation to the rest of the population.

New transportation facilities and new food-based industries made it easier and cheaper for the residents of cities and towns to eat a more varied diet. The fledgling canning industry was spurred by the need to supply food for troops during the Civil War. After the war, the canners turned to the civilian market, and by the 1880s, urban Americans had become accustomed to eating canned meat, condensed milk (invented by Gail Borden in 1856), canned peas, and canned corn. The Heinz company was already supplying bottled ketchup and factory pickles to a vast population, and the Campbell's company was just about to start marketing soups. By 1900, cheese and butter making had become largely a factory operation, made easier and cheaper by the invention of the centrifugal cream separator in 1879.

After the Civil War, the railroads replaced steamboats and canal barges as the principal carriers of farm products (from wheat to hogs, from apples to tobacco), thus both shortening the time required to bring goods to market and sharply lowering the cost of transportation. After the 1880s, when refrigerated transport of various kinds was introduced, this trend accelerated: even more products could be brought to market (butchered meat, for example, or fresh fish) in an even shorter time. New refrigeration techniques transformed beer making from a home to a factory operation; by 1873, there were some 4,000 breweries in the United States with an output of 10 million barrels a year. Commercial baking had also expanded and Americans were becoming fond of factory-made crackers and cookies. In the end, then, another historical feedback loop had been established, a loop connecting industrialization with agricultural change. Industrialization made farming more productive, which made it possible for the population to increase, which created a larger market for manufactured goods, which increased the rate of industrialization.

Conclusion: Industrialization and Technological Systems

By 1920, a majority of Americans had crossed the great divide between preindustrial and industrial societies. The foods they ate, the conditions under which they worked, the places in which they lived—all

had been transformed. The majority of Americans were no longer living on farms. They were eating food that had been carried to them by one technological system (the railroad) after having been processed by machines that were powered by a second (electricity) and lubricated by a third (petroleum). If they wanted to light their domiciles at night or heat their dwelling places during cold weather, they could not avoid interacting with one or another technological system for distributing energy—unless they were willing to manufacture their own candles (even then, they might have ended up buying paraffin from Standard Oil). The social ties that bound individuals and communities together—someone has been elected, someone else has died, young men are about to be drafted, a young woman has given birth—were being carried over, communicated through, and to some extent controlled by technological networks that were owned by large, monopolistically inclined corporations. More people were living longer lives; fewer babies were dying in infancy; the standard of living for many Americans (albeit not for all) was rising. And at the very same time, because of the very same processes, people were becoming more dependent on each other.

Early in the nineteenth century the process of industrialization had appeared (to those who were paying attention) as a rather discrete undertaking: a spinning factory in a neighboring town, a merchant miller up the river, a railroad station a few miles distant. By the end of the century, virtually all Americans must have been aware that it had become something vastly different: a systematic undertaking that had created interlocking physical and social networks in which all Americans—rich or poor, young or old, urban or rural—were increasingly enmeshed.

Suggestions for Further Reading

Alfred D. Chandler, *Strategy and Structure: Chapters in the History of American Enterprise* (Garden City, NY, 1966).

Howard P. Chudacoff and Judith E. Smith, *The Evolution of American Urban Society* (Englewood Cliffs, NJ, 1988).

Robert W. Garnet, *The Telephone Enterprise* (Baltimore, 1985).

Richard F. Hirsh, *Technology and Transformation in the American Electric Utility Industry* (Cambridge, 1989).

August Giebelhaus, *Business and Government in the Oil Industry: A Case Study of Sun Oil* (Greenwich, CT, 1980).

Thomas Parke Hughes, *American Genesis: A Century of Invention and Technological Enthusiasm, 1870—1970* (New York, 1989).

Thomas Parke Hughes, *Networks of Power. The Electrification of Society, 1880—1920* (Baltimore, 1983).

Malcolm MacLaren, *The Rise of the Electrical Industry During the 19th Century* (Princeton, 1943).

Glen Porter, *The Rise of Big Business: 1860—1910* (New York, 1973).

Mark H. Rose, *Cities of Light and Heat: Domesticating Gas and Electricity in American Homes* (University Park, PA, 1995).

Nathan Rosenberg, *Exploring the Black Box: Technology, Economics and History* (Cambridge, 1982).

Nathan Rosenberg, *Technology and American Economic Growth* (New York, 1972).

George David Smith, *The Anatomy of a Business Strategy: Bell, Western Electric and the Origins of 'the American Telephone Industry* (Baltimore, 1985).

Carlene Stephens, *Inventing Standard Time* (Washington, DC, 1983).

John F. Stover, *American Railroads* (Chicago, 1961).

Neil H. Wasserman, *From Invention to Innovation: Long Distance Telephone Transmission at the Turn of the Century* (Baltimore, 1985).

Harold Williamson and Arnold Daum, *The American Petroleum Industry*. 2 vols. (Evanston, IL, 1959).

The Lexus and the Olive Tree

by

Thomas L. Friedman

Anchor Books, Random House, Inc, New York, 1999

Chapter 1, “The New System”

When I say that globalization has replaced the Cold War as the defining international system, what exactly do I mean?

I mean that, as an international system, the Cold War had its own structure of power: the balance between the United States and the U.S.S.R. The Cold War had its own rules: in foreign affairs, neither superpower would encroach on the other’s sphere of influence; in economics, less developed countries would focus on nurturing their own national industries, developing countries on export-led growth, communist countries on autarky and Western economies on regulated trade. The Cold War had its own dominant ideas: the clash between communism and capitalism, as well as detente, nonalignment and perestroika. The Cold War had its own demographic trends: the movement of people from east to west was largely frozen by the Iron Curtain, but the movement from south to north was a more steady flow. The Cold War had its own perspective on the globe: the world was a space divided into the communist camp, the Western camp, and the neutral camp, and everyone’s country was in one of them. The Cold War had its own defining technologies: nuclear weapons and the second Industrial Revolution were dominant, but for many people in developing countries the hammer and sickle were still relevant tools. The Cold War had its own defining measurement: the throw weight of nuclear missiles. And lastly, the Cold War had its own defining anxiety: nuclear annihilation. When taken all together the elements of this Cold War system influenced the domestic politics, commerce and foreign relations of virtually every country in the world. The Cold War system didn’t shape everything, but it shaped many things

Today’s era of globalization is a similar international system, with its own unique attributes, which contrast sharply with those of the Cold War. To begin with the Cold War system was characterized by one overarching feature—division. The world was a divided-up, chopped-up place and both your threats and opportunities in the Cold War system tended to grow out of who you were divided from. Appropriately, this Cold War system was symbolized by a single word: the *wall*—the Berlin Wall. One of my favorite descriptions of that world was provided by Jack Nicholson in the movie *A Few Good Men*. Nicholson plays it Marine colonel who is the commander of the U.S. base in Cuba, at Guantánamo Bay. In the climactic scene of the movie, Nicholson is pressed by Tom Cruise to explain how a certain weak soldier under Nicholson’s command, Santiago, was beaten to death by his own fellow Marines. “You want answers?” shouts Nicholson. “You want answers?” I want the truth, retorts Cruise. “You can’t handle the truth,” says Nicholson. “Son, we live in a world that has walls and those walls have to be guarded by men with guns. Who’s gonna do it? You? You, Lieutenant Weinberg? I have a greater responsibility than you can possibly fathom. You weep for Santiago and you curse the Marines. You have that luxury. You have the luxury of not knowing what I know—that Santiago’s death, while tragic, probably saved lives. And my existence, while grotesque and incomprehensible to you, saves lives. You don’t want the truth because deep down in places you don’t talk about at parties, you want me on that wall. You need me on that wall.”

The globalization system is a bit different. It also has one overarching feature—integration. The world has become an increasingly interwoven place, and today, whether you are a company or a country, your threats and opportunities increasingly derive from who you are connected to. This globalization system is also characterized by a single word: the *Web*. So in the broadest sense we have gone from a system built around division and walls to a system increasingly built around integration and webs. In the Cold War we reached for the “hotline,” which was a symbol that we were all divided but at least two people were in charge—the United States and the Soviet Union—and in the globalization system we reach for the Internet, which is a symbol that we are all increasingly connected and nobody is quite in charge.

This leads to many other differences between the globalization system and the Cold War system.

The globalization system, unlike the Cold War system, is not frozen, but a dynamic ongoing process. That's why I define globalization this way: it is the inexorable integration of markets, nation-states and technologies to a degree never witnessed before—in a way that is enabling individuals, corporations and nation-states to reach around the world further, faster, deeper and cheaper than ever before, and in a way that is enabling the world to reach into individuals, corporations and nation-states farther, faster, deeper, cheaper than ever before. This process of globalization is also producing a powerful backlash from those brutalized or left behind by this new system.

The driving idea behind globalization is free-market capitalism—the more you let market forces rule and the more you open your economy to free trade and competition, the more efficient and flourishing your economy will be. Globalization means the spread of free-market capitalism to virtually every country in the world. Therefore, globalization also has its own set of economic rules—rules that revolve around opening, deregulating and privatizing your economy, in order to make it more competitive and attractive to foreign investment. In 1975, at the height of the Cold War, only 8 percent of countries worldwide had liberal, free-market capital regimes, and foreign direct investment at the time totaled only \$23 billion, according to the World Bank. By 1997, the number of countries with liberal economic regimes constituted 28 percent, and foreign investment totaled \$644 billion.

Unlike the Cold War system, globalization has its own dominant culture, which is why it tends to be homogenizing to a certain degree. In previous eras this sort of cultural homogenization happened on a regional scale—the Romanization of Western Europe and the Mediterranean world, the Islamification of Central Asia, North Africa, Europe and the Middle East by the Arabs and later the Ottomans, or the Russification of Eastern and Central Europe and parts of Eurasia under the Soviets. Culturally speaking, globalization has tended to involve the spread (for better and for worse) of Americanization—from Big Macs to iMacs to Mickey Mouse.

Globalization has its own defining technologies: computerization, miniaturization, digitization, satellite communications, fiber optics and the Internet, which reinforce its defining perspective of integration. Once a country makes the leap into the system of globalization, its elites begin to internalize this perspective of integration, and always try to locate themselves in a global context. I was visiting Amman, Jordan, in the summer of 1998 and having coffee at the Inter-Continental Hotel with my friend Rami Khouri, the leading political columnist in Jordan. We sat down and I asked him what was new. The first thing he said to me was: "Jordan was just added to CNN's worldwide weather highlights." What Rami was saying was that it is important for Jordan to know that those institutions which think globally believe it is now worth knowing what the weather is like in Amman. It makes Jordanians feel more important and holds out the hope that they will be enriched by having more tourists or global investors visiting. The day after seeing Rami I happened to go to Israel and meet with Jacob Frenkel, governor of Israel's Central Bank and a University of Chicago-trained economist. Frenkel remarked that he too was going through a perspective change: "Before, when we talked about macroeconomics, we started by looking at the local markets, local financial systems and the interrelationship between them, and then, as an afterthought, we looked at the international economy. There was a feeling that what we do is primarily our own business and then there are some outlets where we will sell abroad. Now we reverse the perspective. Let's not ask what markets we should export to, after having decided what to produce; rather let's first study the global framework within which we operate and then decide what to produce. It changes your whole perspective."

While the defining measurement of the Cold War was weight—particularly the throw weight of missiles—the defining measurement of the globalization system is speed—speed of commerce, travel, communication and innovation. The Cold War was about Einstein's mass-energy equation, $e = mc^2$. Globalization tends to revolve around Moore's Law, which states that the computing power of silicon chips will double every eighteen to twenty-four months, while the price will halve. In the Cold War, the most frequently asked question was: "Whose side are you on?" In globalization, the most frequently asked question is: "To what extent are you connected to everyone?" In the Cold War, the second most frequently asked question was: "How big is your missile?" In globalization, the second most frequently asked question is: "How fast is your modem?" The defining document of the Cold War system was "The Treaty." The defining document of globalization is "The Deal." The Cold War system even had its own style. In 1961, according to *Foreign Policy* magazine, Cuban President Fidel Castro, wearing his usual olive drab military uniform, made his famous declaration "I shall be a Marxist-Leninist for the rest of my life." In January 1999, Castro put on a business suit for a conference on globalization in Havana, to which financier George Soros and free-market economist Milton Friedman were both invited.

If the defining economists of the Cold War system were Karl Marx and John Maynard Keynes,

who each in his own way wanted to tame capitalism, the defining economists of the globalization system are Joseph Schumpeter and Intel chairman Andy Grove, who prefer to unleash capitalism. Schumpeter, a former Austrian Minister of Finance and Harvard Business School professor, expressed the view in his classic work *Capitalism, Socialism and Democracy* that the essence of capitalism is the process of “creative destruction”—the perpetual cycle of destroying the old and less efficient product or and replacing it with new, more efficient ones. Andy Grove took Schumpeter’s insight that “only the paranoid survive” for the title of his book on life in Silicon Valley, and made it in many ways the business model of globalization capitalism. Grove helped to popularize the view that dramatic, industry-transforming innovations are taking place today faster and faster. Thanks to these technological breakthroughs, the speed by which your latest invention can be made obsolete or turned into a commodity is now lightning quick. Therefore, only the paranoid, only those who are constantly looking over their shoulders to see who is creating something new that will destroy them and then staying just one step ahead of them, will survive. Those countries that are most willing to let capitalism quickly destroy inefficient companies, so that money can be freed up and directed to more innovative ones, will thrive in the era of globalization. Those which rely on their governments to protect them from such creative destruction will fall behind in this era.

James Surowiecki, the business columnist for *Slate* magazine, reviewing Grove’s book, neatly summarized what Schumpeter and Grove have in common, which is the essence of globalization economics. It is the notion that: “Innovation replaces tradition. The present—or perhaps the future—replaces the past. Nothing matters so much as what will come next, and what will come next can only arrive if what is here now gets overturned. While this makes the system a terrific place for innovation, it makes it a difficult place to live, since most people prefer some measure of security about the future to a life lived in almost constant uncertainty ... We are not forced to re-create our relationships with those closest to us on a regular basis. And yet that’s precisely what Schumpeter, and Grove after him, suggest is necessary to prosper [today].”

Indeed, if the Cold War were a sport, it would be sumo wrestling, says Johns Hopkins University foreign affairs professor Michael Mandelbaum. “It would be two big fat guys in a ring, with all sorts of posturing and rituals and stomping of feet, but actually very little contact, until the end of the match, when there is a brief moment of shoving and the loser gets pushed out of the ring, but nobody gets killed.”

By contrast, if globalization were a sport, it would be the 100-meter dash, over and over and over. And no matter how many times you win, you have to race again the next day. And if you lose by just one-hundredth of a second it can be as if you lost by an hour. (Just ask French multinationals. In 1999, French labor laws were changed, requiring every employer to implement a four-hour reduction in the workweek, from 39 hours to 35 hours, with no cut in pay. Many French firms were fighting the move because of the impact it would have on their productivity in a global market. Henri Thierry, human resources director for Thomson-CSF Communications, a high-tech firm in the suburbs of Paris, told *The Washington Post*: “We are in a worldwide competition. If we lose one point of productivity, we lose orders. If we’re obliged to go to 35 hours it would be like requiring French athletes to run the 100 meters wearing flippers. They wouldn’t have much of a chance winning a medal.”)

To paraphrase German political theorist Carl Schmitt, the Cold War was a world of “friends” and “enemies.” The globalization world, by contrast, tends to turn all friends and enemies into “competitors.”

If the defining anxiety of the Cold War was fear of annihilation from an enemy you knew all too well in a world struggle that was fixed and stable, the defining anxiety in globalization is fear of rapid change from an enemy you can’t see, touch or feel—a sense that your job, community or workplace can be changed at any moment by anonymous economic and technological forces that are anything but stable. The defining defense system of the Cold War was radar—to expose the threats coming from the other side of the wall. The defining defense system of the globalization era is the X-ray machine—to expose the threats coming from within.

Globalization also has its own demographic pattern—a rapid acceleration of the movement of people from rural areas and agricultural lifestyles to urban areas and urban lifestyles more intimately linked with global fashion, food, markets and entertainment trends.

Last, and most important, globalization has its own defining structure of power, which is much more complex than the Cold War structure. The Cold War system was built exclusively around nation-states. You acted on the world in that system through your state. The Cold War was primarily a drama of states confronting states, balancing states and aligning with states. And, as a system, the Cold War was balanced at the center by two superstates: the United States and the Soviet Union.

The globalization system, by contrast, is built around three balances, which overlap and affect one

another. The first is the traditional balance between nation-states. In the globalization system, the United States is now the sole and dominant superpower and all other nations are subordinate to it to one degree or another. The balance of power between the United States and the other states, though, still matters for the stability of this system. And it can still explain a lot of the news you read on the front page of the papers, whether it is the containment of Iraq in the Middle East or the expansion of NATO against Russia in Central Europe.

The second balance in the globalization system is between nation-states and global markets. These global markets are made up of millions of investors moving money around the world with the click of a mouse. I call them “the Electronic Herd,” and this herd gathers in key global financial centers, such as Wall Street, Hong Kong, London and Frankfurt, which I call “the Supermarkets.” The attitudes and actions of the Electronic Herd and the Supermarkets can have a huge impact on nation-states today, even to the point of triggering the downfall of governments. Who ousted Suharto in Indonesia in 1998? It wasn’t another state, it was the Supermarkets, by withdrawing their support for, and confidence in, the Indonesian economy. You will not understand the front page of newspapers today unless you bring the Supermarkets into your analysis. Because the United States can destroy you by dropping bombs and the Supermarkets can destroy you by downgrading your bonds. In other words, the United States is the dominant player in maintaining the globalization gameboard, but it is not alone in influencing the moves on that gameboard. This globalization gameboard today is a lot like a Ouija board—sometimes pieces are moved around by the obvious hand of the superpower, and sometimes they are moved around by hidden hands of the Supermarkets.

The third balance that you have to pay attention to in the globalization system—the one that is really the newest of all—is the balance between individuals and nation-states. Because globalization has brought down many of the walls that limited the movement and reach of people, and because it has simultaneously wired the world into networks, it gives more power to individuals to influence both markets and nation-states than at any time in history. Individuals can increasingly act on the world stage directly—unmediated by a state. So you have today not only a superpower, not only Supermarkets, but, as will be demonstrated later in the book, you now have Super-empowered individuals. Some of these Super-empowered individuals are quite angry, some of them quite wonderful—but all of them are now able to act directly on the world stage.

Without the knowledge of the U.S. government, Long-Term Capital Management—a few guys with a hedge fund in Greenwich, Connecticut—amassed more financial bets around the world than all the foreign reserves of China. Osama bin Laden, a Saudi millionaire with his own global network, declared war on the United States in the late 1990s, and the U.S. Air Force retaliated with its cruise missile attack on him (where he resided in Afghanistan) as though he were another nation-state. Think about that. The United States fired 75 cruise missiles, at \$1 million apiece, at a person! That was a superpower against a Super-empowered angry man. Jody Williams won the Nobel Peace Prize in 1997 for her contribution to the international ban on landmines. She achieved that ban not only without much government help, but in the face of opposition from all the major powers. And what did she say was her secret weapon for organizing 1,000 different human rights and arms control groups on six continents? “E-mail.”

Nation-states, and the American superpower in particular, are still hugely important today, but so too now are Supermarkets and Super-empowered individuals. You will never understand the globalization system, or the front page of the morning paper, unless you see it as a complex interaction between all three of these actors: states bumping up against states, states bumping up against Supermarkets, and Supermarkets and states bumping up against Super-empowered individuals.

Unfortunately, for reasons I will explain later, the system of globalization has come upon us far faster than our ability to retrain ourselves to see and comprehend it. Think about just this one fact: Most people had never even heard of the Internet in 1990, and very few people had an E-mail address then. That was just ten years ago! But today the Internet, cell phones and E-mail have become essential tools that many people, and not only in developed countries, cannot imagine living without. It was no different, I am sure, at the start of the Cold War, with the first appearance of nuclear arsenals and deterrence theories. It took a long time for leaders and analysts of that era to fully grasp the real nature and dimensions of the Cold War system. They emerged from World War II thinking that this great war had produced a certain kind of world, but they soon discovered it had laid the foundations for a world very different from the one they anticipated. Much of what came to be seen as great Cold War architecture and strategizing were responses on the fly to changing events and evolving threats. Bit by bit, these Cold War strategists built the institutions, the perceptions and the reflexes that came to be known as the Cold War system.

It will be no different with the globalization system, except that it may take us even longer to get our minds around it, because it requires so much retraining just to see this new system and because it is built not just around superpowers but also around Supermarkets and Super-empowered individuals. I would say that in 2000 we understand as much about how today's system of globalization is going to work as we understood about how the Cold War system was going to work in 1946—the year Winston Churchill gave his speech warning that an “Iron Curtain” was coming down, cutting off the Soviet zone of influence from Western Europe. We barely understood how the Cold War system was going to play out thirty years after Churchill's speech! That was when Routledge published a collection of essays by some of the top Sovietologists, entitled *Soviet Economy Towards the Year 2000*. It was a good seller when it came out. It never occurred at that time to any of the authors that there wouldn't be a Soviet economy in the year 2000.

If you want to appreciate how few people understand exactly how this system works, think about one amusing fact. The two key economists who were advising Long-Term Capital Management, Robert C. Merton and Myron S. Scholes, shared the Nobel Prize for economics in 1997, roughly one year before LTCM so misunderstood the nature of risk in today's highly integrated global marketplace that it racked up the biggest losses in hedge fund history. And what did LTCM's two economists win their Nobel Prize for? For their studies on how complex financial instruments, known as derivatives, can be used by global investors to offset risk! In 1997 they won the Nobel Prize for managing risk. In 1998 they won the booby prize for creating risk. Same guys, same market new world.

Excerpt from Chapter 3, “The Lexus and the Olive Tree”

Once you recognize that globalization is the international system that has replaced the Cold War system, is this all you need to know to explain world affairs today? Not quite. Globalization is what is new. And if the world were made of just microchips and markets, you could probably rely on globalization to explain almost everything. But, alas, the world is made of microchips and markets and men and women, with all their peculiar habits, traditions, longings and unpredictable aspirations. So world affairs today can only be explained as the interaction between what is as new as an Internet Web site and what is as old as a gnarled olive tree on the banks of the river Jordan. I first started thinking about this while riding on a train in Japan in May 1992, eating a sushi box dinner and traveling at 180 miles per hour.

I was in Tokyo on a reporting assignment and had arranged to visit the Lexus luxury car factory outside Toyota City, south of Tokyo. It was one of the most memorable tours I've ever taken. At that time, the factory was producing 300 Lexus sedans each day, made by 66 human beings and 310 robots. From what I could tell, the human beings were there mostly for quality control. Only a few of them were actually screwing in bolts or soldering parts together. The robots were doing all the work. There were even robotic trucks that hauled materials around the floor and could sense when a human was in their path and would “beep, beep, beep” at them to move. I was fascinated watching the robot that applied the rubber seal that held in place the front windshield of each Lexus. The robot arm would neatly paint the hot molten rubber in a perfect rectangle around the window. But what I liked most was that when it finished its application there was always a tiny drop of rubber left hanging from the tip of the robot's finger—like the drop of toothpaste that might be left at the top of the tube after you've squeezed it onto your toothbrush. At the Lexus factory, though, this robot arm would swing around in a wide loop until the tip met a tiny, almost invisible metal wire that would perfectly slice off that last small drop of hot black rubber—leaving nothing left over. I kept staring at this process, thinking to myself how much planning, design and technology it must have taken to get that robot arm to do its job and then swing around each time, at the precise angle, so that this little thumbnail-size wire could snip off the last drop of hot rubber for the robot to start clean on the next window. I was impressed.

After touring the factory, I went back to Toyota City and boarded the bullet train for the ride back to Tokyo. The bullet train is aptly named, for it has both the look and feel of a speeding bullet. As I nibbled away on one of those sushi dinner boxes you can buy in any Japanese train station, I was reading that day's *International Herald Tribune*, and a story caught my eye on the top right corner of page 3. It was about the daily State Department briefing. State Department spokeswoman Margaret D. Tutwiler had given a controversial interpretation of a 1948 United Nations resolution, relating to the right of return for Palestinian refugees to Israel. I don't remember all the details, but whatever her interpretation was, it had clearly agitated both the Arabs and the Israelis and sparked a furor in the Middle East, which this story was reporting.

So there I was speeding along at 180 miles an hour on the most modern train in the world, reading this story about the oldest corner of the world. And the thought occurred to me that these Japanese, whose Lexus factory I had just visited and whose train I was riding, were building the greatest luxury car in the world with robots. And over here, on the top of page 3 of the *Herald Tribune*, the people with whom I had lived for so many years in Beirut and Jerusalem, whom I knew so well, were still fighting over who owned which olive tree. It struck me then that the Lexus and the olive tree were actually pretty good symbols of this post-Cold War era: half the world seemed to be emerging from the Cold War intent on building a better Lexus, dedicated to modernizing, streamlining and privatizing their economies in order to thrive in the system of globalization. And half of the world—sometimes half the same country, sometimes half the same person—was still caught up in the fight over who owns which olive tree.

Olive trees are important. They represent everything that roots us, anchors us, identifies us and locates us in this world—whether it be belonging to a family, a community, a tribe, a nation, a religion or, most of all, a place called home. Olive trees are what give us the warmth of family, the joy of individuality, the intimacy of personal rituals, the depth of private relationships, as well as the confidence and security to reach out and encounter others. We fight so intensely at times over our olive trees because, at their best, they provide the feelings of self-esteem and belonging that are as essential for human survival as food in the belly. Indeed, one reason that the nation-state will never disappear, even if it does weaken, is because it is the ultimate olive tree—the ultimate expression of whom we belong to—linguistically, geographically and historically. You cannot be a complete person alone. You can be a rich person alone. You can be a smart person alone. But you cannot be a complete person alone. For that you must be part of, and rooted in, an olive grove. This truth was once beautifully conveyed by Rabbi Harold S. Kushner in his interpretation of a scene from Gabriel García Márquez’s classic novel *One Hundred Years of Solitude*:

Márquez tells of a village where people were afflicted with a strange plague of forgetfulness, a kind of contagious amnesia. Starting with the oldest inhabitants and working its way through the population, the plague causes people to forget the names of even the most common everyday objects. One young man, still unaffected, tries to limit the damage by putting labels on everything. “This is a table,” “This is a window,” “This is a cow; it has to be milked every morning.” And at the entrance to the town, on the main road, he puts up two large signs. One reads “The name of our village is Macondo,” and the larger one reads “God exists.” The message I get from that story is that we can, and probably will, forget most of what we have learned in life—the math, the history, the chemical formulas, the address and phone number of the first house we lived in when we got married—and all that forgetting will do us no harm. But if we forget whom we belong to, and if we forget that there is a God, something profoundly human in us will be lost.

But while olive trees are essential to our very being, an attachment to one’s olive trees, when taken to excess, can lead us into forging identities, bonds and communities based on the exclusion of others. And when these obsessions really run amok, as with the Nazis in Germany or the murderous Aum Shinrikyo cult in Japan or the Serbs in Yugoslavia, they lead to the extermination of others.

Conflicts between Serbs and Muslims, Jews and Palestinians, Armenians and Azeris over who owns which olive tree are so venomous precisely because they are about who will be at home and anchored in a local world and who will not be. Their underlying logic is: I must control this olive tree, because if the other controls it, not only will I be economically and politically under his thumb, but my whole sense of home will be lost. I’ll never be able to take my shoes off and relax. Few things are more enraging to people than to have their identity or their sense of home stripped away. They will die for it, kill for it, sing for it, write poetry for it and novelize about it. Because without a sense of home and belonging, life becomes barren and rootless. And life as a tumbleweed is no life at all.

In the Cold War system, the most likely threat to your olive tree was from another olive tree. It was from your neighbor coming over, violently digging up your olive tree and planting his in its place. That threat has not been eliminated today, but, for the moment, it has been diminished in many parts of the world. The biggest threat today to your olive tree is likely to come from the Lexus—from all the anonymous, transnational, homogenizing, standardizing market forces and technologies that make up today’s globalizing economic system. There are some things about this system that can make the Lexus so overpowering it can overrun and overwhelm every olive tree in sight—breaking down communities, steamrolling environments and crowding out traditions—and this can produce a real olive tree backlash. But there are other things about this system that empower even the smallest, weakest political community

to actually use the new technologies and markets to preserve their olive trees, their culture and identity. Traveling the world in recent years, again and again I have come on this simultaneous wrestling match, tug-of-war, balancing act between the Lexus and the olive tree.

The Lexus and olive tree wrestling with each other in the new system of globalization was reflected in Norway's 1994 referendum about whether or not to join the European Union. That should have been a slam dunk for Norwegians. After all, Norway is in Europe. It is a rich, developed country and it has a significant amount of Intra-European trade. Joining the EU made all the economic sense in the world for Norway in a world of increasing globalization. But the referendum failed, because too many Norwegians felt joining the EU would mean uprooting too much of their own Norwegian identity and way of life, which, thanks to Norwegian North Sea oil (sold into a global economy), the Norwegians could still afford to preserve—without EU membership. Many Norwegians looked at the EU and said to themselves, “Now let me get this straight. I am supposed to take my Norwegian identity and deposit it into a Euro-Cuisinart, where it will be turned into Euromush by Eurobureaucrats paid in Eurodollars at the Euro-Parliament in the Eurocapital covered by Eurojournalists? Hey, no, thanks. I'd rather be Sten from Norway. I'd rather cling to my own unique olive tree identity and be a little less efficient economically.”

The olive tree backlashing against the Lexus is the August 1999 story from France, by *The Washington Post's* Anne Swardson, about Philippe Folliot, the mayor of the southwestern French village of St. Pierre-de-Trivisy—population 610. Folliot and the St. Pierre-de-Trivisy town council slapped a 100-percent tax on bottles of Coca-Cola sold at the town's camp ground, in retaliation for a tariff that the United States had slapped on Roquefort cheese, which is produced only in the southwestern French region around St. Pierre-de-Trivisy. As he applied some Roquefort to a piece of crusty bread, Folliot told Swardson, “Roquefort is made from the milk of only one breed of sheep, it is made in only one place in France, and it is made in only one special way. It is the opposite of globalization. Coca-Cola you can buy anywhere in the world and it is exactly the same. [Coke] is a symbol of American multinational that wants to uniformize taste all over the planet. That's what we are against.”

The Lexus being exploited by the olive tree was the report in *The Economist* of August 14, 1999, entitled “Cyberthugs.” It stated that “The National Criminal Intelligence Service blamed the increasingly sophisticated nature of football hooligans for the organized violence last weekend between fans of Millwall and Cardiff City. Rival bands of thugs are apparently prepared to cooperate by fixing venues for fights via the Internet. Information is exchanged in closed or open Websites. Some even report the violence as it happens: ‘It's kicking off right now as I speak,’ wrote Paul Dodd, a particularly dopey hooligan known to cyber nerds and police alike. The police now say they surf for such Websites, hoping to discover other planned attacks.”

West Side Story meets the World Wide Web.

The olive tree exploiting the Lexus is the story that came to light in the summer of 1999 about Adolf Hitler's racist manifesto *Mein Kampf*, which is banned in Germany by the German government. You cannot sell it in any German bookstore, or publish it in Germany. But Germans found that they could order the book over the Internet from Amazon.com and it would come in the mail in a way that the German government was powerless to stop. Indeed, so many Germans ordered *Mein Kampf* from Amazon.com that in the summer of 1999 Hitler made Amazon.com's top-ten bestseller list for Germany. Amazon.com at first refused to stop shipping *Mein Kampf* to Germany, insisting that the English translation was not covered by censorship, and that it was not going to get in the business of deciding what its customers were allowed to read. However, after this was publicized, Amazon.com was so bombarded with angry E-mails from all over the world that it stopped selling Hitler's works.

The olive tree trumping the Lexus, and then the Lexus then coming right back to trump the olive tree, was the nuclear-testing saga that unfolded in India in the late 1990s. In the spring of 1998 India's newly elected nationalist Bharatiya Janata Party (BJP) decided to defy the world and resume testing its nuclear weapons. Asserting India's right to test had been a key plank in the BJP's election campaign. I visited India shortly after the tests, where I talked to rich and poor, government and nongovernment types, villagers and city slickers. I kept waiting to meet the Indian who would say to me, “You know, these nuclear tests were really stupid. We didn't get any additional security out of them and they've really cost us with sanctions.” I am sure that sentiment was there—but I couldn't find anyone to express it. Even those Indian politicians who denounced their nuclear tests as a cheap, jingoistic maneuver by India's new Hindu nationalist government would tell you that these tests were the only way for India to get what it wants most from the United States and China: R-E-S-P-E-C-T. I finally realized the depth of this sentiment when I

went to see a saffron-robed Indian human rights campaigner, Swami Agnivesh. As the two of us sat cross-legged on the floor of his living room in his simple Delhi home, I thought, "Surely he will disavow this test." But no sooner did we start talking than he declared to me: "We are India, the second-largest country in the world! You can't just take us for granted. India doesn't feel threatened by Pakistan, but in the whole international game India is being marginalized by the China-U.S. axis." The next day I went out to Dasna, a village north of New Delhi, where I randomly stopped shopkeepers to talk. Dasna is one of the poorest places I have ever seen. Nobody seemed to have shoes. Everyone seemed to be skin and bones. There were more water buffalo and bicycles than cars on the road. The air was heavy with the smell of cow dung used for energy. But they loved their government's nuclear sound-and-light show. "We are nine hundred million people. We will not die from these sanctions," pronounced Pramod Batra, the forty-two-year-old village doctor in Dasna. "This nuclear test was about self-respect, and self-respect is more important than roads, electricity and water. Anyway, what did we do? We exploded our bomb. It was like shooting a gun off into the air. We didn't hurt anybody."

But while India's olive tree impulse seemed to have prevailed over its needs for a Lexus, when this happens in today's globalization system there is always a hidden long-term price. While in New Delhi, I stayed at the Oberoi Hotel, where I swam laps in the pool at the end of each day to recover from the sweltering 100-degree heat. My first day there, while I was doing my breaststrokes, there was an Indian woman also swimming laps in the lane next to me. During a rest stop we started talking and she told me she ran the India office of Salomon Brothers-Smith Barney, the major American investment bank. I told her I was a columnist who had come over to write about the fallout from the Indian nuclear tests.

"Have you heard who's in town?" she asked me as we each trod water. "No," I said, shaking my head. "Who's in town?"

"Moody's," she said. Moody's Investors Service is the international credit-rating agency which rates economies, with grades of A, B and C, so that global investors know who is pursuing sound economics and who is not, and if your economy gets a lower rating it means you will have to pay higher interest rates for international loans. "Moody's has sent a team over to re-rate the Indian economy," she said.

"Have you heard anything about what they decided?"

No, I hadn't, I replied.

"You might want to check," she said, and swam away.

I did check. It turned out that the Moody's team had moved around New Delhi almost as quietly and secretly as India's nuclear scientists had prepared their bomb. I couldn't find out anything about their decisions, but the night I left India, I was listening to the evening news when the fourth item caught my ear. It said that in reaction to the Indian government's new bloated, directionless budget, and in the wake of the Indian nuclear tests and the U.S. sanctions imposed on India for blowing off some nukes, Moody's had downgraded India's economy from "investment grade," which meant it was safe for global investors, to "speculative grade," which meant it was risky. The Standard & Poor's rating agency also changed its outlook on the Indian economy from "stable" to "negative." This meant that any Indian company trying to borrow money from international markets would have to pay higher interest. And because India has a low savings rate, those foreign funds are crucial for a country that needs \$500 billion in new infrastructure over the next decade in order to be competitive.

So yes, the olive tree had had its day in India. But when it pushes out like that in the system of globalization, there is always a price to pay. You can't escape the system. Sooner or later the Lexus always catches up with you. A year and a half after India's nuclear test, I picked up *The Wall Street Journal* (Oct. 7, 1999) to read the following headline: "India's BJP Is Shifting Priority to the Economy." The story noted that the BJP came to power some two years earlier "calling for India to assert its nuclear capability—a pledge it fulfilled two months later with a series of weapons tests that sparked global sanctions and stalled investment." Upon its reelection, though, Prime Minister Atal Bihari Vajpayee wasn't even waiting for the votes to be counted before signaling his new priority: economic reform. "The priority is to build a national consensus on the acceptance of global capital, market norms and whatever goes with it. You have to go out and compete for investments," Vajpayee told the *Indian Express* newspaper.

An example of the Lexus and olive tree forces in balance was the Gulf Air flight I took from Bahrain to London, on which the television monitor on my Business Class seat included a channel which, using a global positioning satellite (GPS) linked into the airplane's antenna, showed passengers exactly where the plane was in relation to the Muslim holy city of Mecca at all times. The screen displayed a diagram of the aircraft with a white dot that moved around the diagram as the plane changed directions.

This enabled Muslim passengers, who are enjoined to pray five times a day facing toward Mecca, to always know which way to face inside the plane when they unrolled their prayer rugs. During the flight, I saw several passengers near me wedge into the galley to perform their prayer rituals, and thanks to the GPS system, they knew just which way to aim.

The Lexus ignoring the olive tree in the era of globalization was a computer part that a friend of mine sent me. On the hack was written: "This part was made in Malaysia, Singapore, the Philippines, China, Mexico, Germany, the U.S., Thailand, Canada and Japan. It was made in so many different places that we cannot specify a country of origin."

The Lexus trumping the olive tree in the era of globalization was the small item that appeared in the August 11, 1997, edition of *Sports Illustrated*. It said: "The 38-year-old Welsh soccer club Llansantffraid has changed its name to "Total Network Solutions" in exchange for \$400,000 from a cellular phone company."

The Lexus and olive tree working together in the era of globalization was on display in a rather unusual *Washington Times* story of September 21, 1997, which reported that Russian counterintelligence officers were complaining about having to pay twice as much to recruit a CIA spy as a double agent than the other way around. An official of Russia's Federal Security Service (the successor to the KGB), speaking on condition of anonymity, told the Itar-Tass news agency that a Russian spy could be bought for a mere \$1 million, while CIA operatives held out for \$2 million to work for the other side.

At roughly the same time that this report appeared, Israel's *Yediot Aharonot* newspaper published what seemed to me to be the first-ever totally free-market intelligence scoop. Yediot editors went to Moscow and bought some Russian spy satellite photographs of new Scud missile bases in Syria. Then *Yediot* hired a private U.S. expert on satellite photos to analyze the pictures. Then *Yediot* published the whole package as a scoop about Syria's new missile threat, without ever having once quoted a government official. Who needs Deep Throat when you have deep pockets'?

Finally, my favorite "Lexus trumps olive tree in the era of globalization" story is about Abu Jihad's son. I was attending the Middle East Economic Summit in Amman, Jordan, in 1995, and was having lunch by myself on the balcony of the Amman Marriott. Out of the blue, a young Arab man approached my table and asked, "Are you Tom Friedman?" I said yes.

"Mr. Friedman," the young man continued politely, "you knew my father."

"Who was your father?" I asked.

"My father was Abu Jihad."

Abu Jihad, whose real name was Khalil al-Wazir, was one of the Palestinians who, with Yasser Arafat, founded el-Fatah and later took over the Palestine Liberation Organization. Abu Jihad, meaning "father of struggle," was his nom de guerre, and he was the overall commander of Palestinian military operations in Lebanon and the West Bank in the days when I was the *New York Times* correspondent in Beirut. I got to know him in Beirut. Palestinians considered him a military hero; Israelis considered him one of the most dangerous Palestinian terrorists. An Israeli hit team assassinated Abu Jihad in his living room in Tunis on April 16, 1988, pumping a hundred bullets into his body.

"Yes, I knew your father very well—I once visited your home in Damascus," I told the young man. "What do you do?"

He handed me his business card. It read: "Jihad al-Wazir, Managing Director, World Trade Center, Gaza, Palestine."

I read that card and thought to myself, "That's amazing. From Jesse James to Michael Milton in one generation."

The challenge in this era of globalization—for countries and individuals—is to find a healthy balance between preserving a sense of identity, home and community and doing what it takes to survive within the globalization system. Any society that wants to thrive economically today must constantly be trying to build a better Lexus and driving it out into the world. But no one should have any illusions that merely participating in this global economy will make a society healthy. If that participation comes at the price of a country's identity, if individuals feel their olive tree roots crushed, or washed out, by this global system, those olive tree roots will rebel. They will rise up and strangle the process. Therefore the survival of globalization as a system will depend, in part, on how well all of us strike this balance. A country without healthy olive trees will never feel rooted or secure enough to open up fully to the world and reach out into it. But a country that is only olive trees, that is only roots, and has no Lexus, will never go, or grow, very far.

Keeping the two in balance is a constant struggle.

Maybe that's why of the many stories you will read in this book my favorite comes from my old college friend Victor Friedman, who teaches business management at the Ruppin Institute in Israel. I telephoned him one day to say hello and he told me he was glad that I called because he no longer had my phone numbers. When I asked why, he explained that he no longer had his handheld computer, in which he kept everything—his friends' addresses, E-mail addresses, phone numbers and his schedule for the next two years. He then told me what happened to it.

"We had a [desktop] computer at home that broke down. I took it to be repaired at a computer shop in Hadera [a town in central Israel]. A couple weeks later the shop called and said my PC was repaired. So I tossed my palm computer into my leather briefcase and drove over to Hadera to pick up my repaired PC. I left the shop carrying the big PC computer and my briefcase, which had my palm computer inside. When I got to the car, I put my briefcase down on the sidewalk, opened the trunk of my car and very carefully placed my repaired PC in the trunk to make sure that it was secure. Then I got in the car and drove off, leaving my briefcase on the sidewalk. Well, as soon as I got to my office and looked for my briefcase I realized what had happened—and what was going to happen next—and I immediately called the Hadera police to tell them '*Don't blow up my briefcase.*' [It is standard Israeli police practice to blow apart any package, briefcase or suspicious object left on a sidewalk, because this is how many Palestinian bombs against Israeli civilians have been set off. Israelis are so well trained to protect against this that if you leave a package for a minute, the police will already have been called.] I knew no one would steal the briefcase. In Israel, a thief wouldn't touch such an object left on the sidewalk. But I was too late. The police dispatcher told me that the bomb squad was already on the scene and had 'dealt with it.' When I got to the police station they handed me back my beautiful leather briefcase with a nice neat bullet hole right through the middle. The only thing it hurt was my handheld computer. My Genius OP9300 took a direct hit. My whole life was in that thing and I had never made a backup. I told the police I felt terrible for causing such a problem, and they said, 'Don't feel bad, it happens to everyone.' For weeks I walked around campus with my briefcase with the bullet hole in it to remind myself to stop and think more often. Most of my management students are in the Israeli Army and as soon as they saw the briefcase and that bullet hole they would immediately crack up laughing, because they knew just what had happened."

After Victor finished telling me this story, he said, "By the way, send me your E-mail address. I need to start a new address book."

Excerpt from Chapter 12, "The Golden Arches Theory of Conflict Prevention"

Every once in a while when I am traveling abroad, I need to indulge in a burger and a bag of McDonald's french fries. For all I know, I have eaten McDonald's burgers and fries in more countries in the world than anyone, and I can testify that *they all really do taste the same*. But as I Quarter-Pounded my way around the world in recent years, I began to notice something intriguing. I don't know when the insight struck me. It was a bolt out of the blue that must have hit somewhere between the McDonald's in Tiananmen Square in Beijing, the McDonald's in Tahrir Square in Cairo and the McDonald's off Zion Square in Jerusalem. And it was this:

No two countries that both had McDonald's had fought a war against each other since each got its McDonald's.

I'm not kidding. It was uncanny. Look at the Middle East: Israel had a kosher McDonald's, Saudi Arabia had McDonald's, which closed five times a day for Muslim prayer, Egypt had McDonald's and both Lebanon and Jordan had become McDonald's countries. None of them have had a war since the Golden Arches went in. Where is the big threat of war in the Middle East today? Israel-Syria, Israel-Iran and Israel-Iraq. Which three Middle East countries don't have McDonald's? Syria, Iran and Iraq.

I was intrigued enough by my own thesis to call McDonald's headquarters in Oak Brook, Illinois, and report it to them. They were intrigued enough by it to invite me to test it out on some of their international executives at Hamburger University, McDonald's in-house research and training facility. The McDonald's folks ran my model past all their international experts and confirmed that they, too, couldn't find an exception. I feared the exception would be the Falklands war, but Argentina didn't get its first McDonald's until 1986, four years after that war with Great Britain. (Civil wars and border skirmishes don't count: McDonald's in Moscow, El Salvador and Nicaragua served burgers to both sides in their respective civil wars.)

Armed with this data, I offered up “The Golden Arches Theory of Conflict Prevention,” which stipulated that when a country reached the level of economic development where it had a middle class big enough to support a McDonald’s network, it became a McDonald’s country. And people in McDonald’s countries didn’t like to fight wars anymore, they preferred to wait in line for burgers.

Others have made similar observations during previous long periods of peace and commerce—using somewhat more conventional metaphors. The French philosopher Montesquieu wrote in the eighteenth century that international trade had created an international “Grand Republic,” which was uniting all merchants and trading nations across boundaries, which would surely lock in a more peaceful world. In *The Spirit of the Laws* he wrote that “two nations who traffic with each other become reciprocally dependent; for if one has an interest in buying, the other has an interest in selling; and thus their union is founded on their mutual necessities.” And in his chapter entitled “How Commerce Broke Through the Barbarism of Europe,” Montesquieu argued for his own Big Mac thesis: “Happy it is for men that they are in a situation in which, though their passions prompt them to be wicked, it is, nevertheless, to their interest to be humane and virtuous.”

In the pre-World War I era of globalization, the British writer Norman Angell observed in his 1910 book, *The Great Illusion*, that the major Western industrial powers, America, Britain, Germany and France, were losing their taste for war-making: “How can modern life, with its overpowering proportion of industrial activities and its infinitesimal proportion of military, keep alive the instincts associated with war as against those developed by peace?” With all the free trade and commercial links tying together major European powers in his day, Angell argued that it would be insane for them to go to war, because it would destroy both the winner and the loser.

Montesquieu and Angell were actually right. Economic integration was making the cost of war much higher for both victor and vanquished, and any nation that chose to ignore that fact would be devastated. But their hope that this truth would somehow end geopolitics was wrong. Montesquieu and Angell, one might say, forgot their Thucydides. Thucydides wrote in his history of the Peloponnesian War that nations are moved to go to war for one of three reasons—“honor, fear and interest”—and globalization, while it raises the costs of going to war for reasons of honor, fear or interest, does not and cannot make any of these instincts obsolete—not as long as the world is made of men not machines, and not as long as olive trees still matter. The struggle for power, the pursuit of material and strategic interests and the ever-present emotional tug of one’s own olive tree continue even in a world of microchips, satellite phones and the Internet. This book isn’t called *The Lexus and the Olive Tree* for nothing. Despite globalization, people are still attached to their culture, their language and a place called home. And they will sing for home, cry for home, fight for home and die for home. Which is why globalization does not, and will not, end geopolitics. Let me repeat that for all the realists who read this book: *Globalization does not end geopolitics.*

But it does affect it. The simple point I was trying to make—using McDonald’s as a metaphor—is that today’s version of globalization significantly raises the costs of countries using war as a means to pursue honor, react to fears or advance their interests. What is new today, compared to when Montesquieu and even Angell were writing, is a *difference in degree*. Today’s version of globalization—with its intensifying economic integration, digital integration, its ever-widening connectivity of individuals and nations, its spreading of capitalist values and networks to the remotest corners of the world and its growing dependence on the Golden Straitjacket and the Electronic Herd—makes for a much stronger web of constraints on the foreign policy behavior of those nations which are plugged into the system. It both increases the incentives for not making war and it increases the costs of going to war in more ways than in any previous era in modern history.

But it can’t guarantee that there will be no more wars. There will always be leaders and nations who, for good reasons and bad reasons, will resort to war, and some nations, such as North Korea, Iraq or Iran, will choose to live outside the constraints of the system. Still, the bottom line is this: If in the previous era of globalization nations in the system thought twice before trying to solve problems through warfare, in this era of globalization they will think about it three times.

Of course, no sooner did the first edition of this book come out, in April 1999, than nineteen McDonald’s-laden NATO countries undertook air strikes against Yugoslavia, which also had McDonald’s. Immediately, all sorts of commentators and reviewers began writing to say that this proved my McDonald’s theory all wrong, and, by implication, the notion that globalization would affect geopolitics. I was both amazed and amused by how much the Golden Arches Theory had gotten around and how intensely certain people wanted to prove it wrong. They were mostly realists and out-of-work Cold

Warriors who insisted that politics, and the never ending struggle between nation-states, were the immutable defining feature of international affairs, and they were both professionally and psychologically threatened by the idea that globalization and economic integration might actually influence geopolitics in some very new and fundamental ways. Many of these critics were particularly obsessed with the Balkans precisely because this old-world saga, in which politics, passion and olive trees always takes precedence over economics and the Lexus, is what they knew. They were so busy elevating the Balkans into a world historical issue, into the paradigm of what world politics is actually about, that they failed to notice just what an exception it was, and how, rather than spreading around the world, the Balkans was isolated by the world. They were so busy debating whether we were in 1917, 1929 or 1939 that they couldn't see that what was happening in 2000 might actually be something fundamentally new—something that doesn't end geopolitics but influences and reshapes it in important ways. These critics, I find, are so busy dwelling on what happened yesterday, and telling you what will happen someday, that they have nothing to say about what is happening today. They are experts at extrapolating the future from the past, while skipping over the present. It's not surprising this group would be threatened by the McDonald's argument, because, if it were even half true, they would have to adapt their worldviews or, even worse, learn to look at the world differently and to bring economics, environment, markets, technology, the Internet and the whole globalization system more into their analyses of geopolitics.

My first reaction to these critics was to defensively point out that NATO isn't a country, that the Kosovo war wasn't even a real war and to the extent that it was a real war it was an intervention by NATO into a civil war between Kosovo Serbs and Albanians. And I pointed out that when I posited my original McDonald's theory I had qualified it in several important ways: the McDonald's theory didn't apply to civil wars, because, I explained, globalization is going to sharpen civil wars within countries between localizers and globalizers—between those who eat the Big Mac and those who fear the Big Mac will eat them. Moreover, the theory was offered with a limited shelf life, because, I said, sooner or later virtually every country would have McDonald's, and sooner or later two of them would go to war.

But I quickly realized that no one was interested in my caveats, the fine print or the idea that McDonald's was simply a metaphor for a larger point about the impact of globalization on geopolitics. They just wanted to drive a stake through this Golden Arches Theory. So the more I thought about the criticism, the more I told people, "You know what, forget all the caveats and the fine print. Let's assume Kosovo is a real test. Let's see how the war ends." And when you look at how the war ended you can see just how much the basic logic of the Golden Arches Theory still applies.

Here's why: As the Pentagon will tell you, airpower alone brought the 1999 Kosovo war to a close in seventy-eight days for one reason—not because NATO made life impossible for the Serb troops in Kosovo. Indeed, the Serbian army ended up driving most of its armor out of Kosovo unscathed. No, this war ended in seventy-eight days, using airpower alone, because NATO made life miserable for the Serb civilians in Belgrade. Belgrade was a modern European city integrated with Western Europe, with a population that wanted to be part of today's main global trends, from the Internet to economic development—which the presence of McDonald's symbolized.

Once NATO turned out the lights in Belgrade, and shut down the power grids and the economy, Belgrade's citizens almost immediately demanded that President Slobodan Milosevic bring an end to the war, as did the residents of Yugoslavia's other major cities. Because the air war forced a choice on them: Do you want to be part of Europe and the broad economic trends and opportunities in the world today or do you want to keep Kosovo and become an isolated, backward tribal enclave: It's McDonald's or Kosovo—you can't have both. And the Serbian people chose McDonald's. Not only did NATO soldiers not want to die for Kosovo—neither did the Serbs of Belgrade. In the end, they wanted to be part of the world, more than they wanted to be part of Kosovo. They wanted McDonald's re-opened, much more than they wanted Kosovo reoccupied. They wanted to stand in line for burgers, much more than they wanted to stand in line for Kosovo. Airpower alone couldn't work in Vietnam because a people who were already in the Stone Age couldn't be bombed back into it. But it could work in Belgrade, because people who were integrated into Europe and the world could be bombed out of it. And when presented by NATO with the choice—your Lexus or your olive tree?—they opted for the Lexus.

So, yes, there is now one exception to the Golden Arches Theory—an exception that, in the end, only proves how powerful is the general rule. Kosovo proves just how much pressure even the most olive-tree-hugging nationalist regimes can come under when the costs of their adventures, and wars of choice, are brought home to their people in the age of globalization. Because in a world where we all increasingly know how each other lives, where governments increasingly have to promise and deliver the

same things, governments can ask their people to sacrifice only so much. When governments do things that make economic integration and a better lifestyle—symbolized by the presence of McDonald's—less possible, people in developed countries simply will not tolerate it for as long as they did in the past. Which is why countries in the system will now think three times before going to war and those that don't will pay three times the price. So let me slightly amend the Golden Arches Theory in light of Kosovo and what are sure to be future Kosovos. I would restate it as follows: People in McDonald's countries don't like to fight wars anymore, they prefer to wait in line for burgers—and *those leaders or countries which ignore that fact will pay a much, much higher price than they think.*

On July 8, 1999, USA "Today ran a story from Belgrade that caught my eye. It was about the economic devastation visited on Yugoslavia as a result of the war. The story contained the following two paragraphs, which, had I written them myself, people would have insisted I made them up:

“Zoran Vukovic, 56, a bus driver in the city of Niš, earns the equivalent of \$62 a month, less than half his salary before the war. The [Serb] government laid off almost half of the roughly 200 drivers last month. The rest had their salaries slashed. With the state controlling the price of food, Vukovic and his eight dependents can survive. But most extras are simply out of the question.

“ ‘McDonald's is now only a dream,’ says Vukovic, who used to take his three grandchildren to the Belgrade outlet. ‘One day, maybe, everything will be O.K. I just don't think it will be in my lifetime.’ ”

A Road Map for Natural Capitalism

by Amory B. Lovins, L. Hunter Lovins, and Paul Hawken



Harvard Business Review

Reprint 99309


Harvard Business Review

MAY–JUNE 1999

Reprint Number

KATHLEEN M. EISENHARDT AND SHONA L. BROWN	PATCHING: RESTITCHING BUSINESS PORTFOLIOS IN DYNAMIC MARKETS	99303
ROBERT SIMONS	HOW RISKY IS YOUR COMPANY?	99311
JAY W. LORSCH AND RAKESH KHURANA	CHANGING LEADERS: THE BOARD'S ROLE IN CEO SUCCESSION A ROUNDTABLE WITH PHILIP CALDWELL, GEORGE D. KENNEDY, G. G. MICHELSON, HENRY WENDT, AND ALFRED M. ZEIEN	99308
JON R. KATZENBACH AND JASON A. SANTAMARIA	FIRING UP THE FRONT LINE	99307
ROSABETH MOSS KANTER	FROM SPARE CHANGE TO REAL CHANGE: THE SOCIAL SECTOR AS BETA SITE FOR BUSINESS INNOVATION	99306
JEFFREY PFEFFER AND ROBERT I. SUTTON	THE SMART-TALK TRAP	99310
AMORY B. LOVINS, L. HUNTER LOVINS, AND PAUL HAWKEN	A ROAD MAP FOR NATURAL CAPITALISM	99309
JACQUELINE VISCHER	HBR CASE STUDY WILL THIS OPEN SPACE WORK?	99312
DAVID DUNNE AND CHAKRAVARTHI NARASIMHAN	THINKING ABOUT... THE NEW APPEAL OF PRIVATE LABELS	99302
DANNY ERTEL	IDEAS AT WORK TURNING NEGOTIATION INTO A CORPORATE CAPABILITY	99304
DONALD F. HASTINGS	FIRST PERSON LINCOLN ELECTRIC'S HARSH LESSONS FROM INTERNATIONAL EXPANSION	99305
NICHOLAS G. CARR	BOOKS IN REVIEW BEING VIRTUAL: CHARACTER AND THE NEW ECONOMY	99301

A ROAD MAP FOR NATURAL CAPITALISM



Business strategies built around
the radically more productive use
of natural resources can solve
many environmental
problems at a profit.

BY AMORY B. LOVINS,
L. HUNTER LOVINS,
AND PAUL HAWKEN

ON SEPTEMBER 16, 1991, a small group of scientists was sealed inside Biosphere II, a glittering 3.2-acre glass and metal dome in Oracle, Arizona. Two years later, when the radical attempt to replicate the earth's main ecosystems in miniature ended, the engineered environment was dying. The gaunt researchers had survived only because fresh air had been pumped in. Despite \$200 million worth of elaborate equipment, Biosphere II had failed to generate breathable air, drinkable water, and adequate food for just eight people. Yet Biosphere I, the planet we all

inhabit, effortlessly performs those tasks every day for 6 billion of us.

Disturbingly, Biosphere I is now itself at risk. The earth's ability to sustain life, and therefore economic activity, is threatened by the way we extract, process, transport, and dispose of a vast flow of resources—some 220 billion tons a year, or more than 20 times the average American's body weight every day. With dangerously narrow focus, our industries

.....

Some very simple changes to the way we run our businesses can yield startling benefits for today's shareholders and for future generations.

.....

look only at the exploitable resources of the earth's ecosystems—its oceans, forests, and plains—and not at the larger services that those systems provide for free. Resources and ecosystem services both come from the earth—even from the same biological systems—but they're two different things. Forests, for instance, not only produce the resource of wood fiber but also provide such ecosystem services as water storage, habitat, and regulation of the atmosphere and climate. Yet companies that earn income from harvesting the wood fiber resource often do so in ways that damage the forest's ability to carry out its other vital tasks.

Unfortunately, the cost of destroying ecosystem services becomes apparent only when the services start to break down. In China's Yangtze basin in 1998, for example, deforestation triggered flooding that killed 3,700 people, dislocated 223 million, and inundated 60 million acres of cropland. That \$30 billion disaster forced a logging moratorium and a \$12 billion crash program of reforestation.

The reason companies (and governments) are so prodigal with ecosystem services is that the value

A MacArthur Fellow, Amory B. Lovins is the research director and CFO of Rocky Mountain Institute (RMI).

L. Hunter Lovins is the CEO of RMI, the nonprofit resource policy center they cofounded in 1982 in Snowmass, Colorado (<http://www.rmi.org>).

*Paul Hawken is the founder of the Smith & Hawken retail and catalog company, cofounder of the knowledge-management software company Datafusion, and author of *Growing a Business* (Simon & Schuster, 1983) and *The Ecology of Commerce* (Harper Collins, 1993).*

*Hawken and the Lovinses consult for businesses worldwide and have coauthored the forthcoming *Natural Capitalism* (Little Brown, September 1999).*

of those services doesn't appear on the business balance sheet. But that's a staggering omission. The economy, after all, is embedded in the environment. Recent calculations published in the journal *Nature* conservatively estimate the value of all the earth's ecosystem services to be at least \$33 trillion a year. That's close to the gross world product, and it implies a capitalized book value on the order of half a quadrillion dollars. What's more, for most of these services, there is no known substitute at any price, and we can't live without them.

This article puts forward a new approach not only for protecting the biosphere but also for improving profits and competitiveness. Some very simple changes to the way we run our businesses, built on advanced techniques for making resources more productive, can yield startling benefits both for today's shareholders and for future generations.

This approach is called *natural capitalism* because it's what capitalism might become if its largest category of capital—the “natural capital” of ecosystem services—were properly valued. The journey to natural capitalism involves four major shifts in business practices, all vitally interlinked:

- **Dramatically increase the productivity of natural resources.** Reducing the wasteful and destructive flow of resources from depletion to pollution represents a major business opportunity. Through fundamental changes in both production design and technology, farsighted companies are developing ways to make natural resources—energy, minerals, water, forests—stretch 5, 10, even 100 times further than they do today. These major resource savings often yield higher profits than small resource savings do—or even saving no resources at all would—and not only pay for themselves over time but in many cases reduce initial capital investments.

- **Shift to biologically inspired production models.** Natural capitalism seeks not merely to reduce waste but to eliminate the very concept of waste. In closed-loop production systems, modeled on nature's designs, every output either is returned harmlessly to the ecosystem as a nutrient, like compost, or becomes an input for manufacturing another product. Such systems can often be designed to eliminate the use of toxic materials, which can hamper nature's ability to reprocess materials.

- **Move to a solutions-based business model.** The business model of traditional manufacturing rests on the sale of goods. In the new model, value is instead delivered as a flow of services—providing illumination, for example, rather than selling lightbulbs. This model entails a new perception of value,



Reducing the wasteful and destructive flow of resources represents a major business opportunity.

a move from the acquisition of goods as a measure of affluence to one where well-being is measured by the continuous satisfaction of changing expectations for quality, utility, and performance. The new relationship aligns the interests of providers and customers in ways that reward them for implementing the first two innovations of natural capitalism—resource productivity and closed-loop manufacturing.

■ **Reinvest in natural capital.** Ultimately, business must restore, sustain, and expand the planet’s ecosystems so that they can produce their vital services and biological resources even more abundantly. Pressures to do so are mounting as human needs expand, the costs engendered by deteriorating ecosystems rise, and the environmental awareness of consumers increases. Fortunately, these pressures all create business value.

Natural capitalism is not motivated by a current scarcity of natural resources. Indeed, although many biological resources, like fish, are becoming scarce, most mined resources, such as copper and oil, seem ever more abundant. Indices of average commodity prices are at 28-year lows, thanks partly to powerful extractive technologies, which are often subsidized and whose damage to natural capital remains unaccounted for. Yet even despite these artificially low prices, using resources manyfold more productively can now be so profitable that pioneering companies—large and small—have already embarked on the journey toward natural capitalism.¹

Still the question arises—if large resource savings are available and profitable, why haven’t they all been captured already? The answer is simple: scores

of the important business opportunities they reveal. But first, let’s map the route toward natural capitalism.

Dramatically Increase the Productivity of Natural Resources

In the first stage of a company’s journey toward natural capitalism, it strives to wring out the waste of energy, water, materials, and other resources throughout its production systems and other operations. There are two main ways companies can do this at a profit. First, they can adopt a fresh approach to design that considers industrial systems as a whole rather than part by part. Second, companies can replace old industrial technologies with new ones, particularly with those based on natural processes and materials.

Implementing Whole-System Design. Inventor Edwin Land once remarked that “people who seem to have had a new idea have often simply stopped having an old idea.” This is particularly true when designing for resource savings. The old idea is one of diminishing returns—the greater the resource saving, the higher the cost. But that old idea is giving way to the new idea that bigger savings can cost less—that saving a large fraction of resources can actually cost less than saving a small fraction of resources. This is the concept of expanding returns, and it governs much of the revolutionary thinking behind whole-system design. Lean manufacturing is an example of whole-system thinking that has helped many companies dramatically reduce such forms of waste as lead times, defect rates, and inventory. Applying whole-system thinking to the productivity of natural resources can achieve even more.

Consider Interface Corporation, a leading maker of materials for commercial interiors. In its new Shanghai carpet factory, a liquid had to be circulated through a standard pumping loop similar to those used in nearly all industries. A top European company designed the system to use pumps requiring a total of 95 horsepower. But before construction began, Interface’s engineer, Jan Schilham, realized that two embarrassingly simple design changes would cut that power requirement to only 7 horsepower—a 92% reduction. His redesigned system cost less to build, involved no new technology, and worked better in all respects.

What two design changes achieved this 12-fold saving in pumping power? First, Schilham chose fatter-than-usual pipes, which create much less friction than thin pipes do and therefore need far

.....
 Saving a large fraction of resources can actually cost less than saving a small fraction of resources. This is the concept of expanding returns.

of common practices in both the private and public sectors systematically reward companies for wasting natural resources and penalize them for boosting resource productivity. For example, most companies expense their consumption of raw materials through the income statement but pass resource-saving investment through the balance sheet. That distortion makes it more tax efficient to waste fuel than to invest in improving fuel efficiency. In short, even though the road seems clear, the compass that companies use to direct their journey is broken. Later we’ll look in more detail at some of the obstacles to resource productivity—and some

less pumping energy. The original designer had chosen thin pipes because, according to the textbook method, the extra cost of fatter ones wouldn't be justified by the pumping energy that they would save. This standard design trade-off optimizes the pipes by themselves but "pessimizes" the larger system. Schilham optimized the *whole* system by counting not only the higher capital cost of the fatter pipes but also the *lower* capital cost of the smaller pumping equipment that would be needed. The pumps, motors, motor controls, and electrical components could all be much smaller because there'd be less friction to overcome. Capital cost would fall far more for the smaller equipment than it would rise for the fatter pipe.

Choosing big pipes and small pumps—rather than small pipes and big pumps—would therefore make the whole system cost less to build, even before counting its future energy savings.

Schilham's second innovation was to reduce the friction even more by making the pipes short and straight rather than long and crooked. He did this by laying out the pipes first, *then* positioning the various tanks, boilers, and other equipment that they connected. Designers normally locate the production equipment in arbitrary positions and then have a pipe fitter connect everything. Awkward placement forces the pipes to make numerous bends that greatly increase friction. The pipe fitters don't mind: they're paid by the hour, they profit from the extra pipes and fittings, and they don't pay for the oversized pumps or inflated electric bills. In addition to reducing those four kinds of costs, Schilham's short, straight pipes were easier to insulate, saving an extra 70 kilowatts of heat loss and repaying the insulation's cost in three months.

This small example has big implications for two reasons. First, pumping is the largest application of motors, and motors use three-quarters of all industrial electricity. Second, the lessons are very widely relevant. Interface's pumping loop shows how simple changes in design mentality can yield huge resource savings and returns on investment. This isn't rocket science; often it's just a rediscovery of good Victorian engineering principles that have been lost because of specialization.

Whole-system thinking can help managers find small changes that lead to big savings that are cheap, free, or even better than free (because they make the whole system cheaper to build). They can do this because often the right investment in one part of the system can produce multiple benefits throughout the system. For example, companies would gain 18 distinct economic benefits—of which

direct energy savings is only one—if they switched from ordinary motors to premium-efficiency motors or from ordinary lighting ballasts (the transformer-like boxes that control fluorescent lamps) to electronic ballasts that automatically dim the lamps to match available daylight. If everyone in America integrated these and other selected technologies into all existing motor and lighting systems in an

.....

**Interface's engineer realized that two
embarrassingly simple design changes would
cut power requirements by 92%.**

.....

optimal way, the nation's \$220-billion-a-year electric bill would be cut in half. The after-tax return on investing in these changes would in most cases exceed 100% per year.

The profits from saving electricity could be increased even further if companies also incorporated the best off-the-shelf improvements into their building structure and their office, heating, cooling, and other equipment. Overall, such changes could cut national electricity consumption by at least 75% and produce returns of around 100% a year on the investments made. More important, because workers would be more comfortable, better able to see, and less fatigued by noise, their productivity and the quality of their output would rise. Eight recent case studies of people working in well-designed, energy-efficient buildings measured labor productivity gains of 6% to 16%. Since a typical office pays about 100 times as much for people as it does for energy, this increased productivity in people is worth about 6 to 16 times as much as eliminating the entire energy bill.

Energy-saving, productivity-enhancing improvements can often be achieved at even lower cost by piggybacking them onto the periodic renovations that all buildings and factories need. A recent proposal for reallocating the normal 20-year renovation budget for a standard 200,000-square-foot glass-clad office tower near Chicago, Illinois, shows the potential of whole-system design. The proposal suggested replacing the aging glazing system with a new kind of window that lets in nearly six times more daylight than the old sun-blocking glass units. The new windows would reduce the flow of heat and noise four times better than traditional windows do. So even though the glass costs slightly more, the overall cost of the renovation would be reduced because the windows would let in cool, glare-free daylight that, when combined with more

efficient lighting and office equipment, would reduce the need for air-conditioning by 75%. Installing a fourfold more efficient, but fourfold smaller, air-conditioning system would cost \$200,000 less than giving the old system its normal 20-year renovation. The \$200,000 saved would, in turn, pay for the extra cost of the new windows and other improvements. This whole-system approach to renovation would not only save 75% of the building's total energy use, it would also greatly improve the

.....

In an experiment at its Swiss headquarters, Dow Europe cut office paper flow by about 30% in six weeks simply by discouraging unneeded information.

.....

building's comfort and marketability. Yet it would cost essentially the same as the normal renovation. There are about 100,000 twenty-year-old glass office towers in the United States that are ripe for such improvement.

Major gains in resource productivity require that the right steps be taken in the right order. Small changes made at the downstream end of a process often create far larger savings further upstream. In almost any industry that uses a pumping system, for example, saving one unit of liquid flow or friction in an exit pipe saves about ten units of fuel, cost, and pollution at the power station.

Of course, the original reduction in flow itself can bring direct benefits, which are often the reason changes are made in the first place. In the 1980s, while California's industry grew 30%, for example, its water use was cut by 30%, largely to avoid increased wastewater fees. But the resulting reduction in pumping energy (and the roughly tenfold larger saving in power-plant fuel and pollution) delivered bonus savings that were at the time largely unanticipated.

To see how downstream cuts in resource consumption can create huge savings upstream, consider how reducing the use of wood fiber disproportionately reduces the pressure to cut down forests. In round numbers, half of all harvested wood fiber is used for such structural products as lumber; the other half is used for paper and cardboard. In both cases, the biggest leverage comes from reducing the amount of the retail product used. If it takes, for example, three pounds of harvested trees to produce one pound of product, then saving one pound of product will save three pounds of trees—plus all the

environmental damage avoided by not having to cut them down in the first place.

The easiest savings come from not using paper that's unwanted or unneeded. In an experiment at its Swiss headquarters, for example, Dow Europe cut office paper flow by about 30% in six weeks simply by discouraging unneeded information. For instance, mailing lists were eliminated and senders of memos got back receipts indicating whether each recipient had wanted the information. Taking

those and other small steps, Dow was also able to increase labor productivity by a similar proportion because people could focus on what they really needed to read. Similarly, Danish hearing-aid maker Oticon saved upwards of 30% of its paper as a by-product of redesigning its business processes to produce better decisions faster. Setting the default on office printers and copiers to double-sided mode re-

duced AT&T's paper costs by about 15%. Recently developed copiers and printers can even strip off old toner and printer ink, permitting each sheet to be reused about ten times.

Further savings can come from using thinner but stronger and more opaque paper, and from designing packaging more thoughtfully. In a 30-month effort at reducing such waste, Johnson & Johnson saved 2,750 tons of packaging, 1,600 tons of paper, \$2.8 million, and at least 330 acres of forest annually. The downstream savings in paper use are multiplied by the savings further upstream, as less need for paper products (or less need for fiber to make each product) translates into less raw paper, less raw paper means less pulp, and less pulp requires fewer trees to be harvested from the forest. Recycling paper and substituting alternative fibers such as wheat straw will save even more.

Comparable savings can be achieved for the wood fiber used in structural products. Pacific Gas and Electric, for example, sponsored an innovative design developed by Davis Energy Group that used engineered wood products to reduce the amount of wood needed in a stud wall for a typical tract house by more than 70%. These walls were stronger, cheaper, more stable, and insulated twice as well. Using them enabled the designers to eliminate heating and cooling equipment in a climate where temperatures range from freezing to 113°F. Eliminating the equipment made the whole house much less expensive both to build and to run while still maintaining high levels of comfort. Taken together, these and many other savings in the paper and construction industries could make our use of wood fiber so much more productive that, in principle,

the entire world's present wood fiber needs could probably be met by an intensive tree farm about the size of Iowa.

Adopting Innovative Technologies. Implementing whole-system design goes hand in hand with introducing alternative, environmentally friendly technologies. Many of these are already available and profitable but not widely known. Some, like the "designer catalysts" that are transforming the chemical industry, are already runaway successes. Others are still making their way to market, delayed by cultural rather than by economic or technical barriers.

The automobile industry is particularly ripe for technological change. After a century of development, motorcar technology is showing signs of age. Only 1% of the energy consumed by today's cars is actually used to move the driver: only 15% to 20% of the power generated by burning gasoline reaches the wheels (the rest is lost in the engine and drivetrain) and 95% of the resulting propulsion moves the car, not the driver. The industry's infrastructure is hugely expensive and inefficient. Its convergent products compete for narrow niches in saturated core markets at commoditylike prices. Auto making is capital intensive, and product cycles are long. It is profitable in good years but subject to large losses in bad years. Like the typewriter industry just before the advent of personal computers, it is vulnerable to displacement by something completely different.

Enter the Hypercar. Since 1993, when Rocky Mountain Institute placed this automotive concept in the public domain, several dozen current and potential auto manufacturers have committed billions of dollars to its development and commercialization. The Hypercar integrates the best existing technologies to reduce the consumption of fuel as much as 85% and the amount of materials used up to 90% by introducing four main innovations.

First, making the vehicle out of advanced polymer composites, chiefly carbon fiber, reduces its weight by two-thirds while maintaining crashworthiness. Second, aerodynamic design and better tires reduce air resistance by as much as 70% and rolling resistance by up to 80%. Together, these innovations save about two-thirds of the fuel. Third, 30% to 50% of the remaining fuel is saved by using a "hybrid-electric" drive. In such a system, the wheels are turned by electric motors whose power is made onboard by a small engine or turbine, or even more efficiently by a fuel cell. The fuel cell generates electricity directly by chemically com-

binning stored hydrogen with oxygen, producing pure hot water as its only by-product. Interactions between the small, clean, efficient power source and the ultralight, low-drag auto body then further reduce the weight, cost, and complexity of both. Fourth, much of the traditional hardware – from transmissions and differentials to gauges and certain parts of the suspension – can be replaced by electronics controlled with highly integrated, customizable, and upgradable software.

These technologies make it feasible to manufacture pollution-free, high-performance cars, sport utilities, pickup trucks, and vans that get 80 to 200 miles per gallon (or its energy equivalent in other fuels). These improvements will not require any compromise in quality or utility. Fuel savings will not come from making the vehicles small, sluggish, unsafe, or unaffordable, nor will they depend on government fuel taxes, mandates, or subsidies. Rather, Hypercars will succeed for the same reason that people buy compact discs instead of phonograph records: the CD is a superior product that redefines market expectations. From the manufacturers' perspective, Hypercars will cut cycle times, capital needs, body part counts, and assembly effort and space by as much as tenfold. Early adopters will have a huge competitive advantage – which is why dozens of corporations, including most automakers, are now racing to bring Hypercar-like products to market.²

In the long term, the Hypercar will transform industries other than automobiles. It will displace about an eighth of the steel market directly and most of the rest eventually, as carbon fiber becomes far cheaper. Hypercars and their cousins could ultimately save as much oil as OPEC now sells. Indeed,



We could use wood fiber so much more productively that, in principle, the entire world's wood fiber needs could probably be met by an intensive tree farm about the size of Iowa.



oil may well become uncompetitive as a fuel long before it becomes scarce and costly. Similar challenges face the coal and electricity industries because the development of the Hypercar is likely to accelerate greatly the commercialization of inexpensive hydrogen fuel cells. These fuel cells will help shift power production from centralized coal-fired and nuclear power stations to networks of decentralized, small-scale generators. In fact, fuel-

cell-powered Hypercars could themselves be part of these networks. They'd be, in effect, 20-kilowatt power plants on wheels. Given that cars are left parked – that is, unused – more than 95% of the time, these Hypercars could be plugged into a grid and could then sell back enough electricity to repay as much as half the predicted cost of leasing them. A national Hypercar fleet could ultimately have five to ten times the generating capacity of the national electric grid.

As radical as it sounds, the Hypercar is not an isolated case. Similar ideas are emerging in such industries as chemicals, semiconductors, general manufacturing, transportation, water and waste-water treatment, agriculture, forestry, energy, real estate, and urban design. For example, the amount of carbon dioxide released for each microchip manufactured can be reduced almost 100-fold through improvements that are now profitable or soon will be.

Some of the most striking developments come from emulating nature's techniques. In her book, *Biomimicry*, Janine Benyus points out that spiders convert digested crickets and flies into silk that's as strong as Kevlar without the need for boiling sulfuric acid and high-temperature extruders. Using no furnaces, abalone can convert seawater into an inner shell twice as tough as our best ceramics. Trees turn sunlight, water, soil, and air into cellulose, a

.....

Only about 1% of all materials mobilized to serve America is actually made into products and still in use six months after sale.

.....

sugar stronger than nylon but one-fourth as dense. They then bind it into wood, a natural composite with a higher bending strength than concrete, aluminum alloy, or steel. We may never become as skillful as spiders, abalone, or trees, but smart designers are already realizing that nature's environmentally benign chemistry offers attractive alternatives to industrial brute force.

Whether through better design or through new technologies, reducing waste represents a vast business opportunity. The U.S. economy is not even 10% as energy efficient as the laws of physics allow. Just the energy thrown off as waste heat by U.S. power stations equals the total energy use of Japan. Materials efficiency is even worse: only about 1% of all the materials mobilized to serve America is actually made into products and still in use six months after sale. In every sector, there are opportunities for reducing the amount of resources that go into

a production process, the steps required to run that process, and the amount of pollution generated and by-products discarded at the end. These all represent avoidable costs and hence profits to be won.

Redesign Production According to Biological Models

In the second stage on the journey to natural capitalism, companies use closed-loop manufacturing to create new products and processes that can totally prevent waste. This plus more efficient production processes could cut companies' long-term materials requirements by more than 90% in most sectors.

The central principle of closed-loop manufacturing, as architect Paul Bierman-Lytle of the engineering firm CH2M Hill puts it, is "waste equals food." Every output of manufacturing should be either composted into natural nutrients or remanufactured into technical nutrients – that is, it should be returned to the ecosystem or recycled for further production. Closed-loop production systems are designed to eliminate any materials that incur disposal costs, especially toxic ones, because the alternative – isolating them to prevent harm to natural systems – tends to be costly and risky. Indeed, meeting EPA and OSHA standards by eliminating harmful

materials often makes a manufacturing process cost less than the hazardous process it replaced. Motorola, for example, formerly used chlorofluorocarbons for cleaning printed circuit boards after soldering. When CFCs were outlawed because they destroy stratospheric ozone, Motorola at first explored such alterna-

tives as orange-peel terpenes. But it turned out to be even cheaper – and to produce a better product – to redesign the whole soldering process so that it needed no cleaning operations or cleaning materials at all.

Closed-loop manufacturing is more than just a theory. The U.S. remanufacturing industry in 1996 reported revenues of \$53 billion – more than consumer-durables manufacturing (appliances; furniture; audio, video, farm, and garden equipment). Xerox, whose bottom line has swelled by \$700 million from remanufacturing, expects to save another \$1 billion just by remanufacturing its new, entirely reusable or recyclable line of "green" photocopiers. What's more, policy makers in some countries are already taking steps to encourage industry to think along these lines. German law, for example, makes many manufacturers responsible for their products forever, and Japan is following suit.

Combining closed-loop manufacturing with resource efficiency is especially powerful. DuPont, for example, gets much of its polyester industrial film back from customers after they use it and recycles it into new film. DuPont also makes its polyester film ever stronger and thinner so it uses less material and costs less to make. Yet because the film performs better, customers are willing to pay more for it. As DuPont chairman Jack Krol noted in 1997, "Our ability to continually improve the inherent properties [of our films] enables this process [of developing more productive materials, at lower cost, and higher profits] to go on indefinitely."


Interface is leading the way to this next frontier of industrial ecology. While its competitors are "down cycling" nylon-and-PVC-based carpet into less valuable carpet backing, Interface has invented a new floor-covering material called Solenium, which can be completely remanufactured into identical new product. This fundamental innovation emerged from a clean-sheet redesign. Executives at Interface didn't ask how they could sell more carpet of the familiar kind; they asked how they could create a dream product that would best meet their customers' needs while protecting and nourishing natural capital.

Solenium lasts four times longer and uses 40% less material than ordinary carpets – an 86% reduction in materials intensity. What's more, Solenium is free of chlorine and other toxic materials, is virtually stainproof, doesn't grow mildew, can easily be cleaned with water, and offers aesthetic advantages over traditional carpets. It's so superior in every respect that Interface doesn't market it as an environmental product – just a better one.

Solenium is only one part of Interface's drive to eliminate every form of waste. Chairman Ray C. Anderson defines waste as "any measurable input that does not produce customer value," and he considers all inputs to be waste until shown otherwise. Between 1994 and 1998, this zero-waste approach led to a systematic treasure hunt that helped to keep resource inputs constant while revenues rose by \$200 million. Indeed, \$67 million of the revenue

increase can be directly attributed to the company's 60% reduction in landfill waste.

Subsequently, president Charlie Eitel expanded the definition of waste to include all fossil fuel inputs, and now many customers are eager to buy products from the company's recently opened solar-



The central principle of closed-loop manufacturing is "waste equals food." Every output of manufacturing should either be composted into natural nutrients and returned to the ecosystem or be remanufactured into new products.

powered carpet factory. Interface's green strategy has not only won plaudits from environmentalists, it has also proved a remarkably successful business strategy. Between 1993 and 1998, revenue has more than doubled, profits have more than tripled, and the number of employees has increased by 73%.

Change the Business Model

In addition to its drive to eliminate waste, Interface has made a fundamental shift in its business model – the third stage on the journey toward natural capital-

ism. The company has realized that clients want to walk on and look at carpets—but not necessarily to own them. Traditionally, broadloom carpets in office buildings are replaced every decade because some portions look worn out. When that happens, companies suffer the disruption of shutting down their offices and removing their furniture. Billions of pounds of carpets are removed each year and sent to landfills, where they will last up to 20,000 years. To escape this unproductive and wasteful cycle, Interface is transforming itself from a company that sells and fits carpets into one that provides floor-covering services.

Under its Evergreen Lease, Interface no longer sells carpets but rather leases a floor-covering service for a monthly fee, accepting responsibility for keeping the carpet fresh and clean. Monthly inspections detect and replace worn carpet tiles. Since at most 20% of an area typically shows at least 80% of the wear, replacing only the worn parts reduces the consumption of carpeting material by about 80%. It also minimizes the disruption that customers experience—worn tiles are seldom found under furniture. Finally, for the customer, leasing carpets can provide a tax advantage by turning a capital expenditure into a tax-deductible expense. The result: the customer gets cheaper and better

selling products. The more products sold, the better—at least for the company, if not always for the customer or the earth. But any model that wastes natural resources also wastes money. Ultimately, that model will be unable to compete with a service model that emphasizes solving problems and building long-term relationships with customers rather than making and selling products. The shift to what James Womack of the Lean Enterprise Institute calls a “solutions economy” will almost always improve customer value *and* providers’ bottom lines because it aligns both parties’ interests, offering rewards for doing more and better with less.

Interface is not alone. Elevator giant Schindler, for example, prefers leasing vertical transportation services to selling elevators because leasing lets it capture the savings from its elevators’ lower energy and maintenance costs. Dow Chemical and Safety-Kleen prefer leasing dissolving services to selling solvents because they can reuse the same solvent scores of times, reducing costs. United Technologies’ Carrier division, the world’s largest manufacturer of air conditioners, is shifting its mission from selling air conditioners to leasing comfort. Making its air conditioners more durable and efficient may compromise future equipment sales, but it provides what customers want and will pay for—

.....

Elevator giant Schindler prefers leasing vertical transportation services to selling elevators because leasing lets it capture the savings from its elevators’ lower energy and maintenance costs.

.....

services that cost the supplier far less to produce. Indeed, the energy saved from not producing a whole new carpet is in itself enough to produce all the carpeting that the new business model requires. Taken together, the 5-fold savings in carpeting material that Interface achieves through the Evergreen Lease and the 7-fold materials savings achieved through the use of Solenium deliver a stunning 35-fold reduction in the flow of materials needed to sustain a superior floor-covering service. Remanufacturing, and even making carpet initially from renewable materials, can then reduce the extraction of virgin resources essentially to the company’s goal of zero.

Interface’s shift to a service-leasing business reflects a fundamental change from the basic model of most manufacturing companies, which still look on their businesses as machines for producing and

better comfort at lower cost. But Carrier is going even further. It’s starting to team up with other companies to make buildings more efficient so that they need less air-conditioning, or even none at all, to yield the same level of comfort. Carrier will get paid to provide the agreed-upon level of comfort, however that’s delivered. Higher profits will come from providing better solutions rather than from selling more equipment. Since comfort with little or no air-conditioning (via better building design) works better and costs less than comfort with copious air-conditioning, Carrier is smart to capture this opportunity itself before its competitors do. As they say at 3M: “We’d rather eat our *own* lunch, thank you.”

The shift to a service business model promises benefits not just to participating businesses but to the entire economy as well. Womack points out that by helping customers reduce their need for capital goods such as carpets or elevators, and by rewarding suppliers for extending and maximizing asset values rather than for churning them, adoption of the service model will reduce the volatility in the turnover of capital goods that lies at the heart of the business cycle. That would significantly reduce the overall volatility of the world’s economy. At present, the producers of capital goods face feast

or famine because the buying decisions of households and corporations are extremely sensitive to fluctuating income. But in a continuous-flow-of-services economy, those swings would be greatly reduced, bringing a welcome stability to businesses. Excess capacity—another form of waste and source of risk—need no longer be retained for meeting peak demand. The result of adopting the new model would be an economy in which we grow and get richer by using less and become stronger by being leaner and more stable.

Reinvest in Natural Capital

The foundation of textbook capitalism is the prudent reinvestment of earnings in productive capital. Natural capitalists who have dramatically raised their resource productivity, closed their loops, and shifted to a solutions-based business model have one key task remaining. They must reinvest in restoring, sustaining, and expanding the most important form of capital—their own natural habitat and biological resource base.

This was not always so important. Until recently, business could ignore damage to the ecosystem because it didn't affect production and didn't increase costs. But that situation is changing. In 1998 alone, violent weather displaced 300 million people and caused upwards of \$90 billion worth of damage, representing more weather-related destruction than was reported through the entire decade of the 1980s. The increase in damage is strongly linked to deforestation and climate change, factors that accelerate the frequency and severity of natural disasters and are the consequences of inefficient industrialization. If the flow of services from industrial systems is to be sustained or increased in the future for a growing population, the vital flow of services from living systems will have to be maintained or increased as well. Without reinvestment in natural capital, shortages of ecosystem services are likely to become the limiting factor to prosperity in the next century. When a manufacturer realizes that a supplier of key components is overextended and running behind on deliveries, it takes immediate action lest its own production lines come to a halt. The ecosystem is a supplier of key components for the life of the planet, and it is now falling behind on its orders.

Failure to protect and reinvest in natural capital can also hit a company's revenues indirectly. Many companies are discovering that public perceptions of environmental responsibility, or its lack thereof, affect sales. MacMillan Bloedel, targeted by environmental activists as an emblematic clear-cutter

and chlorine user, lost 5% of its sales almost overnight when dropped as a U.K. supplier by Scott Paper and Kimberly-Clark. Numerous case studies show that companies leading the way in implementing changes that help protect the environment tend to gain disproportionate advantage, while companies perceived as irresponsible lose their franchise, their legitimacy, and their shirts. Even businesses that claim to be committed to the concept of sustainable development but whose strategy is seen as mistaken, like Monsanto, are encountering stiffening public resistance to their products. Not surprisingly, University of Oregon business professor Michael Russo, along with many other analysts, has found that a strong environmental rating is "a consistent predictor of profitability."

The pioneering corporations that have made re-investments in natural capital are starting to see some interesting paybacks. The independent power producer AES, for example, has long pursued a policy of planting trees to offset the carbon emissions of its power plants. That ethical stance, once thought quixotic, now looks like a smart investment because a dozen brokers are now starting to create markets in carbon reduction. Similarly, certification by the Forest Stewardship Council of certain sustainably grown and harvested products has given Collins Pine the extra profit margins that enabled its U.S. manufacturing operations to survive brutal competition. Taking an even longer view, Swiss Re and other European reinsurers are seeking to cut their storm-damage losses by pressing for international public policy to protect the climate and by investing in climate-safe technologies that also promise good profits. Yet most companies still do not realize that a vibrant ecological web underpins their survival and their business success. Enriching natural capital is not just a public good—it is vital to every company's longevity.

It turns out that changing industrial processes so that they actually replenish and magnify the stock of natural capital can prove especially profitable because nature does the production; people need just step back and let life flourish. Industries that directly harvest living resources, such as forestry, farming, and fishing, offer the most suggestive examples. Here are three:

- Allan Savory of the Center for Holistic Management in Albuquerque, New Mexico, has redesigned cattle ranching to raise the carrying capacity of rangelands, which have often been degraded not by overgrazing but by undergrazing and grazing the wrong way. Savory's solution is to keep the cattle moving from place to place, grazing intensively but briefly at each site, so that they mimic the dense

but constantly moving herds of native grazing animals that coevolved with grasslands. Thousands of ranchers are estimated to be applying this approach, improving both their range and their profits. This “management-intensive rotational grazing” method, long standard in New Zealand, yields such clearly superior returns that over 15% of Wisconsin’s dairy farms have adopted it in the past few years.

- The California Rice Industry Association has discovered that letting nature’s diversity flourish can be more profitable than forcing it to produce a single product. By flooding 150,000 to 200,000 acres of Sacramento valley rice fields – about 30% of California’s rice-growing area – after harvest, farmers are able to create seasonal wetlands that support millions of wildfowl, replenish groundwater, improve fertility, and yield other valuable benefits. In addition, the farmers bale and sell the rice straw, whose high silica content – formerly an air-pollution hazard when the straw was burned – adds insect resistance and hence value as a construction material when it’s resold instead.

- John Todd of Living Technologies in Burlington, Vermont, has used biological Living Machines – linked tanks of bacteria, algae, plants, and other organisms – to turn sewage into clean water. That not only yields cleaner water at a reduced cost, with no toxicity or odor, but it also produces commercially valuable flowers and makes the plant compatible with its residential neighborhood. A similar plant

duction processes more to biological ones. There is evidence that many business leaders are starting to think this way. The consulting firm Arthur D. Little surveyed a group of North American and European business leaders and found that 83% of them already believe that they can derive “real business value [from implementing a] sustainable-development approach to strategy and operations.”

A Broken Compass?

If the road ahead is this clear, why are so many companies straying or falling by the wayside? We believe the reason is that the instruments companies use to set their targets, measure their performance, and hand out rewards are faulty. In other words, the markets are full of distortions and perverse incentives. Of the more than 60 specific forms of misdirection that we have identified,³ the most obvious involve the ways companies allocate capital and the way governments set policy and impose taxes. Merely correcting these defective practices would uncover huge opportunities for profit.

Consider how companies make purchasing decisions. Decisions to buy small items are typically based on their initial cost rather than their full life-cycle cost, a practice that can add up to major waste. Distribution transformers that supply electricity to buildings and factories, for example, are a minor item at just \$320 apiece, and most companies try to

save a quick buck by buying the lowest-price models. Yet nearly all the nation’s electricity must flow through transformers, and using the cheaper but less efficient models wastes \$1 billion a year. Such examples are legion. Equipping standard new office-lighting circuits with fatter wire that reduces electrical resistance could gener-

ate after-tax returns of 193% a year. Instead, wire as thin as the National Electrical Code permits is usually selected because it costs less up-front. But the code is meant only to prevent fires from overheated wiring, not to save money. Ironically, an electrician who chooses fatter wire – thereby reducing long-term electricity bills – doesn’t get the job. After paying for the extra copper, he’s no longer the low bidder.

Some companies do consider more than just the initial price in their purchasing decisions but still don’t go far enough. Most of them use a crude payback estimate rather than more accurate metrics like discounted cash flow. A few years ago, the median simple payback these companies were demanding from energy efficiency was 1.9 years. That’s equivalent to requiring an after-tax return of around 71% per year – about six times the marginal cost of capital.

.....

Many executives think they already “did” efficiency in the 1970s, but with today’s far better technologies, it’s profitable to start over again.

.....

at the Ethel M Chocolates factory in Las Vegas, Nevada, not only handles difficult industrial wastes effectively but is showcased in its public tours.

Although such practices are still evolving, the broad lessons they teach are clear. In almost all climates, soils, and societies, working with nature is more productive than working against it. Reinvesting in nature allows farmers, fishermen, and forest managers to match or exceed the high yields and profits sustained by traditional input-intensive, chemically driven practices. Although much of mainstream business is still headed the other way, the profitability of sustainable, nature-emulating practices is already being proven. In the future, many industries that don’t now consider themselves dependent on a biological resource base will become more so as they shift their raw materials and pro-

Most companies also miss major opportunities by treating their facilities costs as an overhead to be minimized, typically by laying off engineers, rather than as profit center to be optimized—by using those engineers to save resources. Deficient measurement and accounting practices also prevent companies from allocating costs—and waste—with any accuracy. For example, only a few semiconductor plants worldwide regularly and accurately measure how much energy they're using to produce a unit of chilled water or clean air for their clean-room production facilities. That makes it hard for them to improve efficiency. In fact, in an effort to save time, semiconductor makers frequently build new plants as exact copies of previous ones—a design method nicknamed “infectious repetitis.”

Many executives pay too little attention to saving resources because they are often a small percentage of total costs (energy costs run to about 2% in most industries). But those resource savings drop straight to the bottom line and so represent a far greater percentage of profits. Many executives also think they already “did” efficiency in the 1970s, when the oil shock forced them to rethink old habits. They're forgetting that with today's far better technologies, it's profitable to start all over again. Malden Mills, the Massachusetts maker of such products as Polartec, was already using “efficient” metal-halide lamps in the mid-1990s. But a recent warehouse retrofit reduced the energy used for lighting by another 93%, improved visibility, and paid for itself in 18 months.

The way people are rewarded often creates perverse incentives. Architects and engineers, for example, are traditionally compensated for what they spend, not for what they save. Even the striking economics of the retrofit design for the Chicago office tower described earlier wasn't incentive enough actually to implement it. The property was controlled by a leasing agent who earned a commission every time she leased space, so she didn't want to wait the few extra months needed to refit the building. Her decision to reject the efficiency-quadrupling renovation proved costly for both her and her client. The building was so uncomfortable and expensive to occupy that it didn't lease, so ultimately the owner had to unload it at a firesale price. Moreover, the new owner will for the next 20 years be deprived of the opportunity to save capital cost.

If corporate practices obscure the benefits of natural capitalism, government policy positively undermines it. In nearly every country on the planet, tax laws penalize what we want more of—jobs and

income—while subsidizing what we want less of—resource depletion and pollution. In every state but Oregon, regulated utilities are rewarded for selling more energy, water, and other resources, and penalized for selling less, even if increased production would cost more than improved customer efficiency. In most of America's arid western states, use-it-or-lose-it water laws encourage inefficient water con-

.....

In nearly every country on the planet, tax laws penalize jobs and income while subsidizing resource depletion and pollution.

.....

sumption. Additionally, in many towns, inefficient use of land is enforced through outdated regulations, such as guidelines for ultrawide suburban streets recommended by 1950s civil-defense planners to accommodate the heavy equipment needed to clear up rubble after a nuclear attack.

The costs of these perverse incentives are staggering: \$300 billion in annual energy wasted in the United States, and \$1 trillion already misallocated to unnecessary air-conditioning equipment and the power supplies to run it (about 40% of the nation's peak electric load). Across the entire economy, unneeded expenditures to subsidize, encourage, and try to remedy inefficiency and damage that should not have occurred in the first place probably account for most, if not all, of the GDP growth of the past two decades. Indeed, according to former World Bank economist Herman Daly and his colleague John Cobb (along with many other analysts), Americans are hardly better off than they were in 1980. But if the U.S. government and private industry could redirect the dollars currently earmarked for remedial costs toward reinvestment in natural and human capital, they could bring about a genuine improvement in the nation's welfare. Companies, too, are finding that wasting resources also means wasting money and people. These intertwined forms of waste have equally intertwined solutions. Firing the unproductive tons, gallons, and kilowatt-hours often makes it possible to keep the people, who will have more and better work to do.

Recognizing the Scarcity Shift

In the end, the real trouble with our economic compass is that it points in exactly the wrong direction. Most businesses are behaving as if people were still scarce and nature still abundant—the conditions that helped to fuel the first Industrial

Revolution. At that time, people were relatively scarce compared with the present-day population. The rapid mechanization of the textile industries caused explosive economic growth that created labor shortages in the factory and the field. The Industrial Revolution, responding to those shortages and mechanizing one industry after another, made people a hundred times more productive than they had ever been.

The logic of economizing on the scarcest resource, because it limits progress, remains correct. But the pattern of scarcity is shifting: now people aren't scarce but nature is. This shows up first in industries that depend directly on ecological health. Here, production is increasingly constrained by fish rather than by boats and nets, by forests rather than by chain saws, by fertile topsoil rather than by plows. Moreover, unlike the traditional factors of industrial production—capital and labor—the biological limiting factors cannot be substituted for one another. In the industrial system, we can easily exchange machinery for labor. But no technology or amount of money can substitute for a stable cli-

mate and a productive biosphere. Even proper pricing can't replace the priceless.

Natural capitalism addresses those problems by reintegrating ecological with economic goals. Because it is both necessary and profitable, it will subsume traditional industrialism within a new economy and a new paradigm of production, just as industrialism previously subsumed agrarianism. The companies that first make the changes we have described will have a competitive edge. Those that don't make that effort won't be a problem because ultimately they won't be around. In making that choice, as Henry Ford said, "Whether you believe you can, or whether you believe you can't, you're absolutely right."

1. Our book, *Natural Capitalism*, provides hundreds of examples of how companies of almost every type and size, often through modest shifts in business logic and practice, have dramatically improved their bottom lines.

2. Nonproprietary details are posted at <http://www.hypercar.com>.

3. Summarized in the report "Climate: Making Sense and Making Money" at <http://www.rmi.org/catalog/climate.htm>.

Reprint 99309

To place an order, call 1-800-988-0086.

The Tragedy of the Commons

The population problem has no technical solution;
it requires a fundamental extension in morality.

Garrett Hardin

Science, vol. 162, no. 3859, pp. 1243–1248, 13 December 1968

The author is professor of biology, University of California, Santa Barbara. This article is based on a presidential address presented before the meeting of the Pacific Division of the American Association for the Advancement of Science at Utah State University, Logan, 25 June 1968.

At the end of a thoughtful article on the future of nuclear war, Wiesner and York (1) concluded that: “Both sides in the arms race are ... confronted by the dilemma of steadily increasing military power and steadily decreasing national security. *It is our considered professional judgment that this dilemma has no technical solution.* If the great powers continue to look for solutions in the area of science and technology only, the result will be to worsen the situation.”

I would like to focus your attention not on the subject of the article (national security in a nuclear world) but on the kind of conclusion they reached, namely that there is no technical solution to the problem. An implicit and almost universal assumption of discussions published in professional and semipopular scientific journals is that the problem under discussion has a technical solution. A technical solution may be defined as one that requires a change only in the techniques of the natural sciences, demanding little or nothing in the way of change in human values or ideas of morality.

In our day (though not in earlier times) technical solutions are always welcome. Because of previous failures in prophecy, it takes courage to assert that a desired technical solution is not possible. Wiesner and York exhibited this courage; publishing in a science journal, they insisted that the solution to the problem was not to be found in the natural sciences. They cautiously qualified their statement with the phrase, “It is our considered professional judgment. ...” Whether they were right or not is not the concern of the present article. Rather, the concern here is with the important concept of a class of human problems which can be called “no technical solution problems,” and, more specifically, with the identification and discussion of one of these.

It is easy to show that the class is not a null class. Recall the game of tick-tack-toe. Consider the problem, “How can I win the game of tick-tack-toe?” It is well known that I cannot, if I assume (in keeping with the conventions of game theory) that my opponent understands the game perfectly. Put another way, there is no “technical solution” to the problem. I can win only by giving a radical meaning to the word “win.” I can hit my opponent over the head; or I can drug him; or I can falsify the records. Every way in which I “win” involves, in some sense, an abandonment of the game, as we intuitively understand it. (I can also, of course, openly abandon the game—refuse to play it. This is what most adults do.)

The class of “No technical solution problems” has members. My thesis is that the “population problem,” as conventionally conceived, is a member of this class. How it is conventionally conceived needs some comment. It is fair to say that most people who anguish over the population problem are trying to find a way to avoid the evils of overpopulation without relinquishing any of the privileges they now enjoy. They think that farming the seas or developing new strains of wheat will solve the problem—technologically. I try to show here that the solution they seek cannot be found. The population problem cannot be solved in a technical way, any more than can the problem of winning the game of tick-tack-toe.

What Shall We Maximize?

Population, as Malthus said, naturally tends to grow “geometrically,” or, as we would now say, exponentially. In a finite world this means that the per capita share of the world’s goods must steadily decrease. Is ours a finite world?

A fair defense can be put forward for the view that the world is infinite; or that we do not know that it is not. But, in terms of the practical problems that we must face in the next few generations with the foreseeable technology, it is clear that we will greatly increase human misery if we do not, during the immediate future, assume that the world available to the terrestrial human population is finite. "Space" is no escape (2).

A finite world can support only a finite population; therefore, population growth must eventually equal zero. (The case of perpetual wide fluctuations above and below zero is a trivial variant that need not be discussed.) When this condition is met, what will be the situation of mankind? Specifically, can Bentham's goal of "the greatest good for the greatest number" be realized?

No—for two reasons, each sufficient by itself. The first is a theoretical one. It is not mathematically possible to maximize for two (or more) variables at the same time. This was clearly stated by von Neumann and Morgenstern (3), but the principle is implicit in the theory of partial differential equations, dating back at least to D'Alembert (1717-1783).

The second reason springs directly from biological facts. To live, any organism must have a source of energy (for example, food). This energy is utilized for two purposes: mere maintenance and work. For man, maintenance of life requires about 1600 kilocalories a day ("maintenance calories"). Anything that he does over and above merely staying alive will be defined as work, and is supported by "work calories" which he takes in. Work calories are used not only for what we call work in common speech; they are also required for all forms of enjoyment, from swimming and automobile racing to playing music and writing poetry. If our goal is to maximize population it is obvious what we must do: We must make the work calories per person approach as close to zero as possible. No gourmet meals, no vacations, no sports, no music, no literature, no art. ... I think that everyone will grant, without argument or proof, that maximizing population does not maximize goods. Bentham's goal is impossible.

In reaching this conclusion I have made the usual assumption that it is the acquisition of energy that is the problem. The appearance of atomic energy has led some to question this assumption. However, given an infinite source of energy, population growth still produces an inescapable problem. The problem of the acquisition of energy is replaced by the problem of its dissipation, as J. H. Fremlin has so wittily shown (4). The arithmetic signs in the analysis are, as it were, reversed; but Bentham's goal is still unobtainable.

The optimum population is, then, less than the maximum. The difficulty of defining the optimum is enormous; so far as I know, no one has seriously tackled this problem. Reaching an acceptable and stable solution will surely require more than one generation of hard analytical work—and much persuasion.

We want the maximum good per person; but what is good? To one person it is wilderness, to another it is ski lodges for thousands. To one it is estuaries to nourish ducks for hunters to shoot; to another it is factory land. Comparing one good with another is, we usually say, impossible because goods are incommensurable. Incommensurables cannot be compared.

Theoretically this may be true; but in real life incommensurables *are* commensurable. Only a criterion of judgment and a system of weighting are needed. In nature the criterion is survival. Is it better for a species to be small and hideable, or large and powerful? Natural selection commensurates the incommensurables. The compromise achieved depends on a natural weighting of the values of the variables.

Man must imitate this process. There is no doubt that in fact he already does, but unconsciously. It is when the hidden decisions are made explicit that the arguments begin. The problem for the years ahead is to work out an acceptable theory of weighting. Synergistic effects, nonlinear variation, and difficulties in discounting the future make the intellectual problem difficult, but not (in principle) insoluble.

Has any cultural group solved this practical problem at the present time, even on an intuitive level? One simple fact proves that none has: there is no prosperous population in the world today that has, and has had for some time, a growth rate of zero. Any people that has intuitively identified its optimum point will soon reach it, after which its growth rate becomes and remains zero.

Of course, a positive growth rate might be taken as evidence that a population is below its optimum. However, by any reasonable standards, the most rapidly growing populations on earth today are (in general) the most miserable. This association (which need not be invariable) casts doubt on the optimistic assumption that the positive growth rate of a population is evidence that it has yet to reach its optimum.

We can make little progress in working toward optimum population size until we explicitly exorcize the spirit of Adam Smith in the field of practical demography. In economic affairs, *The Wealth of Nations* (1776) popularized the "invisible hand," the idea that an individual who "intends only his own

gain,” is, as it were, “led by an invisible hand to promote . . . the public interest| (5). Adam Smith did not assert that this was invariably true, and perhaps neither did any of his followers. But he contributed to a dominant tendency of thought that has ever since interfered with positive action based on rational analysis, namely, the tendency to assume that decisions reached individually will, in fact, be the best decisions for an entire society. If this assumption is correct it justifies the continuance of our present policy of *laissez-faire* in reproduction. If it is correct we can assume that men will control their individual fecundity so as to produce the optimum population. If the assumption is not correct, we need to reexamine our individual freedoms to see which ones are defensible.

Tragedy of Freedom in a Commons

The rebuttal to the invisible hand in population control is to be found in a scenario first sketched in a little-known pamphlet (6) in 1833 by a mathematical amateur named William Forster Lloyd (1794-1852). We may well call it “the tragedy of the commons,” using the word “tragedy” as the philosopher Whitehead used it (7): “The essence of dramatic tragedy is not unhappiness. It resides in the solemnity of the remorseless working of things.” He then goes on to say, “This inevitableness of destiny can only be illustrated in terms of human life by incidents which in fact involve unhappiness. For it is only by them that the futility of escape can be made evident in the drama.”

The tragedy of the commons develops in this way. Picture a pasture open to all. It is to be expected that each herdsman will try to keep as many cattle as possible on the commons. Such an arrangement may work reasonably satisfactorily for centuries because tribal wars, poaching, and disease keep the numbers of both man and beast well below the carrying capacity of the land. Finally, however, comes the day of reckoning, that is, the day when the long-desired goal of social stability becomes a reality. At this point, the inherent logic of the commons remorselessly generates tragedy.

As a rational being, each herdsman seeks to maximize his gain. Explicitly or implicitly, more or less consciously, he asks, “What is the utility *to me* of adding one more animal to my herd?” This utility has one negative and one positive component.

1) The positive component is a function of the increment of one animal. Since the herdsman receives all the proceeds from the sale of the additional animal, the positive utility is nearly +1.

2) The negative component is a function of the additional overgrazing created by one more animal. Since, however, the effects of overgrazing are shared by all the herdsmen, the negative utility for any particular decision-making herdsman is only a fraction off -1.

Adding together the component partial utilities, the rational herdsman concludes that the only sensible course for him to pursue is to add another animal to his herd. And another; and another. . . . But this is the conclusion reached by each and every rational herdsman sharing a commons. Therein is the tragedy. Each man is locked into a system that compels him to increase his herd without limit—in a world that is limited. Ruin is the destination toward which all men rush, each pursuing his own best interest in a society that believes in the freedom of the commons. Freedom in a commons brings ruin to all.

Some would say that this is a platitude. Would that it were! In a sense, it was learned thousands of years ago, but natural selection favors the forces of psychological denial (8). The individual benefits as an individual from his ability to deny the truth even though society as a whole, of which he is a part, suffers.

Education can counteract the natural tendency to do the wrong thing, but the inexorable succession of generations requires that the basis for this knowledge be constantly refreshed.

A simple incident that occurred a few years ago in Leominster, Massachusetts, shows how perishable the knowledge is. During the Christmas shopping season the parking meters downtown were covered with plastic bags that bore tags reading: “Do not open until after Christmas. Free parking courtesy of the mayor and city council.” In other words, facing the prospect of an increased demand for already scarce space, the city fathers reinstated the system of the commons. (Cynically, we suspect that they gained more votes than they lost by this retrogressive act.)

In an approximate way, the logic of the commons has been understood for a long time, perhaps since the discovery of agriculture or the invention of private property in real estate. But it is understood mostly only in special cases which are not sufficiently generalized. Even at this late date, cattlemen leasing national land on the western ranges demonstrate no more than an ambivalent understanding, in constantly pressuring federal authorities to increase the head count to the point where overgrazing produces erosion and weed-dominance. Likewise, the oceans of the world continue to suffer from the survival of the

philosophy of the commons. Maritime nations still respond automatically to the shibboleth of the “freedom of the seas.” Professing to believe in the “inexhaustible resources of the oceans,” they bring species after species of fish and whales closer to extinction (9).

The National Parks present another instance of the working out of the tragedy of the commons. At present, they are open to all, without limit. The parks themselves are limited in extent—there is only one Yosemite Valley—whereas population seems to grow without limit. The values that visitors seek in the parks are steadily eroded. Plainly, we must soon cease to treat the parks as commons or they will be of no value to anyone.

What shall we do? We have several options. We might sell them off as private property. We might keep them as public property, but allocate the right to enter them. The allocation might be on the basis of wealth, by the use of an auction system. It might be on the basis of merit, as defined by some agreed-upon standards. It might be by lottery. Or it might be on a first-come, first-served basis, administered to long queues. These, I think, are all the reasonable possibilities. They are all objectionable. But we must choose—or acquiesce in the destruction of the commons that we call our National Parks.

Pollution

In a reverse way, the tragedy of the commons reappears in problems of pollution. Here it is not a question of taking something out of the commons, but of putting something in—sewage, or chemical, radioactive, and heat wastes into water; noxious and dangerous fumes into the air, and distracting and unpleasant advertising signs into the line of sight. The calculations of utility are much the same as before. The rational man finds that his share of the cost of the wastes he discharges into the commons is less than the cost of purifying his wastes before releasing them. Since this is true for everyone, we are locked into a system of “fouling our own nest,” so long as we behave only as independent, rational, free-enterprisers.

The tragedy of the commons as a food basket is averted by private property, or something formally like it. But the air and waters surrounding us cannot readily be fenced, and so the tragedy of the commons as a cesspool must be prevented by different means, by coercive laws or taxing devices that make it cheaper for the polluter to treat his pollutants than to discharge them untreated. We have not progressed as far with the solution of this problem as we have with the first. Indeed, our particular concept of private property, which deters us from exhausting the positive resources of the earth, favors pollution. The owner of a factory on the bank of a stream—whose property extends to the middle of the stream, often has difficulty seeing why it is not his natural right to muddy the waters flowing past his door. The law, always behind the times, requires elaborate stitching and fitting to adapt it to this newly perceived aspect of the commons.

The pollution problem is a consequence of population. It did not much matter how a lonely American frontiersman disposed of his waste. “Flowing water purifies itself every 10 miles,” my grandfather used to say, and the myth was near enough to the truth when he was a boy, for there were not too many people. But as population became denser, the natural chemical and biological recycling processes became overloaded, calling for a redefinition of property rights.

How To Legislate Temperance?

Analysis of the pollution problem as a function of population density uncovers a not generally recognized principle of morality, namely: *the morality of an act is a function of the state of the system at the time it is performed* (10). Using the commons as a cesspool does not harm the general public under frontier conditions, because there is no public, the same behavior in a metropolis is unbearable. A hundred and fifty years ago a plainsman could kill an American bison, cut out only the tongue for his dinner, and discard the rest of the animal. He was not in any important sense being wasteful. Today, with only a few thousand bison left, we would be appalled at such behavior.

In passing, it is worth noting that the morality of an act cannot be determined from a photograph. One does not know whether a man killing an elephant or setting fire to the grassland is harming others until one knows the total system in which his act appears. “One picture is worth a thousand words.” said an ancient Chinese; but it may take 10,000 words to validate it. It is as tempting to ecologists as it is to reformers in general to try to persuade others by way of the photographic shortcut. But the essence of an argument cannot be photographed: it must be presented rationally—in words.

That morality is system-sensitive escaped the attention of most codifiers of ethics in the past. “Thou shalt not . . .” is the form of traditional ethical directives which make no allowance for particular circumstances. The laws of our society follow the pattern of ancient ethics, and therefore are poorly suited to governing a complex, crowded, changeable world. Our epicyclic solution is to augment statutory law with administrative law. Since it is practically impossible to spell out all the conditions under which it is safe to burn trash in the back yard or to run an automobile without smog-control, by law we delegate the details to bureaus. The result is administrative law, which is rightly feared for an ancient reason—*Quis custodiet ipsos custodes?*—“Who shall watch the watchers themselves?” John Adams said that we must have “a government of laws and not men.” Bureau administrators, trying to evaluate the morality of acts in the total system, are singularly liable to corruption, producing a government by men, not laws.

Prohibition is easy to legislate (though not necessarily to enforce); but how do we legislate temperance? Experience indicates that it can be accomplished best through the mediation of administrative law. We limit possibilities unnecessarily if we suppose that the sentiment of *Quis custodiet* denies us the use of administrative law. We should rather retain the phrase as a perpetual reminder of fearful dangers we cannot avoid. The great challenge facing us now is to invent the corrective feedbacks that are needed to keep custodians honest. We must find ways to legitimate the needed authority of both the custodians and the corrective feedbacks.

Freedom To Breed Is Intolerable

The tragedy of the commons is involved in population problems in another way. In a world governed solely by the principle of “dog eat dog”—if indeed there ever was such a world—how many children a family had would not be a matter of public concern. Parents who bred too exuberantly would leave fewer descendants, not more, because they would be unable to care adequately for their children. David Lack and others have found that such a negative feedback demonstrably controls the fecundity of birds (11). But men are not birds, and have not acted like them for millenniums, at least.

If each human family were dependent only on its own resources; *if* the children of improvident parents starved to death; *if*, thus, overbreeding brought its own “punishment” to the germ line—*then* there would be no public interest in controlling the breeding of families. But our society is deeply committed to the welfare state (12), and hence is confronted with another aspect of the tragedy of the commons.

In a welfare state, how shall we deal with the family, the religion, the race, or the class (or indeed any distinguishable and cohesive group) that adopts overbreeding as a policy to secure its own aggrandizement (13)? To couple the concept of freedom to breed with the belief that everyone born has an equal right to the commons is to lock the world into a tragic course of action.

Unfortunately this is just the course of action that is being pursued by the United Nations. In late 1967, some 30 nations agreed to the following (14):

The Universal Declaration of Human Rights describes the family as the natural and fundamental unit of society. It follows that any choice and decision with regard to the size of the family must irrevocably rest with the family itself, and cannot be made by anyone else.

It is painful to have to deny categorically the validity of this right; denying it, one feels as uncomfortable as a resident of Salem, Massachusetts, who denied the reality of witches in the 17th century. At the present time, in liberal quarters, something like a taboo acts to inhibit criticism of the United Nations. There is a feeling that the United Nations is “our last and best hope,” that we shouldn’t find fault with it; we shouldn’t play into the hands of the archconservatives. However, let us not forget what Robert Louis Stevenson said: “The truth that is suppressed by friends is the readiest weapon of the enemy.” If we love the truth we must openly deny the validity of the Universal Declaration of Human Rights, even though it is promoted by the United Nations. We should also join with Kingsley Davis (15) in attempting to get Planned Parenthood-World Population to see the error of its ways in embracing the same tragic ideal.

Conscience Is Self-Eliminating

It is a mistake to think that we can control the breeding of mankind in the long run by an appeal to conscience. Charles Galton Darwin made this point when he spoke on the centennial of the publication of his grandfather’s great book. The argument is straightforward and Darwinian.

People vary. Confronted with appeals to limit breeding, some people will undoubtedly respond to the plea more than others. Those who have more children will produce a larger fraction of the next generation than those with more susceptible consciences. The difference will be accentuated, generation by generation.

In C. G. Darwin's words: "It may well be that it would take hundreds of generations for the progenitive instinct to develop in this way, but if it should do so, nature would have taken her revenge, and the variety *Homo contracpiens* would become extinct and would be replaced by the variety *Homo progenitivus*" (16).

The argument assumes that conscience or the desire for children (no matter which) is hereditary—but hereditary only in the most general formal sense. The result will be the same whether the attitude is transmitted through germ cells, or exosomatically, to use A. J. Lotka's term. (If one denies the latter possibility as well as the former, then what's the point of education?) The argument has here been stated in the context of the population problem, but it applies equally well to any instance in which society appeals to an individual exploiting a commons to restrain himself for the general good—by means of his conscience. To make such an appeal is to set up a selective system that works toward the elimination of conscience from the race.

Pathogenic Effects of Conscience

The long-term disadvantage of an appeal to conscience should be enough to condemn it; but has serious short-term disadvantages as well. If we ask a man who is exploiting a commons to desist "in the name of conscience," what are we saying to him? What does he hear?—not only at the moment but also in the wee small hours of the night when, half asleep, he remembers not merely the words we used but also the nonverbal communication cues we gave him unawares? Sooner or later, consciously or subconsciously, he senses that he has received two communications, and that they are contradictory: (i) (intended communication) "If you don't do as we ask, we will openly condemn you for not acting like a responsible citizen"; (ii) (the unintended communication) "If you do behave as we ask, we will secretly condemn you for a simpleton who can be shamed into standing aside while the rest of us exploit the commons."

Everyman then is caught in what Bateson has called a "double bind." Bateson and his co-workers have made a plausible case for viewing the double bind as an important causative factor in the genesis of schizophrenia (17). The double bind may not always be so damaging, but it always endangers the mental health of anyone to whom it is applied. "A bad conscience," said Nietzsche, "is a kind of illness."

To conjure up a conscience in others is tempting to anyone who wishes to extend his control beyond the legal limits. Leaders at the highest level succumb to this temptation. Has any President during the past generation failed to call on labor unions to moderate voluntarily their demands for higher wages, or to steel companies to honor voluntary guidelines on prices? I can recall none. The rhetoric used on such occasions is designed to produce feelings of guilt in noncooperators.

For centuries it was assumed without proof that guilt was a valuable, perhaps even an indispensable, ingredient of the civilized life. Now, in this post-Freudian world, we doubt it.

Paul Goodman speaks from the modern point of view when he says: "No good has ever come from feeling guilty, neither intelligence, policy, nor compassion. The guilty do not pay attention to the object but only to themselves, and not even to their own interests, which might make sense, but to their anxieties" (18).

One does not have to be a professional psychiatrist to see the consequences of anxiety. We in the Western world are just emerging from a dreadful two-centuries-long Dark Ages of Eros that was sustained partly by prohibition laws, but perhaps more effectively by the anxiety-generating mechanism of education. Alex Comfort has told the story well in *The Anxiety Makers* (19); it is not a pretty one.

Since proof is difficult, we may even concede that the results of anxiety may sometimes, from certain points of view, be desirable. The larger question we should ask is whether, as a matter of policy, we should ever encourage the use of a technique the tendency (if not the intention) of which is psychologically pathogenic. We hear much talk these days of responsible parenthood; the coupled words are incorporated into the titles of some organizations devoted to birth control. Some people have proposed massive propaganda campaigns to instill responsibility into the nation's (or the world's) breeders. But what is the meaning of the word responsibility in this context? Is it not merely a synonym for the word conscience? When we use the word responsibility in the absence of substantial sanctions are we not trying to browbeat a

free man in a commons into acting against his own interest? Responsibility is a verbal counterfeit for a substantial *quid pro quo*. It is an attempt to get something for nothing.

If the word responsibility is to be used at all, I suggest that it be in the sense Charles Frankel uses it (20). “Responsibility,” says this philosopher, “is the product of definite social arrangements.” Notice that Frankel calls for social arrangements—not propaganda.

Mutual Coercion Mutually Agreed upon

The social arrangements that produce responsibility are arrangements that create coercion, of some sort. Consider bank-robbing. The man who takes money from a bank acts as if the bank were a commons. How do we prevent such action? Certainly not by trying to control his behavior solely by a verbal appeal to his sense of responsibility. Rather than rely on propaganda we follow Frankel’s lead and insist that a bank is not a commons; we seek the definite social arrangements that will keep it from becoming a commons. That we thereby infringe on the freedom of would-be robbers we neither deny nor regret.

The morality of bank-robbing is particularly easy to understand because we accept complete prohibition of this activity. We are willing to say “Thou shalt not rob banks,” without providing for exceptions. But temperance also can be created by coercion. Taxing is a good coercive device. To keep downtown shoppers temperate in their use of parking space we introduce parking meters for short periods, and traffic fines for longer ones. We need not actually forbid a citizen to park as long as he wants to; we need merely make it increasingly expensive for him to do so. Not prohibition, but carefully biased options are what we offer him. A Madison Avenue man might call this persuasion; I prefer the greater candor of the word coercion.

Coercion is a dirty word to most liberals now, but it need not forever be so. As with the four-letter words, its dirtiness can be cleansed away by exposure to the light, by saying it over and over without apology or embarrassment. To many, the word coercion implies arbitrary decisions of distant and irresponsible bureaucrats; but this is not a necessary part of its meaning. The only kind of coercion I recommend is mutual coercion, mutually agreed upon by the majority of the people affected.

To say that we mutually agree to coercion is not to say that we are required to enjoy it, or even to pretend we enjoy it. Who enjoys taxes? We all grumble about them. But we accept compulsory taxes because we recognize that voluntary taxes would favor the conscienceless. We institute and (grumblingly) support taxes and other coercive devices to escape the horror of the commons.

An alternative to the commons need not be perfectly just to be preferable. With real estate and other material goods, the alternative we have chosen is the institution of private property coupled with legal inheritance. Is this system perfectly just? As a genetically trained biologist I deny that it is. It seems to me that, if there are to be differences in individual inheritance, legal possession should be perfectly correlated with biological inheritance—that those who are biologically more fit to be the custodians of property and power should legally inherit more. But genetic recombination continually makes a mockery of the doctrine of “like father, like son” implicit in our laws of legal inheritance. An idiot can inherit millions, and a trust fund can keep his estate intact. We must admit that our legal system of private property plus inheritance is unjust—but we put up with it because we are not convinced, at the moment, that anyone has invented a better system. The alternative of the commons is too horrifying to contemplate. Injustice is preferable to total ruin.

It is one of the peculiarities of the warfare between reform and the status quo that it is thoughtlessly governed by a double standard. Whenever a reform measure is proposed it is often defeated when its opponents triumphantly discover a flaw in it. As Kingsley Davis has pointed out (21), worshippers of the status quo sometimes imply that no reform is possible without unanimous agreement, an implication contrary to historical fact. As nearly as I can make out, automatic rejection of proposed reforms is based on one of two unconscious assumptions: (i) that the status quo is perfect; or (ii) that the choice we face is between reform and no action; if the proposed reform is imperfect, we presumably should take no action at all, while we wait for a perfect proposal.

But we can never do nothing. That which we have done for thousands of years is also action. It also produces evils. Once we are aware that the status quo is action, we can then compare its discoverable advantages and disadvantages with the predicted advantages and disadvantages of the proposed reform, discounting as best we can for our lack of experience. On the basis of such a comparison, we can make a rational decision which will not involve the unworkable assumption that only perfect systems are tolerable.

Recognition of Necessity

Perhaps the simplest summary of this analysis of man's population problems is this: the commons, if justifiable at all, is justifiable only under conditions of low-population density. As the human population has increased, the commons has had to be abandoned in one aspect after another.

First we abandoned the commons in food gathering, enclosing farm land and restricting pastures and hunting and fishing areas. These restrictions are still not complete throughout the world.

Somewhat later we saw that the commons as a place for waste disposal would also have to be abandoned. Restrictions on the disposal of domestic sewage are widely accepted in the Western world; we are still struggling to close the commons to pollution by automobiles, factories, insecticide sprayers, fertilizing operations, and atomic energy installations.

In a still more embryonic state is our recognition of the evils of the commons in matters of pleasure. There is almost no restriction on the propagation of sound waves in the public medium. The shopping public is assaulted with mindless music, without its consent. Our government is paying out billions of dollars to create supersonic transport which will disturb 50,000 people for every one person who is whisked from coast to coast 3 hours faster. Advertisers muddy the airwaves of radio and television and pollute the view of travelers. We are a long way from outlawing the commons in matters of pleasure. Is this because our Puritan inheritance makes us view pleasure as something of a sin, and pain (that is, the pollution of advertising) as the sign of virtue?

Every new enclosure of the commons involves the infringement of somebody's personal liberty. Infringements made in the distant past are accepted because no contemporary complains of a loss. It is the newly proposed infringements that we vigorously oppose; cries of "rights" and "freedom" fill the air. But what does "freedom" mean? When men mutually agreed to pass laws against robbing, mankind became more free, not less so. Individuals locked into the logic of the commons are free only to bring on universal ruin once they see the necessity of mutual coercion; they become free to pursue other goals. I believe it was Hegel who said, "Freedom is the recognition of necessity."

The most important aspect of necessity that we must now recognize, is the necessity of abandoning the commons in breeding. No technical solution can rescue us from the misery of overpopulation. Freedom to breed will bring ruin to all. At the moment, to avoid hard decisions many of us are tempted to propagandize for conscience and responsible parenthood. The temptation must be resisted, because an appeal to independently acting consciences selects for the disappearance of all conscience in the long run, and an increase in anxiety in the short.

The only way we can preserve and nurture other and more precious freedoms is by relinquishing the freedom to breed, and that very soon. "Freedom is the recognition of necessity"—and it is the role of education to reveal to all the necessity of abandoning the freedom to breed. Only so, can we put an end to this aspect of the tragedy of the commons.

References

1. J. B. Wiesner and H. F. York, *Sci. Amer.* **211** (No. 4), 27 (1964).
2. G. Hardin, *J. Hered.* **50**, 68 (1959); S. von Hoernor, *Science* **137**, 18 (1962).
3. J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton Univ. Press, Princeton, N.J., 1947), p.11.
4. J. H. Fremlin, *New Sci.*, No. **415** (1964), p. 285.
5. A. Smith, *The Wealth of Nations* (Modern Library, New York, 1937), p. 423.
6. W. F. Lloyd, *Two Lectures on the Checks to Population* (Oxford Univ. Press, Oxford, England, 1833), reprinted (in part) in *Population, Evolution, and Birth Control*, G. Hardin, Ed. (Freeman, San Francisco, 1964), p. 37.
7. A. N. Whitehead, *Science and the Modern World* (Mentor, New York, 1948), p. 17.
8. G. Hardin, *Population, Evolution, and Birth Control* (Freeman, San Francisco, 1964), p. 56.
9. S. McVay, *Sci. Amer.* **216** (No. 8), 13 (1966).
10. J. Fletcher, *Situation Ethics* (Westminster, Philadelphia, 1966).
11. D. Lack, *The Natural Regulation of Animal Numbers* (Clarendon Press, Oxford, 1954).
12. H. Girvetz, *From Wealth to Welfare* (Stanford Univ. Press, Stanford, Calif., 1950).
13. G. Hardin, *Perspec. Biol. Med.* **6**, 366 (1963).
14. U. Thant, *Int. Planned Parenthood News*, No. 168 (February 1968), p. 3.
15. K. Davis, *Science* **158**, 730 (1967).
16. S. Tax, Ed., *Evolution after Darwin* (Univ. of Chicago Press, Chicago, 1960), vol. 2, p. 469.
17. G. Bateson, D. D. Jackson, J. Haley, J. Weakland, *Behav. Sci.* **1**, 251 (1956).
18. P. Goodman, *New York Rev. Books* **10**(8), 22 (23 May 1968).
19. A. Comfort, *The Anxiety Makers* (Nelson, London, 1967).
20. C. Frankel, *The Case for Modern Man* (Harper, New York, 1955), p. 203.
21. J. D. Roslansky, *Genetics and the Future of Man* (Appleton-Century-Crofts, New York, 1966), p. 177.

CHAPTER 16

The World as a Polder: What Does It All Mean to Us Today?

Introduction ■ The most serious problems ■ If we don't solve them . . . ■ Life in Los Angeles ■ One-liner objections ■ The past and the present ■ Reasons for hope ■

The chapters of this book have discussed why past or present societies succeed or fail at solving their environmental problems. Now, this final chapter considers the book's practical relevance: what does it all mean to us today?

I shall begin by explaining the major sets of environmental problems facing modern societies, and the time scale on which they pose threats. As a specific example of how these problems play out, I examine the area where I have spent most of the last 39 years of my life, Southern California. I then consider the objections most often raised to dismiss the significance of environmental problems today. Since half of this book was devoted to ancient societies because of the lessons that they might hold for modern societies, I look at differences between the ancient and the modern worlds that affect what lessons we can draw from the past. Finally, for anyone who asks, "What can I do as an individual?" I offer suggestions in the Further Readings section.

It seems to me that the most serious environmental problems facing past and present societies fall into a dozen groups. Eight of the 12 were significant already in the past, while four (numbers 5, 7, 8, and 10: energy, the photosynthetic ceiling, toxic chemicals, and atmospheric changes) became serious only recently. The first four of the 12 consist of destruction or losses of natural resources; the next three involve ceilings on natural resources; the three after that consist of harmful things that we produce or move around; and the last two are population issues. Let's begin with the natural resources

that we are destroying or losing: natural habitats, wild food sources, biological diversity, and soil.

1. At an accelerating rate, we are destroying natural habitats or else converting them to human-made habitats, such as cities and villages, farmlands and pastures, roads, and golf courses. The natural habitats whose losses have provoked the most discussion are forests, wetlands, coral reefs, and the ocean bottom. As I mentioned in the preceding chapter, more than half of the world's original area of forest has already been converted to other uses, and at present conversion rates one-quarter of the forests that remain will become converted within the next half-century. Those losses of forests represent losses for us humans, especially because forests provide us with timber and other raw materials, and because they provide us with so-called ecosystem services such as protecting our watersheds, protecting soil against erosion, constituting essential steps in the water cycle that generates much of our rainfall, and providing habitat for most terrestrial plant and animal species. Deforestation was a or *the* major factor in all the collapses of past societies described in this book. In addition, as discussed in Chapter 1 in connection with Montana, issues of concern to us are not only forest destruction and conversion, but also changes in the structure of wooded habitats that do remain. Among other things, that changed structure results in changed fire regimes that put forests, chaparral woodlands, and savannahs at greater risk of infrequent but catastrophic fires.

Other valuable natural habitats besides forests are also being destroyed. An even larger fraction of the world's original wetlands than of its forests has already been destroyed, damaged, or converted. Consequences for us arise from wetlands' importance in maintaining the quality of our water supplies and the existence of commercially important freshwater fisheries, while even ocean fisheries depend on mangrove wetlands to provide habitat for the juvenile phase of many fish species. About one-third of the world's coral reefs—the oceanic equivalent of tropical rainforests, because they are home to a disproportionate fraction of the ocean's species—have already been severely damaged. If current trends continue, about half of the remaining reefs would be lost by the year 2030. That damage and destruction result from the growing use of dynamite as a fishing method, reef overgrowth by algae ("seaweeds") when the large herbivorous fish that normally graze on the algae become fished out, effects of sediment runoff and pollutants from adjacent lands cleared or converted to agriculture, and coral

bleaching due to rising ocean water temperatures. It has recently become appreciated that fishing by trawling is destroying much or most of the shallow ocean bottom and the species dependent on it.

2. Wild foods, especially fish and to a lesser extent shellfish, contribute a large fraction of the protein consumed by humans. In effect, this is protein that we obtain for free (other than the cost of catching and transporting the fish), and that reduces our needs for animal protein that we have to grow ourselves in the form of domestic livestock. About two billion people, most of them poor, depend on the oceans for protein. If wild fish stocks were managed appropriately, the stock levels could be maintained, and they could be harvested perpetually. Unfortunately, the problem known as the tragedy of the commons (Chapter 14) has regularly undone efforts to manage fisheries sustainably, and the great majority of valuable fisheries already either have collapsed or are in steep decline (Chapter 15). Past societies that overfished included Easter Island, Mangareva, and Henderson.

Increasingly, fish and shrimp are being grown by aquaculture, which in principle has a promising future as the cheapest way to produce animal protein. In several respects, though, aquaculture as commonly practiced today is making the problem of declining wild fisheries worse rather than better. Fish grown by aquaculture are mostly fed wild-caught fish and thereby usually consume more wild fish meat (up to 20 times more) than they yield in meat of their own. They contain higher toxin levels than do wild-caught fish. Cultured fish regularly escape, interbreed with wild fish, and thereby harm wild fish stocks genetically, because cultured fish strains have been selected for rapid growth at the expense of poor survival in the wild (50 times worse survival for cultured salmon than for wild salmon). Aquaculture runoff causes pollution and eutrophication. The lower costs of aquaculture than of fishing, by driving down fish prices, initially drive fishermen to exploit wild fish stocks even more heavily in order to maintain their incomes constant when they are receiving less money per pound of fish.

3. A significant fraction of wild species, populations, and genetic diversity has already been lost, and at present rates a large fraction of what remains will be lost within the next half-century. Some species, such as big edible animals, or plants with edible fruits or good timber, are of obvious value to us. Among the many past societies that harmed themselves by exterminating such species were the Easter and Henderson Islanders whom we have discussed.

But biodiversity losses of small inedible species often provoke the response, "Who cares? Do you really care less for humans than for some lousy

useless little fish or weed, like the snail darter or Furbish lousewort?" This response misses the point that the entire natural world is made up of wild species providing us for free with services that can be very expensive, and in many cases impossible, for us to supply ourselves. Elimination of lots of lousy little species regularly causes big harmful consequences for humans, just as does randomly knocking out many of the lousy little rivets holding together an airplane. The literally innumerable examples include: the role of earthworms in regenerating soil and maintaining its texture (one of the reasons that oxygen levels dropped inside the Biosphere 2 enclosure, harming its human inhabitants and crippling a colleague of mine, was a lack of appropriate earthworms, contributing to altered soil/atmosphere gas exchange); soil bacteria that fix the essential crop nutrient nitrogen, which otherwise we have to spend money to supply in fertilizers; bees and other insect pollinators (they pollinate our crops for free, whereas it's expensive for us to pollinate every crop flower by hand); birds and mammals that disperse wild fruits (foresters still haven't figured out how to grow from seed the most important commercial tree species of the Solomon Islands, whose seeds are naturally dispersed by fruit bats, which are becoming hunted out); elimination of whales, sharks, bears, wolves, and other top predators in the wild and on the land, changing the whole food chain beneath them; and wild plants and animals that decompose wastes and recycle nutrients, ultimately providing us with clean water and air.

4. Soils of farmlands used for growing crops are being carried away by water and wind erosion at rates between 10 and 40 times the rates of soil formation, and between 500 and 10,000 times soil erosion rates on forested land. Because those soil erosion rates are so much higher than soil formation rates, that means a net loss of soil. For instance, about half of the topsoil of Iowa, the state whose agriculture productivity is among the highest in the U.S., has been eroded in the last 150 years. On my most recent visit to Iowa, my hosts showed me a churchyard offering a dramatically visible example of those soil losses. A church was built there in the middle of farmland during the 19th century and has been maintained continuously as a church ever since, while the land around it was being farmed. As a result of soil being eroded much more rapidly from fields than from the churchyard, the yard now stands like a little island raised 10 feet above the surrounding sea of farmland.

Other types of soil damage caused by human agricultural practices include salinization, as discussed for Montana, China, and Australia in Chapters 1, 12, and 13; losses of soil fertility, because farming removes nutrients

much more rapidly than they are restored by weathering of the underlying rocks; and soil acidification in some areas, or its converse, alkalization, in other areas. All of these types of harmful impacts have resulted in a fraction of the world's farmland variously estimated at between 20% and 80% having become severely damaged, during an era in which increasing human population has caused us to need more farmland rather than less farmland. Like deforestation, soil problems contributed to the collapses of all past societies discussed in this book.

The next three problems involve ceilings—on energy, freshwater, and photosynthetic capacity. In each case the ceiling is not hard and fixed but soft: we can obtain more of the needed resource, but at increasing costs.

5. The world's major energy sources, especially for industrial societies, are fossil fuels: oil, natural gas, and coal. While there has been much discussion about how many big oil and gas fields remain to be discovered, and while coal reserves are believed to be large, the prevalent view is that known and likely reserves of readily accessible oil and natural gas will last for a few more decades. This view should not be misinterpreted to mean that all of the oil and natural gas within the Earth will have been used up by then. Instead, further reserves will be deeper underground, dirtier, increasingly expensive to extract or process, or will involve higher environmental costs. Of course, fossil fuels are not our sole energy sources, and I shall consider problems raised by the alternatives below.

6. Most of the world's freshwater in rivers and lakes is already being utilized for irrigation, domestic and industrial water, and in situ uses such as boat transportation corridors, fisheries, and recreation. Rivers and lakes that are not already utilized are mostly far from major population centers and likely users, such as in Northwestern Australia, Siberia, and Iceland. Throughout the world, freshwater underground aquifers are being depleted at rates faster than they are being naturally replenished, so that they will eventually dwindle. Of course, freshwater can be made by desalination of seawater, but that costs money and energy, as does pumping the resulting desalinated water inland for use. Hence desalination, while it is useful locally, is too expensive to solve most of the world's water shortages. The Anasazi and Maya were among the past societies to be undone by water problems, while today over a billion people lack access to reliable safe drinking water.

7. It might at first seem that the supply of sunlight is infinite, so one

might reason that the Earth's capacity to grow crops and wild plants is also infinite. Within the last 20 years, it has been appreciated that that is not the case, and that's not only because plants grow poorly in the world's Arctic regions and deserts unless one goes to the expense of supplying heat or water. More generally, the amount of solar energy fixed per acre by plant photosynthesis, hence plant growth per acre, depends on temperature and rainfall. At any given temperature and rainfall the plant growth that can be supported by the sunlight falling on an acre is limited by the geometry and biochemistry of plants, even if they take up the sunlight so efficiently that not a single photon of light passes through the plants unabsorbed to reach the ground. The first calculation of this photosynthetic ceiling, carried out in 1986, estimated that humans then already used (e.g., for crops, tree plantations, and golf courses) or diverted or wasted (e.g., light falling on concrete roads and buildings) about half of the Earth's photosynthetic capacity. Given the rate of increase of human population, and especially of population impact (see point 12 below), since 1986, we are projected to be utilizing most of the world's terrestrial photosynthetic capacity by the middle of this century. That is, most energy fixed from sunlight will be used for human purposes, and little will be left over to support the growth of natural plant communities, such as natural forests.

The next three problems involve harmful things that we generate or move around: toxic chemicals, alien species, and atmospheric gases.

8. The chemical industry and many other industries manufacture or release into the air, soil, oceans, lakes, and rivers many toxic chemicals, some of them "unnatural" and synthesized only by humans, others present naturally in tiny concentrations (e.g., mercury) or else synthesized by living things but synthesized and released by humans in quantities much larger than natural ones (e.g., hormones). The first of these toxic chemicals to achieve wide notice were insecticides, pesticides, and herbicides, whose effects on birds, fish, and other animals were publicized by Rachel Carson's 1962 book *Silent Spring*. Since then, it has been appreciated that the toxic effects of even greater significance for us humans are those on ourselves. The culprits include not only insecticides, pesticides, and herbicides, but also mercury and other metals, fire-retardant chemicals, refrigerator coolants, detergents, and components of plastics. We swallow them in our food and water, breathe them in our air, and absorb them through our skin. Often in very low concentrations, they variously cause birth defects, mental

retardation, and temporary or permanent damage to our immune and reproductive systems. Some of them act as endocrine disruptors, i.e., they interfere with our reproductive systems by mimicking or blocking effects of our own sex hormones. They probably make the major contribution to the steep decline in sperm count in many human populations over the last several decades, and to the apparently increasing frequency with which couples are unable to conceive, even when one takes into account the increasing average age of marriage in many societies. In addition, deaths in the U.S. from air pollution alone (without considering soil and water pollution) are conservatively estimated at over 130,000 per year.

Many of these toxic chemicals are broken down in the environment only slowly (e.g., DDT and PCBs) or not at all (mercury), and they persist in the environment for long times before being washed out. Thus, cleanup costs of many polluted sites in the U.S. are measured in the billions of dollars (e.g., Love Canal, the Hudson River, Chesapeake Bay, the *Exxon Valdez* oil spill, and Montana copper mines). But pollution at those worst sites in the U.S. is mild compared to that in the former Soviet Union, China, and many Third World mines, whose cleanup costs no one even dares to think about.

9. The term "alien species" refers to species that we transfer, intentionally or inadvertently, from a place where they are native to another place where they are not native. Some alien species are obviously valuable to us as crops, domestic animals, and landscaping. But others devastate populations of native species with which they come in contact, either by preying on, parasitizing, infecting, or outcompeting them. The aliens cause these big effects because the native species with which they come in contact had no previous evolutionary experience of them and are unable to resist them (like human populations newly exposed to smallpox or AIDS). There are by now literally hundreds of cases in which alien species have caused one-time or annually recurring damages of hundreds of millions of dollars or even billions of dollars. Modern examples include Australia's rabbits and foxes, agricultural weeds like Spotted Knapweed and Leafy Spurge (Chapter 1), pests and pathogens of trees and crops and livestock (like the blights that wiped out American chestnut trees and devastated American elms), the water hyacinth that chokes waterways, the zebra mussels that choke power plants, and the lampreys that devastated the former commercial fisheries of the North American Great Lakes (Plates 30, 31). Ancient examples include the introduced rats that contributed to the extinction of Easter Island's palm tree by gnawing its nuts, and that ate the eggs and chicks of nesting birds on Easter, Henderson, and all other Pacific islands previously without rats.

10. Human activities produce gases that escape into the atmosphere, where they either damage the protective ozone layer (as do formerly widespread refrigerator coolants) or else act as greenhouse gases that absorb sunlight and thereby lead to global warming. The gases contributing to global warming include carbon dioxide from combustion and respiration, and methane from fermentation in the intestines of ruminant animals. Of course, there have always been natural fires and animal respiration producing carbon dioxide, and wild ruminant animals producing methane, but our burning of firewood and of fossil fuels has greatly increased the former, and our herds of cattle and of sheep have greatly increased the latter.

For many years, scientists debated the reality, cause, and extent of global warming: are world temperatures really historically high now, and, if so, by how much, and are humans the leading cause? Most knowledgeable scientists now agree that, despite year-to-year ups and downs of temperature that necessitate complicated analyses to extract warming trends, the atmosphere really has been undergoing an unusually rapid rise in temperature recently, and that human activities are the or a major cause. The remaining uncertainties mainly concern the future expected magnitude of the effect: e.g., whether average global temperatures will increase by "just" 1.5 degrees Centigrade or by 5 degrees Centigrade over the next century. Those numbers may not sound like a big deal, until one reflects that average global temperatures were "only" 5 degrees cooler at the height of the last Ice Age.

While one might at first think that we should welcome global warming on the grounds that warmer temperatures mean faster plant growth, it turns out that global warming will produce both winners and losers. Crop yields in cool areas with temperatures marginal for agriculture may indeed increase, while crop yields in already warm or dry areas may decrease. In Montana, California, and many other dry climates, the disappearance of mountain snowpacks will decrease the water available for domestic uses, and for irrigation that actually limits crop yields in those areas. The rise in global sea levels as a result of snow and ice melting poses dangers of flooding and coastal erosion for densely populated low-lying coastal plains and river deltas already barely above or even below sea level. The areas thereby threatened include much of the Netherlands, Bangladesh, and the seaboard of the eastern U.S., many low-lying Pacific islands, the deltas of the Nile and Mekong Rivers, and coastal and riverbank cities of the United Kingdom (e.g., London), India, Japan, and the Philippines. Global warming will also produce big secondary effects that are difficult to predict exactly in advance and that are likely to cause huge problems, such as further climate changes

resulting from changes in ocean circulation resulting in turn from melting of the Arctic ice cap.

The remaining two problems involve the increase in human population:

11. The world's human population is growing. More people require more food, space, water, energy, and other resources. Rates and even the direction of human population change vary greatly around the world, with the highest rates of population growth (4% per year or higher) in some Third World countries, low rates of growth (1% per year or less) in some First World countries such as Italy and Japan, and negative rates of growth (i.e., decreasing populations) in countries facing major public health crises, such as Russia and AIDS-affected African countries. Everybody agrees that the world population is increasing, but that its annual percentage rate of increase is not as high as it was a decade or two ago. However, there is still disagreement about whether the world's population will stabilize at some value above its present level (double the present population?), and (if so) how many years (30 years? 50 years?) it will take for population to reach that level, or whether population will continue to grow.

There is long built-in momentum to human population growth because of what is termed the "demographic bulge" or "population momentum," i.e., a disproportionate number of children and young reproductive-age people in today's population, as a result of recent population growth. That is, suppose that every couple in the world decided tonight to limit themselves to two children, approximately the correct number of children to yield an unchanging population in the long run by exactly replacing their two parents who will eventually die (actually, 2.1 children when one considers childless couples and children who won't marry). The world's population would nevertheless continue to increase for about 70 years, because more people today are of reproductive age or entering reproductive age than are old and post-reproductive. The problem of human population growth has received much attention in recent decades and has given rise to movements such as Zero Population Growth, which aim to slow or halt the increase in the world's population.

12. What really counts is not the number of people alone, but their impact on the environment. If most of the world's 6 billion people today were in cryogenic storage and neither eating, breathing, nor metabolizing, that large population would cause no environmental problems. Instead, our numbers pose problems insofar as we consume resources and generate

wastes. That per-capita impact—the resources consumed, and the wastes put out, by each person—varies greatly around the world, being highest in the First World and lowest in the Third World. On the average, each citizen of the U.S., western Europe, and Japan consumes 32 times more resources such as fossil fuels, and puts out 32 times more wastes, than do inhabitants of the Third World (Plate 35).

But low-impact people are becoming high-impact people for two reasons: rises in living standards in Third World countries whose inhabitants see and covet First World lifestyles; and immigration, both legal and illegal, of individual Third World inhabitants into the First World, driven by political, economic, and social problems at home. Immigration from low-impact countries is now the main contributor to the increasing populations of the U.S. and Europe. By the same token, the overwhelmingly most important human population problem for the world as a whole is not the high rate of population increase in Kenya, Rwanda, and some other poor Third World countries, although that certainly does pose a problem for Kenya and Rwanda themselves, and although that is the population problem most discussed. Instead, the biggest problem is the increase in total human impact, as the result of rising Third World living standards, and of Third World individuals moving to the First World and adopting First World living standards.

There are many "optimists" who argue that the world could support double its human population, and who consider only the increase in human numbers and not the average increase in per-capita impact. But I have not met anyone who seriously argues that the world could support 12 times its current impact, although an increase of that factor would result from all Third World inhabitants adopting First World living standards. (That factor of 12 is less than the factor of 32 that I mentioned in the preceding paragraph, because there are already First World inhabitants with high-impact lifestyles, although they are greatly outnumbered by Third World inhabitants.) Even if the people of China alone achieved a First World living standard while everyone else's living standard remained constant, that would double our human impact on the world (Chapter 12).

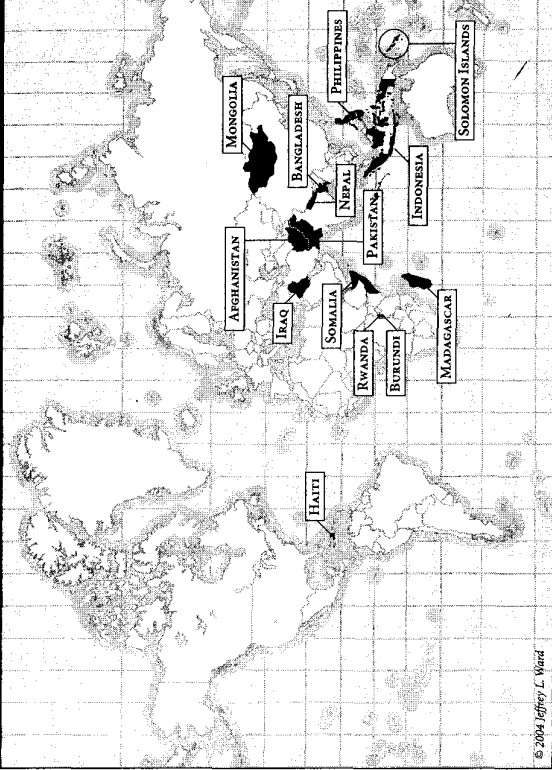
People in the Third World aspire to First World living standards. They develop that aspiration through watching television, seeing advertisements for First World consumer products sold in their countries, and observing First World visitors to their countries. Even in the most remote villages and refugee camps today, people know about the outside world. Third World citizens are encouraged in that aspiration by First World and United

Nations development agencies, which hold out to them the prospect of achieving their dream if they will only adopt the right policies, like balancing their national budgets, investing in education and infrastructure, and so on.

But no one at the U.N. or in First World governments is willing to acknowledge the dream's impossibility: the unsustainability of a world in which the Third World's large population were to reach and maintain current First World living standards. It is impossible for the First World to resolve that dilemma by blocking the Third World's efforts to catch up: South Korea, Malaysia, Singapore, Hong Kong, Taiwan, and Mauritius have already succeeded or are close to success; China and India are progressing rapidly by their own efforts; and the 15 rich Western European countries making up the European Union have just extended Union membership to 10 poorer countries of Eastern Europe, in effect thereby pledging to help those 10 countries catch up. Even if the human populations of the Third World did not exist, it would be impossible for the First World alone to maintain its present course, because it is not in a steady state but is depleting its own resources as well as those imported from the Third World. At present, it is untenable politically for First World leaders to propose to their own citizens that they lower their living standards, as measured by lower resource consumption and waste production rates. What will happen when it finally dawns on all those people in the Third World that current First World standards are unreachable for them, and that the First World refuses to abandon those standards for itself? Life is full of agonizing choices based on trade-offs, but that's the cruelest trade-off that we shall have to resolve: encouraging and helping all people to achieve a higher standard of living, without thereby undermining that standard through overstressing global resources.

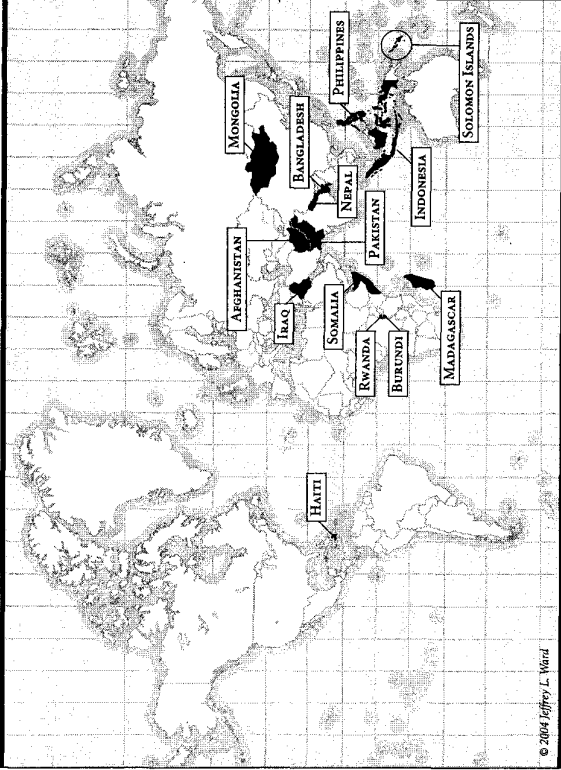
I have described these 12 sets of problems as separate from each other. In fact, they are linked: one problem exacerbates another or makes its solution more difficult. For example, human population growth affects all 11 other problems: more people means more deforestation, more toxic chemicals, more demand for wild fish, etc. The energy problem is linked to other problems because use of fossil fuels for energy contributes heavily to greenhouse gases, the combating of soil fertility losses by using synthetic fertilizers requires energy to make the fertilizers, fossil fuel scarcity increases our interest in nuclear energy which poses potentially the biggest "toxic" problem of all in case of an accident, and fossil fuel scarcity also makes it more expensive to solve our freshwater problems by using energy to desalinate ocean

POLITICAL TROUBLE SPOTS OF THE MODERN WORLD



© 2004 Jeffrey L. Ward

ENVIRONMENTAL TROUBLE SPOTS OF THE MODERN WORLD



© 2004 Jeffrey L. Ward

water. Depletion of fisheries and other wild food sources puts more pressure on livestock, crops, and aquaculture to replace them, thereby leading to more topsoil losses and more eutrophication from agriculture and aquaculture. Problems of deforestation, water shortage, and soil degradation in the Third World foster wars there and drive legal asylum seekers and illegal emigrants to the First World from the Third World.

Our world society is presently on a non-sustainable course, and any of our 12 problems of non-sustainability that we have just summarized would suffice to limit our lifestyle within the next several decades. They are like time bombs with fuses of less than 50 years. For example, destruction of accessible lowland tropical rainforest outside national parks is already virtually complete in Peninsular Malaysia, will be complete at current rates within less than a decade in the Solomon Islands, the Philippines, on Sumatra, and on Sulawesi, and will be complete around the world except perhaps for parts of the Amazon Basin and Congo Basin within 25 years. At current rates, we shall have depleted or destroyed most of the world's remaining marine fisheries, depleted clean or cheap or readily accessible reserves of oil and natural gas, and approached the photosynthetic ceiling within a few decades. Global warming is projected to have reached a degree Centigrade or more, and a substantial fraction of the world's wild animal and plant species are projected to be endangered or past the point of no return, within half a century. People often ask, "What is the single most important environmental/population problem facing the world today?" A flip answer would be, "The single most important problem is our misguided focus on identifying the single most important problem!" That flip answer is essentially correct, because any of the dozen problems if unsolved would do us grave harm, and because they all interact with each other. If we solved 11 of the problems, but not the 12th, we would still be in trouble, whichever was the problem that remained unsolved. We have to solve them all.

Thus, because we are rapidly advancing along this non-sustainable course, the world's environmental problems *will* get resolved, in one way or another, within the lifetimes of the children and young adults alive today. The only question is whether they will become resolved in pleasant ways of our own choice, or in unpleasant ways not of our choice, such as warfare, genocide, starvation, disease epidemics, and collapses of societies. While all of those grim phenomena have been endemic to humanity throughout our history, their frequency increases with environmental degradation, population pressure, and the resulting poverty and political instability.

Examples of those unpleasant solutions to environmental and popula-

tion problems abound in both the modern world and the ancient world. The examples include the recent genocides in Rwanda, Burundi, and the former Yugoslavia; war, civil war, or guerrilla war in the modern Sudan, Philippines, and Nepal, and in the ancient Maya homeland; cannibalism on prehistoric Easter Island and Mangareva and among the ancient Anasazi; starvation in many modern African countries and on prehistoric Easter Island; the AIDS epidemic already in Africa, and incipiently elsewhere; and the collapse of state government in modern Somalia, the Solomon Islands, and Haiti, and among the ancient Maya. An outcome less drastic than a worldwide collapse might "merely" be the spread of Rwanda-like or Haiti-like conditions to many more developing countries, while we First World inhabitants retain many of our First World amenities but face a future with which we are unhappy, beset by more chronic terrorism, wars, and disease outbreaks. But it is doubtful that the First World could retain its separate lifestyle in the face of desperate waves of immigrants fleeing from collapsing Third World countries, in numbers much larger than the current unstoppable influx. I'm reminded again of how I picture the end of Gardar Cathedral Farm and its splendid cattle barn on Greenland, overwhelmed by the influx of Norse from poorer farms where all the livestock had died or been eaten.

But before we let ourselves give way to this one-sidedly pessimistic scenario, let's examine further the problems facing us, and their complexities. This will bring us, I feel, to a position of cautious optimism.

To make the preceding discussion less abstract, I shall now illustrate how those dozen environmental problems affect lifestyles in the part of the world with which I am most familiar: the city of Los Angeles in Southern California, where I live. After growing up on the East Coast of the United States and living for several years in Europe, I first visited California in 1964. It immediately appealed to me, and I moved here in 1966.

Thus, I have seen how Southern California has changed over the last 39 years, mostly in ways that make it less appealing. By world standards, Southern California's environmental problems are relatively mild. Jokes of East Coast Americans to the contrary, this is not an area at imminent risk of a societal collapse. By world standards and even by U.S. standards, its human population is exceptionally rich and environmentally educated. Los Angeles is well known for some problems, especially its smog, but most of its environmental and population problems are modest or typical compared to

those of other leading First World cities. How do those problems affect the lives of my fellow Angelenos and me?

The complaints voiced by virtually everybody in Los Angeles are those directly related to our growing and already high population: our incurable traffic jams; the very high price of housing (Plate 36), as a result of millions of people working in a few centers of employment, and only limited residential space near those centers; and, as a consequence, the long distances, of up to two hours and 60 miles one way, over which people commute daily in their cars between home and work. Los Angeles became the U.S. city with the worst traffic in 1987 and has remained so every year since then. Everyone recognizes that these problems have gotten worse within the last decade. They are now the biggest single factor hurting the ability of Los Angeles employers to attract and retain employees, and they affect our willingness to drive to events and to visit friends. For the 12-mile trip from my home to downtown Los Angeles or its airport, I now allow an hour and 15 minutes. The average Angeleno spends 368 hours per year, or the equivalent of fifteen 24-hour days, commuting to and from work, without considering time spent driving for other purposes (Plate 37).

No cure is even under serious discussion for these problems, which will only get worse. Such highway construction as is now proposed or under way aims only at smoothing a few of the tightest points of congestion and will be overwhelmed by the increasing number of cars. There is no end in sight to how much worse Los Angeles's problems of congestion will become, because millions of people put up with far worse traffic in other cities. For example, my friends in Bangkok, the capital of Thailand, now carry a portable small chemical toilet in their car because travel can be so prolonged and slow; they once set off to go out of town on a holiday weekend but gave up and returned home after 17 hours, when they had advanced only three miles through the traffic jam. While there are optimists who explain in the abstract why increased population will be good and how the world can accommodate it, I have never met an Angeleno (and very few people anywhere in the world) who personally expressed a desire for increased population in the area where he or she personally lived.

The contribution of Southern California to the ongoing increase in the world's average per-capita human impact, as a result of transfers of people from the Third World to the First World, has for years been the most explosive issue in California politics. California's population growth is accelerating, due almost entirely to immigration and to the large average family sizes of the immigrants after their arrival. The border between California and

Mexico is long and impossible to patrol effectively against people from Central America seeking to immigrate here illegally in search of jobs and personal safety. Every month, one reads of would-be immigrants dying in the desert or being robbed or shot, but that does not deter them. Other illegal immigrants come from as far away as China and Central Asia, in ships that unload them just off the coast. California residents are of two minds about all those Third World immigrants seeking to come here to attain the First World lifestyle. On the one hand, our economy is utterly dependent on them to fill jobs in the service and construction industries and on farms. On the other hand, California residents complain that the immigrants compete with unemployed residents for many jobs, depress wages, and burden our already overcrowded hospitals and public education system. A measure (Proposition 187) on the 1994 state election ballot, overwhelmingly approved by voters but then gutted by the courts on constitutional grounds, would have deprived illegal immigrants of most state-funded benefits. No California resident or elected official has suggested a practical solution to the long-standing contradiction, reminiscent of Dominicans' attitude towards Haitians, between needing immigrants as workers and otherwise resenting their presence and their own needs.

Southern California is a leading contributor to the energy crisis. Our city's former network of electric streetcars collapsed in bankruptcies in the 1920s and 1930s, and the rights of way were bought up by automobile manufacturers and subdivided so as to make it impossible to rebuild the network (which competed with automobiles). Angelenos' preference for living in houses rather than in high-rise apartments, and the long distances and diverse routes over which employees working in any given district commute, have made it impossible to design systems of public transportation that would satisfy the needs of most residents. Hence Los Angelenos are dependent on motorcars.

Our high gas consumption, the mountains ringing much of the Los Angeles basin, and prevailing wind directions generate the smog problem that is our city's most notorious drawback (Plate 38). Despite progress in combating smog in recent decades, and despite seasonal variation (smog worst in the late summer and early autumn) and local variation (smog generally worse as one precedes inland), Los Angeles on the average continues to rank near the bottom of American cities for air quality. After years of improvement, our air quality has again been deteriorating in recent years. Another toxic problem that affects lifestyle and health is the spread of the disease-causing organism giardia in California's rivers and lakes over the last several

decades. When I first moved here in the 1960s and went hiking in the mountains, it was safe to drink water from streams; today the guaranteed result would be giardia infection.

The problem of habitat management of which we are most conscious is the fire risk in Southern California's two predominant habitats, chaparral (a scrub woodland similar to the *machhia* of the Mediterranean) and oak woodland. Under natural conditions both habitats experienced occasional fires from lightning strikes, like the situation in Montana forests that I discussed in Chapter 1. Now that people are living in and next to those highly flammable habitats, Angelinos demand that fires be suppressed immediately. Each year, the late summer and early fall, which are the hottest and driest and windiest time of year in Southern California, are the fire season, when somewhere or other hundreds of homes will go up in flames. The canyon in which I live has not had a fire get out of control since 1961, when there was a big fire that burned 600 houses. A theoretical solution to this problem, as in Montana forests, might be frequent controlled small-scale fires to reduce the fuel load, but such fires would be absurdly dangerous in this densely populated urban area, and the public would not stand for it.

Introduced alien species are a big threat and economic burden to California agriculture, the current leading threat being the Mediterranean fruit fly. Non-agricultural threats are introduced pathogens threatening to kill our oak trees and pine trees. Because one of my two sons became interested as a child in amphibians (frogs and salamanders), I have learned that most species of native amphibians have been exterminated from two-thirds of the streams in Los Angeles County, as the result of the spread of three alien predators on amphibians (a crayfish, bullfrog, and mosquitofish) against which Southern California amphibians are helpless because they never evolved to avoid those threats.

The major soil problem affecting California agriculture is salinization as a result of irrigation agriculture, ruining expanses of agricultural land in California's Central Valley, the richest farmland in the United States.

Because rainfall is low in Southern California, Los Angeles depends for its water on long aqueducts, principally from the Sierra Nevada mountain range and adjacent valleys of Northern California, and from the Colorado River on the eastern border of our state. With the growth of California's population, there has been increasing competition for those water supplies among farmers and cities. With global warming, the Sierra snowpack that provides most of our water will decrease, just as in Montana, increasing the likelihood of water shortages in Los Angeles.

As for collapses of fisheries, the sardine fishery of Northern California collapsed early in the 20th century, the abalone industry of Southern California collapsed a few decades ago soon after my arrival, and the rockfish fishery of Southern California is now collapsing and has become subject to severe restrictions or closure within the last year. Fish prices in Los Angeles supermarkets have increased by a factor of 4 since I moved here.

Finally, losses of biodiversity have affected Southern California's most distinctive species. The symbol of the state of California, and of my university (the University of California), is the California Golden Bear, but it is now extinct. (What dreadful symbolism for one's state and university!) Southern California's population of sea otters was exterminated in the last century, and the outcome of recent attempts at reintroduction is uncertain. Within the time that I've lived in Los Angeles, populations of two of our most characteristic bird species, the Roadrunner and the California Quail, have crashed. Southern California amphibians whose numbers have plummeted are the California Newt and the California Tree Frog.

Thus, environmental and population problems have been undermining the economy and the quality of life in Southern California. They are in large measure ultimately responsible for our water shortages, power shortages, garbage accumulation, school crowding, housing shortages and price rises, and traffic congestion. In most of these respects except for our especially bad traffic jams and air quality, we are no worse off than many other areas of the United States.

Most environmental problems involve detailed uncertainties that are legitimate subjects for debate. In addition, however, there are many reasons that are commonly advanced to dismiss the importance of environmental problems, and that are in my opinion not well informed. These objections are often posed in the form of simplistic "one-liners." Here are a dozen of the commonest ones:

"The environment has to be balanced against the economy." This quote portrays environmental concerns as a luxury, views measures to solve environmental problems as incurring a net cost, and considers leaving environmental problems unsolved to be a money-saving device. This one-liner puts the truth exactly backwards. Environmental messes cost us huge sums of money both in the short run and in the long run; cleaning up or preventing those messes saves us huge sums in the long run, and often in the short run as well. In caring for the health of our surroundings, just as of our bodies, it

is cheaper and preferable to avoid getting sick than to try to cure illnesses after they have developed. Just think of the damage caused by agricultural weeds and pests, non-agricultural pests like water hyacinths and zebra mussels, the recurrent annual costs of combating those pests, the value of lost time when we are stuck in traffic, the financial costs resulting from people getting sick or dying from environmental toxins, cleanup costs for toxic chemicals, the steep increase in fish prices due to depletion of fish stocks, and the value of farmland damaged or ruined by erosion and salinization. It adds up to a few hundred million dollars per year here, tens of billions of dollars there, another billion dollars over here, and so on for hundreds of different problems. For instance, the value of "one statistical life" in the U.S.—i.e., the cost to the U.S. economy resulting from the death of an average American whom society has gone to the expense of rearing and educating but who dies before a lifetime of contributing to the national economy—is usually estimated at around \$5 million. Even if one takes the conservative estimate of annual U.S. deaths due to air pollution as 130,000, then deaths due to air pollution cost us about \$650 billion per year. That illustrates why the U.S. Clean Air Act of 1970, although its cleanup measures do cost money, has yielded estimated net health savings (benefits in excess of costs) of about \$1 trillion per year, due to saved lives and reduced health costs.

"Technology will solve our problems." This is an expression of faith about the future, and therefore based on a supposed track record of technology having solved more problems than it created in the recent past. Underlying this expression of faith is the implicit assumption that, from tomorrow onwards, technology will function primarily to solve existing problems and will cease to create new problems. Those with such faith also assume that the new technologies now under discussion will succeed, and that they will do so quickly enough to make a big difference soon. In extended conversations that I had with two of America's most successful and best-known businessmen and financiers, both of them eloquently described to me emerging technologies and financial instruments that differ fundamentally from those of the past and that, they confidently predicted, would solve our environmental problems.

But actual experience is the opposite of this assumed track record. Some dreamed-of new technologies succeed, while others don't. Those that do succeed typically take a few decades to develop and phase in widely: think of gas heating, electric lighting, cars and airplanes, television, computers, and

so on. New technologies, whether or not they succeed in solving the problem that they were designed to solve, regularly create unanticipated new problems. Technological solutions to environmental problems are routinely far more expensive than preventive measures to avoid creating the problem in the first place: for example, the billions of dollars of damages and cleanup costs associated with major oil spills, compared to the modest cost of safety measures effective at minimizing the risks of a major oil spill.

Most of all, advances in technology just increase our ability to do things, which may be either for the better or for the worse. All of our current problems are unintended negative consequences of our existing technology. The rapid advances in technology during the 20th century have been creating difficult new problems faster than they have been solving old problems: that's why we're in the situation in which we now find ourselves. What makes you think that, as of January 1, 2006, for the first time in human history, technology will miraculously stop causing new unanticipated problems while it just solves the problems that it previously produced?

From thousands of examples of unforeseen harmful side effects of new technological solutions, two must suffice: CFCs (chlorofluorocarbons) and motor vehicles. The coolant gases formerly used in refrigerators and air conditioners were toxic ones (like ammonia) that could prove fatal if those appliances leaked while the homeowner was asleep at night. Hence it was hailed as a great advance when CFCs (alias freons) were developed as synthetic refrigerant gases. They are odorless, non-toxic, and highly stable under ordinary conditions at the Earth's surface, so that initially no bad side effects were observed or expected. Within a short time they became viewed as miracle substances and adopted throughout the world as refrigerator and air-conditioner coolants, foam-blowing agents, solvents, and propellants in aerosol cans. But in 1974 it was discovered that in the stratosphere they are broken down by intense ultraviolet radiation to yield highly reactive chlorine atoms that destroy a significant fraction of the ozone layer protecting us and all other living things against lethal ultraviolet effects. That discovery provoked vigorous denial by some corporate interests, fueled not only by the \$200 billion value of CFC-based industrial efforts but also by genuine doubts because of scientific complications involved. Hence the phasing-out of CFCs has taken a long time: not until 1988 did the DuPont Company (the largest manufacturer of CFCs) decide to stop manufacturing them, in 1992 industrialized countries agreed to cease CFC production by 1995, and China and some other developing countries are still producing them.

Unfortunately, the amounts of CFCs already in the atmosphere are sufficiently large, and their breakdown sufficiently slow, that they will continue to be present for many decades after the eventual end of all CFC production.

The other example involves the introduction of the motor vehicle. When I was a child in the 1940s, some of my teachers were old enough to remember the first decades of the 20th century, when motor vehicles were in the process of replacing horse-drawn carriages and trams on city streets of the United States. The two biggest immediate consequences experienced by urban Americans, my teachers recall, were that American cities became wonderfully cleaner and quieter. No longer were streets constantly polluted with horse manure and urine, and no longer was there the constant din of horse hoofs clicking on the pavement. Today, after a century's experience of cars and buses, it strikes us as ludicrous or inconceivable that anyone could praise them for being non-polluting and quiet. While no one is advocating a return to the horse as a solution to smog from engine emissions, the example does serve to illustrate the unanticipated negative side effects even of technologies that (unlike CFCs) we choose to retain.

"If we exhaust one resource, we can always switch to some other resource meeting the same need." Optimists who make such claims ignore the unforeseen difficulties and long transition times regularly involved. For instance, one area in which switching based on not-yet-perfected new technologies has repeatedly been touted as promising to solve a major environmental problem is automobiles. The current hope for breakthrough involves hydrogen cars and fuel cells, which are technologically in their infancy as applied to motor transport. Thus, there is not a track record justifying faith in the hydrogen-car solution to our fossil fuel problem. However, we do have a track record of a long series of other proposed new car technologies touted as breakthroughs, such as rotary engines and (most recently) electric cars, that aroused much discussion and even sales of production models, only to decline or disappear because of unforeseen problems.

Equally instructive is the automobile industry's recent development of fuel-efficient hybrid gas/electric cars, which have been enjoying increasing sales. However, it would be unfair for a believer in switching to mention hybrid cars without also mentioning the automobile industry's simultaneous development of SUVs, which have been outselling hybrids by a big margin and more than offsetting their fuel savings. The net result of these two technological breakthroughs has been that the fuel consumption and exhaust production of our national car fleet has been going up rather than down.

Nobody has figured out a method to ensure that technology will yield only increasingly environment-friendly effects and products (e.g., hybrid cars), without also yielding environment-unfriendly effects and products (e.g., SUVs).

Another example of faith in switching and substitution is the hope that renewable energy sources, such as wind and solar energy, may solve the energy crisis. These technologies do indeed exist; many Californians now use solar energy to heat their swimming pools, and wind generators are already supplying about one-sixth of Denmark's energy needs. However, wind and solar energy have limited applicability because they can be used only at locations with reliable winds or sunlight. In addition, the recent history of technology shows that conversion times for adoption of major switches—e.g., from candles to oil lamps to gas lamps to electric lights for lighting, or from wood to coal to petroleum for energy—require several decades, because so many institutions and secondary technologies associated with the former technology have to be changed. It is indeed likely that energy sources other than fossil fuels will make increasing contributions to our motor transport and energy generation, but this is a long-term prospect. We'll also need to solve our fuel and energy problems for the next several decades, before new technologies become widespread. All too often, a focus by politicians or industries on the promise of hydrogen cars and wind energy for the distant future distracts attention from all the obvious measures needed right now to decrease driving and fuel consumption by existing cars, and to decrease consumption by fossil fuel generating plants.

"There really isn't a world food problem; there is already enough food; we only need to solve the transportation problem of distributing that food to places that need it." (The same thing could be said for energy.) Or else: *"The world's food problem is already being solved by the Green Revolution, with its new high-yield varieties of rice and other crops, or else it will be solved by genetically modified crops."* This argument notes two things: that First World citizens enjoy on the average greater per-capita food consumption than do Third World citizens; and that some First World countries, such as the United States, do or can produce more food than their citizens consume. If food consumption could be equalized over the world, or if surplus First World food could be exported to the Third World, might that alleviate Third World starvation?

The obvious flaw in the first half of this argument is that First World citizens show no interest in eating less, in order that Third World citizens could eat more. The flaw in the second half of the argument is that, while

First World countries are willing occasionally to export food to mitigate starvation occasioned by some crisis (such as a drought or war) in certain Third World countries, First World citizens have shown no interest in paying on a regular basis (via their tax dollars that support foreign aid and subsidies to farmers) to feed billions of Third World citizens on a chronic basis. If that did happen but without effective overseas family planning programs, which the U.S. government currently opposes on principle, the result would just be Malthus's dilemma, i.e., an increase in population proportional to an increase in available food. Population increase and Malthus's dilemma also contribute to explaining why, after decades of hope and money invested in the Green Revolution and high-yield varieties, starvation is still widespread in the world. All of these considerations mean that genetically modified (GM) food varieties by themselves are equally unlikely to solve the world's food problems (while world population supposedly remains stationary?). In addition, virtually all GM crop production at present is of just four crops (soybeans, corn, canola, and cotton) not eaten directly by humans but used for animal fodder, oil, or clothing, and grown in six temperate-zone countries or regions. Reasons are the strong consumer resistance to eating GM foods; and the cruel fact that companies developing GM crops can make money by selling their products to rich farmers in mostly affluent temperate-zone countries, but not by selling to poor farmers in developing tropical countries. Hence the companies have no interest in investing heavily to develop GM cassava, millet, or sorghum for Third World farmers.

"As measured by commonsense indicators such as human lifespan, health, and wealth (in economists' terms, per-capita gross national product or GNP), conditions have actually been getting better for many decades." Or: "Just look around you: the grass is still green, there is plenty of food in the supermarkets, clean water still flows from the taps, and there is absolutely no sign of imminent collapse." For affluent First World citizens, conditions have indeed been getting better, and public health measures have on the average lengthened lifespans in the Third World as well. But lifespan alone is not a sufficient indicator: billions of Third World citizens, constituting about 80% of the world's population, still live in poverty, near or below the starvation level. Even in the United States, an increasing fraction of the population is at the poverty level and lacks affordable medical care, and all proposals to change this situation (e.g., "just provide everyone with health insurance paid by the government") have been politically unacceptable.

In addition, all of us know as individuals that we don't measure our economic well-being just by the present size of our bank accounts: we also look

at our *direction* of cash flow. When you look at your bank statement and you see a positive \$5,000 balance, you don't smile if you then realize that you have been experiencing a net cash drain of \$200 per month for the last several years, and at that rate you have just two years and one month left before you have to file for bankruptcy. The same principle holds for our national economy, and for environmental and population trends. The prosperity that the First World enjoys at present is based on spending down its environmental capital in the bank (its capital non-renewable energy sources, fish stocks, topsoil, forests, etc.). Spending capital should not be misrepresented as making money. It makes no sense to be content with our present comfort when it is clear that we are currently on a non-sustainable course.

In fact, one of the main lessons to be learned from the collapses of the Maya, Anasazi, Easter Islanders, and those other past societies (as well as from the recent collapse of the Soviet Union) is that a society's steep decline may begin only a decade or two after the society reaches its peak numbers, wealth, and power. In that respect, the trajectories of the societies that we have discussed are unlike the usual courses of individual human lives, which decline in a prolonged senescence. The reason is simple: maximum population, wealth, resource consumption, and waste production mean maximum environmental impact, approaching the limit where impact outstrips resources. On reflection, it's no surprise that declines of societies tend to follow swiftly on their peaks.

"Look at how many times in the past the gloom-and-doom predictions of fearmongering environmentalists have proved wrong. Why should we believe them this time?" Yes, some predictions by environmentalists have proved incorrect, favorite examples of critics being a prediction made in 1980 by Paul Ehrlich, John Harte, and John Holdren about rises in prices of five metals, and predictions made in the Club of Rome forecast of 1972. But it is misleading to look selectively for environmentalist predictions that proved wrong, and not also to look for environmentalist predictions that proved right, or anti-environmentalist predictions that proved wrong. There is an abundance of errors of the latter sort: e.g., overly optimistic predictions that the Green Revolution would already have solved the world's hunger problems; the prediction of the economist Julian Simon that we could feed the world's population as it continues to grow for the next 7 billion years; and Simon's prediction "Copper can be made from other elements" and thus there is no risk of a copper shortage. As regards the first of Simon's two predictions, continuation of our current population growth rate would yield

10 people per square yard of land in 774 years, a mass of people equal to the Earth's mass in slightly under 2,000 years, and a mass of people equal to the universe's mass in 6,000 years, long before Simon's forecast of 7 billion years without such problems. As regards his second prediction, we learn in our first course of chemistry that copper is an element, which means that by definition it cannot be made from other elements. My impression is that pessimistic predictions that have proved incorrect, such as Ehrlich's, Harte's, and Holdren's about metal prices or the Club of Rome's about future food supplies, have on the average been much more realistic possibilities at the time that they were made than were Simon's two predictions.

Basically, the one-liner about some environmentalist predictions proving wrong boils down to a complaint about false alarms. In other spheres of our lives, such as fires, we adopt a commonsense attitude towards false alarms. Our local governments maintain expensive firefighting forces, even though in some small towns they are rarely called on to put out fires. Of the fire alarms phoned in to fire departments, many prove to be false alarms, and many others involve small fires that the property owner himself then succeeds in putting out before the fire engines arrive. We comfortably accept a certain frequency of such false alarms and extinguished fires, because we understand that fire risks are uncertain and hard to judge when a fire has just started, and that a fire that does rage out of control may exact high costs in property and human lives. No sensible person would dream of abolishing the town fire department, whether manned by full-time professionals or volunteers, just because a few years went by without a big fire. Nor would anyone blame a homeowner for calling the fire department on detecting a small fire, only to succeed in quenching the fire before the fire truck's arrival. Only if false alarms become an inordinately high proportion of all fire alarms do we feel that something is wrong. In effect, the proportion of false alarms that we tolerate is based on subconsciously comparing the frequency and destructive costs of big fires with the frequency and wasted-services costs of false alarms. A very low frequency of false alarms proves that too many homeowners are being too cautious, waiting too long to call the fire department, and consequently losing their homes.

By the same reasoning, we must expect some environmentalist warnings to turn out to be false alarms, otherwise we would know that our environmental warning systems were much too conservative. The multibillion-dollar costs of many environmental problems justify a moderate frequency of false alarms. In addition, the reason that alarms proved false is often that they convinced us to adopt successful countermeasures. For example, it's

true that our air quality here in Los Angeles today is not as bad as some gloom-and-doom predictions of 50 years ago. However, that's entirely because Los Angeles and the state of California were thereby aroused to adopt many countermeasures (such as vehicle emission standards, smog certificates, and lead-free gas), not because initial predictions of the problem were exaggerated.

"The population crisis is already solving itself, because the rate of increase of the world's population is decreasing, such that world population will level off at less than double its present level." While the prediction that world population will level off at less than double its present level may or may not prove true, it is at present a realistic possibility. However, we can take no comfort in this possibility, for two reasons: by many criteria, even the world's present population is living at a non-sustainable level; and, as explained earlier in this chapter, the larger danger that we face is not just of a two-fold increase in population, but of a much larger increase in human impact if the Third World's population succeeds in attaining a First World living standard. It is surprising to hear some First World citizens nonchalantly mentioning the world's adding "only" 2½ billion more people (the lowest estimate that anyone would forecast) as if that were acceptable, when the world already holds that many people who are malnourished and living on less than \$3 per day.

"The world can accommodate human population growth indefinitely. The more people, the better, because more people mean more inventions and ultimately more wealth." Both of these ideas are associated especially with Julian Simon but have been espoused by many others, especially by economists. The statement about our ability to absorb current rates of population growth indefinitely is not to be taken seriously, because we have already seen that that would mean 10 people per square yard in the year 2779. Data on national wealth demonstrate that the claim that more people mean more wealth is the opposite of correct. The 10 countries with the most people (over 100 million each) are, in descending order of population, China, India, the U.S., Indonesia, Brazil, Pakistan, Russia, Japan, Bangladesh, and Nigeria. The 10 countries with the highest affluence (per-capita real GDP) are, in descending order, Luxembourg, Norway, the U.S., Switzerland, Denmark, Iceland, Austria, Canada, Ireland, and the Netherlands. The only country on both lists is the U.S.

Actually, the countries with large populations are disproportionately poor: eight of the 10 have per-capita GDP under \$8,000, and five of them under \$3,000. The affluent countries have disproportionately few people: seven of the 10 have populations below 9,000,000, and two of them under

500,000. Instead, what does distinguish the two lists is population growth rates: all 10 of the affluent countries have very low relative population growth rates (1% per year or less), while eight of the 10 most populous countries have higher relative population growth rates than any of the most affluent countries, except for two large countries that achieved low population growth in unpleasant ways: China, by government order and enforced abortion, and Russia, whose population is actually decreasing because of catastrophic health problems. Thus, as an empirical fact, more people and a higher population growth rate mean more poverty, not more wealth.

"Environmental concerns are a luxury affordable just by affluent First World yuppies, who have no business telling desperate Third World citizens what they should be doing." This view is one that I have heard mainly from affluent First World yuppies lacking experience of the Third World. In all my experience of Indonesia, Papua New Guinea, East Africa, Peru, and other Third World countries with growing environmental problems and populations, I have been impressed that their people know very well how they are being harmed by population growth, deforestation, overfishing, and other problems. They know it because they immediately pay the penalty, in forms such as loss of free timber for their houses, massive soil erosion, and (the tragic complaint that I hear incessantly) their inability to afford clothes, books, and school fees for their children. The reason why the forest behind their village is nevertheless being logged is usually either that a corrupt government has ordered it logged over their often-violent protest, or else that they signed a logging lease with great reluctance because they saw no other way to get the money needed next year for their children. My best friends in the Third World, with families of 4 to 8 children, lament that they have heard of the benign forms of contraception widespread in the First World, and they want those measures desperately for themselves, but they can't afford or obtain them, due in part to the refusal of the U.S. government to fund family planning in its foreign aid programs.

Another view that is widespread among affluent First World people, but which they will rarely express openly, is that they themselves are managing just fine at carrying on with their lifestyles despite all those environmental problems, which really don't concern them because the problems fall mainly on Third World people (though it is not politically correct to be so blunt). Actually, the rich are not immune to environmental problems. CEOs of big First World companies eat food, drink water, breathe the air, and have (or try to conceive) children, like the rest of us. While they can usually avoid problems of water quality by drinking bottled water, they find it much more

difficult to avoid being exposed to the same problems of food and air quality as the rest of us. Living disproportionately high on the food chain, at levels at which toxic substances become concentrated, they are at more rather than less risk of reproductive impairment due to ingestion of or exposure to toxic materials, possibly contributing to their higher infertility rates and the increasing frequency with which they require medical assistance in conceiving. In addition, one of the conclusions that we saw emerging from our discussion of Maya kings, Greenland Norse chieftains, and Easter Island chiefs is that, in the long run, rich people do not secure their own interests and those of their children if they rule over a collapsing society and merely buy themselves the privilege of being the last to starve or die. As for First World society as a whole, its resource consumption accounts for most of the world's total consumption that has given rise to the impacts described at the beginning of this chapter. Our totally unsustainable consumption means that the First World could not continue for long on its present course, even if the Third World didn't exist and weren't trying to catch up to us.

"If those environmental problems become desperate, it will be at some time far off in the future, after I die, and I can't take them seriously." In fact, at current rates most or all of the dozen major sets of environmental problems discussed at the beginning of this chapter will become acute within the lifetime of young adults now alive. Most of us who have children consider the securing of our children's future as the highest priority to which to devote our time and our money. We pay for their education and food and clothes, make wills for them, and buy life insurance for them, all with the goal of helping them to enjoy good lives 50 years from now. It makes no sense for us to do these things for our individual children, while simultaneously doing things undermining the world in which our children will be living 50 years from now.

This paradoxical behavior is one of which I personally was guilty, because I was born in the year 1937, hence before the birth of my children I too could not take seriously any event (like global warming or the end of the tropical rainforests) projected for the year 2037. I shall surely be dead before that year, and even the date 2037 struck me as unreal. However, when my twin sons were born in 1987, and when my wife and I then started going through the usual parental obsessions about schools, life insurance, and wills, I realized with a jolt: 2037 is the year in which my kids will be my own age of 50 (then)! It's not an imaginary year! What's the point of willing our property to our kids if the world will be in a mess then anyway?

Having lived for five years in Europe shortly after World War II, and then having married into a Polish family with a Japanese branch, I saw at first hand what can happen when parents take good care of their individual children but not of their children's future world. The parents of my Polish, German, Japanese, Russian, British, and Yugoslav friends also bought life insurance, made wills, and obsessed about the schooling of their children, as my wife and I have been doing more recently. Some of them were rich and would have had valuable property to will to their children. But they did not take good care of their children's world, and they blundered into the disaster of World War II. As a result, most of my European and Japanese friends born in the same year as I had their lives blighted in various ways, such as being orphaned, separated from one or both parents during their childhood, bombed out of their houses, deprived of schooling opportunities, deprived of their family estates, or raised by parents burdened with memories of war and concentration camps. The worst-case scenarios that today's children face if we too blunder about their world are different, but equally unpleasant.

This leaves us with two other common one-liners that we have not considered: "*There are big differences between modern societies and those past societies of Easter Islanders, Maya, and Anasazi who collapsed, so that we can't straightforwardly apply lessons from the past.*" And: "*What can I, as an individual, do, when the world is really being shaped by unstoppable powerful juggernauts of governments and big businesses?*" In contrast to the previous one-liners, which upon examination can be quickly dismissed, these two concerns are valid and cannot be dismissed. I shall devote the remainder of this chapter to the former question, and a section of the Further Readings (pp. 555-59) to the latter question.

Are the parallels between the past and present sufficiently close that the collapses of the Easter Islanders, Henderson Islanders, Anasazi, Maya, and Greenland Norse could offer any lessons for the modern world? At first, a critic, noting the obvious differences, might be tempted to object, "It's ridiculous to suppose that the collapses of all those ancient peoples could have broad relevance today, especially to the modern U.S. Those ancients didn't enjoy the wonders of modern technology, which benefits us and which lets us solve problems by inventing new environment-friendly technologies. Those ancients had the misfortune to suffer from effects of climate change. They behaved stupidly and ruined their own environment by

doing obviously dumb things, like cutting down their forests, overharvesting wild animal sources of their protein, watching their topsoil erode away, and building cities in dry areas likely to run short of water. They had foolish leaders who didn't have books and so couldn't learn from history, and who embroiled them in expensive and destabilizing wars, cared only about staying in power, and didn't pay attention to problems at home. They got overwhelmed by desperate starving immigrants, as one society after another collapsed, sending floods of economic refugees to tax the resources of the societies that weren't collapsing. In all those respects, we moderns are fundamentally different from those primitive ancients, and there is nothing that we could learn from them. Especially we in the U.S., the richest and most powerful country in the world today, with the most productive environment and wise leaders and strong loyal allies and only weak insignificant enemies—none of those bad things could possibly apply to us."

Yes, it's true that there are big differences between the situations of those past societies and our modern situation today. The most obvious difference is that there are far more people alive today, packing far more potent technology that impacts the environment, than in the past. Today we have over 6 billion people equipped with heavy metal machinery such as bulldozers and nuclear power, whereas the Easter Islanders had at most a few tens of thousands of people with stone chisels and human muscle power. Yet the Easter Islanders still managed to devastate their environment and bring their society to the point of collapse. That difference greatly increases, rather than decreases, the risks for us today.

A second big difference stems from globalization. Leaving out of this discussion for the moment the question of environmental problems within the First World itself, let's just ask whether the lessons from past collapses might apply anywhere in the Third World today. First ask some ivory-tower academic ecologist, who knows a lot about the environment but never reads a newspaper and has no interest in politics, to name the overseas countries facing some of the worst problems of environmental stress, overpopulation, or both. The ecologist would answer: "That's a no-brainer, it's obvious. Your list of environmentally stressed or overpopulated countries should surely include Afghanistan, Bangladesh, Burundi, Haiti, Indonesia, Iraq, Madagascar, Mongolia, Nepal, Pakistan, the Philippines, Rwanda, the Solomon Islands, and Somalia, plus others" (map, p. 497).

Then go ask a First World politician, who knows nothing and cares less about the environment and population problems, to name the world's worst trouble spots: countries where state government has already been

overwhelmed and has collapsed, or is now at risk of collapsing, or has been wracked by recent civil wars; and countries that, as a result of those problems of their own, are also creating problems for us rich First World countries, which may end up having to provide foreign aid for them, or may face illegal immigrants from them, or may decide to provide them with military assistance to deal with rebellions and terrorists, or may even have to send in our own troops. The politician would answer, "That's a no-brainer, it's obvious. Your list of political trouble spots should surely include Afghanistan, Bangladesh, Burundi, Haiti, Indonesia, Iraq, Madagascar, Mongolia, Nepal, Pakistan, the Philippines, Rwanda, the Solomon Islands, and Somalia, plus others."

Surprise, surprise: the two lists are very similar. The connection between the two lists is transparent: it's the problems of the ancient Maya, Anasazi, and Easter Islanders playing out in the modern world. Today, just as in the past, countries that are environmentally stressed, overpopulated, or both become at risk of getting politically stressed, and of their governments collapsing. When people are desperate, undernourished, and without hope, they blame their governments, which they see as responsible for or unable to solve their problems. They try to emigrate at any cost. They fight each other over land. They kill each other. They start civil wars. They figure that they have nothing to lose, so they become terrorists, or they support or tolerate terrorism.

The results of these transparent connections are genocides such as the ones that already exploded in Bangladesh, Burundi, Indonesia, and Rwanda; civil wars or revolutions, as in most of the countries on the lists; calls for the dispatch of First World troops, as to Afghanistan, Haiti, Indonesia, Iraq, the Philippines, Rwanda, the Solomon Islands, and Somalia; the collapse of central government, as has already happened in Somalia and the Solomon Islands; and overwhelming poverty, as in all of the countries on these lists. Hence the best predictors of modern "state failures"—i.e., revolutions, violent regime change, collapse of authority, and genocide—prove to be measures of environmental and population pressure, such as high infant mortality, rapid population growth, a high percentage of the population in their late teens and 20s, and hordes of unemployed young men without job prospects and ripe for recruitment into militias. Those pressures create conflicts over shortages of land (as in Rwanda), water, forests, fish, oil, and minerals. They create not only chronic internal conflict, but also emigration of political and economic refugees, and wars between coun-

tries arising when authoritarian regimes attack neighboring nations in order to divert popular attention from internal stresses.

In short, it is not a question open for debate whether the collapses of past societies have modern parallels and offer any lessons to us. That question is settled, because such collapses have actually been happening recently, and others appear to be imminent. Instead, the real question is how many more countries will undergo them.

As for terrorists, you might object that many of the political murderers, suicide bombers, and 9/11 terrorists were educated and moneyed rather than uneducated and desperate. That's true, but they still depended on a desperate society for support and toleration. Any society has its murderous fanatics; the U.S. produced its own Timothy McVeigh and its Harvard-educated Theodore Kaczinski. But well-nourished societies offering good job prospects, like the U.S., Finland, and South Korea, don't offer broad support to their fanatics.

The problems of all these environmentally devastated, overpopulated, distant countries become our own problems because of globalization. We are accustomed to thinking of globalization in terms of us rich advanced First Worlders sending our good things, such as the Internet and Coca-Cola, to those poor backward Third Worlders. But globalization means nothing more than improved worldwide communications, which can convey many things in either direction; globalization is not restricted to good things carried only from the First to the Third World.

Among bad things transported from the First World to developing countries, we already mentioned the millions of tons of electronic garbage intentionally transported each year from industrialized nations to China. To grasp the worldwide scale of unintentional garbage transport, consider the garbage collected on the beaches of tiny Oeno and Ducie Atolls in the Southeast Pacific Ocean (see map on p. 122): uninhabited atolls, without freshwater, rarely visited even by yachts, and among the world's most remote bits of land, each over a hundred miles even from remote uninhabited Henderson Island. Surveys there detected, for each linear yard of beach, on the average one piece of garbage, which must have drifted from ships or else from Asian and American countries on the Pacific Rim thousands of miles distant. The commonest items proved to be plastic bags, buoys, glass and plastic bottles (especially Suntory whiskey bottles from Japan), rope, shoes, and lightbulbs, along with oddities such as footballs, toy soldiers and airplanes, bike pedals, and screwdrivers.

A more sinister example of bad things transported from the First World to developing countries is that the highest blood levels of toxic industrial chemicals and pesticides reported for any people in the world are for Eastern Greenland's and Siberia's Inuit people (Eskimos), who are also among the most remote from sites of chemical manufacture or heavy use. Their blood mercury levels are nevertheless in the range associated with acute mercury poisoning, while the levels of toxic PCBs (polychlorinated biphenyls) in Inuit mothers' breast milk fall in a range high enough to classify the milk as "hazardous waste." Effects on the women's babies include hearing loss, altered brain development, and suppressed immune function, hence high rates of ear and respiratory infections.

Why should levels of these poisonous chemicals from remote industrial nations of the Americas and Europe be higher in the Inuit than even in urban Americans and Europeans? It's because staples of the Inuit diet are whales, seals, and seabirds that eat fish, molluscs, and shrimp, and the chemicals become concentrated at each step as they pass up this food chain. All of us in the First World who occasionally consume seafood are also ingesting these chemicals, but in smaller amounts. (However, that doesn't mean that you will be safe if you stop eating seafood, because you now can't avoid ingesting such chemicals no matter what you eat.)

Still other bad impacts of the First World on the Third World include deforestation, Japan's imports of wood products currently being a leading cause of deforestation in the tropical Third World; and overfishing, due to fishing fleets of Japan, Korea, Taiwan and the heavily subsidized fleets of the European Union scouring the world's oceans. Conversely, people in the Third World can now, intentionally or unintentionally, send us their own bad things: their diseases like AIDS, SARS, cholera, and West Nile fever, carried inadvertently by passengers on transcontinental airplanes; unstoppable numbers of legal and illegal immigrants arriving by boat, truck, train, plane, and on foot; terrorists; and other consequences of their Third World problems. We in the U.S. are no longer the isolated Fortress America to which some of us aspired in the 1930s; instead, we are tightly and irreversibly connected to overseas countries. The U.S. is the world's leading importer nation: we import many necessities (especially oil and some rare metals) and many consumer products (cars and consumer electronics), as well as being the world's leading importer of investment capital. We are also the world's leading exporter, particularly of food and of our own manufactured products. Our own society opted long ago to become interlocked with the rest of the world.

That's why political instability anywhere in the world now affects us, our trade routes, and our overseas markets and suppliers. We are so dependent on the rest of the world that if, 30 years ago, you had asked a politician to name the countries most geopolitically irrelevant to our interests because of their being so remote, poor, and weak, the list would surely have begun with Afghanistan and Somalia, yet they subsequently became recognized as important enough to warrant our dispatching U.S. troops. Today the world no longer faces just the circumscribed risk of an Easter Island society or Maya homeland collapsing in isolation, without affecting the rest of the world. Instead, societies today are so interconnected that the risk we face is of a worldwide decline. That conclusion is familiar to any investor in stock markets: instability of the U.S. stock market, or the post-9/11 economic downturn in the U.S., affects overseas stock markets and economies as well, and vice versa. We in the U.S. (or else just affluent people in the U.S.) can no longer get away with advancing our own self-interests, at the expense of the interests of others.

A good example of a society minimizing such clashes of interest is the Netherlands, whose citizens have perhaps the world's highest level of environmental awareness and of membership in environmental organizations. I never understood why, until on a recent trip to the Netherlands I posed the question to three of my Dutch friends while driving through their countryside (Plates 39, 40). Their answer was one that I shall never forget:

"Just look around you here. All of this farmland that you see lies below sea level. One-fifth of the total area of the Netherlands is below sea level, as much as 22 feet below, because it used to be shallow bays, and we reclaimed it from the sea by surrounding the bays with dikes and then gradually pumping out the water. We have a saying, 'God created the Earth, but we Dutch created the Netherlands.' These reclaimed lands are called 'polders.' We began draining them nearly a thousand years ago. Today, we still have to keep pumping out the water that gradually seeps in. That's what our windmills used to be for, to drive the pumps to pump out the polders. Now we use steam, diesel, and electric pumps instead. In each polder there are lines of pumps, starting with those farthest from the sea, pumping the water in sequence until the last pump finally pumps it out into a river or the ocean. In the Netherlands, we have another expression, 'You have to be able to get along with your enemy, because he may be the person operating the neighboring pump in your polder.' And we're all down in the polders together. It's not the case that rich people live safely up on tops of the dikes while poor people live down in the polder bottoms below sea level. If the dikes and

pumps fail, we'll all drown together. When a big storm and high tides swept inland over Zeeland Province on February 1, 1953, nearly 2,000 Dutch people, both rich and poor, drowned. We swore that we would never let that happen again, and the whole country paid for an extremely expensive set of tide barriers. If global warming causes polar ice melting and a world rise in sea level, the consequences will be more severe for the Netherlands than for any other country in the world, because so much of our land is already under sea level. That's why we Dutch are so aware of our environment. We've learned through our history that we're all living in the same polder, and that our survival depends on each other's survival."

That acknowledged interdependence of all segments of Dutch society contrasts with current trends in the United States, where wealthy people increasingly seek to insulate themselves from the rest of society, aspire to create their own separate virtual polders, use their own money to buy services for themselves privately, and vote against taxes that would extend those amenities as public services to everyone else. Those private amenities include living inside gated walled communities (Plate 36), relying on private security guards rather than on the police, sending one's children to well-funded private schools with small classes rather than to the underfunded crowded public schools, purchasing private health insurance or medical care, drinking bottled water instead of municipal water, and (in Southern California) paying to drive on toll roads competing with the jammed public freeways. Underlying such privatization is a misguided belief that the elite can remain unaffected by the problems of society around them: the attitude of those Greenland Norse chiefs who found that they had merely bought themselves the privilege of being the last to starve.

Throughout human history, most peoples have been connected to some other peoples, living together in small virtual polders. The Easter Islanders comprised a dozen clans, dividing their island polder into a dozen territories, and isolated from all other islands, but sharing among clans the Rano Raraku statue quarry, the Puna Pau pukao quarry, and a few obsidian quarries. As Easter Island society disintegrated, all the clans disintegrated together, but nobody else in the world knew about it, nor was anybody else affected. Southeast Polynesia's polder consisted of three interdependent islands, such that the decline of Mangareva's society was disastrous also for the Pitcairn and Henderson Islanders but for no one else. To the ancient Maya, their polder consisted at most of the Yucatán Peninsula and neighboring areas. When the Classic Maya cities collapsed in the southern Yucatán, refugees may have reached the northern Yucatán, but certainly not

Florida. In contrast today our whole world has become one polder, such that events anywhere affect Americans. When distant Somalia collapsed, in went American troops; when the former Yugoslavia and Soviet Union collapsed, out went streams of refugees over all of Europe and the rest of the world; and when changed conditions of society, settlement, and lifestyle spread new diseases in Africa and Asia, those diseases moved over the globe. The whole world today is a self-contained and isolated unit, as Tikopia Island and Tokugawa Japan used to be. We need to realize, as did the Tikopians and Japanese, that there is no other island/other planet to which we can turn for help, or to which we can export our problems. Instead, we need to learn, as they did, to live within our means.

I introduced this section by acknowledging that there are important differences between the ancient world and the modern world. The differences that I then went on to mention—today's larger population and more potent destructive technology, and today's interconnectedness posing the risk of a global rather than a local collapse—may seem to suggest a pessimistic outlook. If the Easter Islanders couldn't solve their milder local problems in the past, how can the modern world hope to solve its big global problems?

People who get depressed at such thoughts often then ask me, "Jared, are you optimistic or pessimistic about the world's future?" I answer, "I'm a cautious optimist." By that, I mean that, on the one hand, I acknowledge the seriousness of the problems facing us. If we don't make a determined effort to solve them, and if we don't succeed at that effort, the world as a whole within the next few decades will face a declining standard of living, or perhaps something worse. That's the reason why I decided to devote most of my career efforts at this stage of my life to convincing people that our problems have to be taken seriously and won't go away otherwise. On the other hand, we shall be able to solve our problems—if we choose to do so. That's why my wife and I did decide to have children 17 years ago: because we did see grounds for hope.

One basis for hope is that, realistically, we are not beset by insoluble problems. While we do face big risks, the most serious ones are not ones beyond our control, like a possible collision with an asteroid of a size that hits the Earth every hundred million years or so. Instead, they are ones that we are generating ourselves. Because we are the cause of our environmental problems, we are the ones in control of them, and we can choose or not choose to stop causing them and start solving them. The future is up for

grabs, lying in our own hands. We don't need new technologies to solve our problems; while new technologies can make some contribution, for the most part we "just" need the political will to apply solutions already available. Of course, that's a big "just." But many societies did find the necessary political will in the past. Our modern societies have already found the will to solve some of our problems, and to achieve partial solutions to others.

Another basis for hope is the increasing diffusion of environmental thinking among the public around the world. While such thinking has been with us for a long time, its spread has accelerated, especially since the 1962 publication of *Silent Spring*. The environmental movement has been gaining adherents at an increasing rate, and they act through a growing diversity of increasingly effective organizations, not only in the United States and Europe but also in the Dominican Republic and other developing countries. At the same time as the environmental movement is gaining strength at an increasing rate, so too are the threats to our environment. That's why I referred earlier in this book to our situation as that of being in an exponentially accelerating horse race of unknown outcome. It's neither impossible, nor is it assured, that our preferred horse will win the race.

What are the choices that we must make if we are now to succeed, and not to fail? There are many specific choices, of which I discuss examples in the Further Readings section, that any of us can make as individuals. For our society as a whole, the past societies that we have examined in this book suggest broader lessons. Two types of choices seem to me to have been crucial in tipping their outcomes towards success or failure: long-term planning, and willingness to reconsider core values. On reflection, we can also recognize the crucial role of these same two choices for the outcomes of our individual lives.

One of those choices has depended on the courage to practice long-term thinking, and to make bold, courageous, anticipatory decisions at a time when problems have become perceptible but before they have reached crisis proportions. This type of decision-making is the opposite of the short-term reactive decision-making that too often characterizes our elected politicians—the thinking that my politically well-connected friend decried as "90-day thinking," i.e., focusing only on issues likely to blow up in a crisis within the next 90 days. Set against the many depressing bad examples of such short-term decision-making are the encouraging examples of courageous long-term thinking in the past, and in the contemporary world of NGOs, business, and government. Among past societies faced with the prospect of ruinous deforestation, Easter Island and Mangareva chiefs

succumbed to their immediate concerns, but Tokugawa shoguns, Inca emperors, New Guinea highlanders, and 16th-century German landowners adopted a long view and reafforested. China's leaders similarly promoted reforestation in recent decades and banned logging of native forests in 1998. Today, many NGOs exist specifically for the purpose of promoting sane long-term environmental policies. In the business world the American corporations that remain successful for long times (e.g., Procter and Gamble) are ones that don't wait for a crisis to force them to reexamine their policies, but that instead look for problems on the horizon and act before there is a crisis. I already mentioned Royal Dutch Shell Oil Company as having an office devoted just to envisioning scenarios decades off in the future.

Courageous, successful, long-term planning also characterizes some governments and some political leaders, some of the time. Over the last 30 years a sustained effort by the U.S. government has reduced levels of the six major air pollutants nationally by 25%, even though our energy consumption and population increased by 40% and our vehicle miles driven increased by 150% during those same decades. The governments of Malaysia, Singapore, Taiwan, and Mauritius all recognized that their long-term economic well-being required big investments in public health to prevent tropical diseases from sapping their economies; those investments proved to be a key to those countries' spectacular recent economic growth. Of the former two halves of the overpopulated nation of Pakistan, the eastern half (independent since 1971 as Bangladesh) adopted effective family planning measures to reduce its rate of population growth, while the western half (still known as Pakistan) did not and is now the world's sixth most populous country. Indonesia's former environmental minister Emil Salim, and the Dominican Republic's former president Joaquín Balaguer, exemplify government leaders whose concern about chronic environmental dangers made a big impact on their countries. All of these examples of courageous long-term thinking in both the public sector and the private sector contribute to my hope.

The other crucial choice illuminated by the past involves the courage to make painful decisions about values. Which of the values that formerly served a society well can continue to be maintained under new changed circumstances? Which of those treasured values must instead be jettisoned and replaced with different approaches? The Greenland Norse refused to jettison part of their identity as a European, Christian, pastoral society, and they died as a result. In contrast, Tikopia Islanders did have the courage to eliminate their ecologically destructive pigs, even though pigs are the sole

large domestic animal and a principal status symbol of Melanesian societies. Australia is now in the process of reappraising its identity as a British agricultural society. The Icelanders and many traditional caste societies of India in the past, and Montana ranchers dependent on irrigation in recent times, did reach agreement to subordinate their individual rights to group interests. They thereby succeeded in managing shared resources and avoiding the tragedy of the commons that has befallen so many other groups. The government of China restricted the traditional freedom of individual reproductive choice, rather than let population problems spiral out of control. The people of Finland, faced with an ultimatum by their vastly more powerful Russian neighbor in 1939, chose to value their freedom over their lives, fought with a courage that astonished the world, and won their gamble, even while losing the war. While I was living in Britain from 1958 to 1962, the British people were coming to terms with the outdatedness of cherished long-held values based on Britain's former role as the world's dominant political, economic, and naval power. The French, Germans, and other European countries have advanced even further in subordinating to the European Union their national sovereignties for which they used to fight so dearly.

All of these past and recent reappraisals of values that I have just mentioned were achieved despite being agonizingly difficult. Hence they also contribute to my hope. They may inspire modern First World citizens with the courage to make the most fundamental reappraisal now facing us: how much of our traditional consumer values and First World living standard can we afford to retain? I already mentioned the seeming political impossibility of inducing First World citizens to lower their impact on the world. But the alternative, of continuing our current impact, is more impossible. This dilemma reminds me of Winston Churchill's response to criticisms of democracy: "It has been said that Democracy is the worst form of government except all those other forms that have been tried from time to time." In that spirit, a lower-impact society is the most impossible scenario for our future—except for all other conceivable scenarios.

Actually, while it won't be easier to reduce our impact, it won't be impossible either. Remember that impact is the product of two factors: population, multiplied times impact per person. As for the first of those two factors, population growth has recently declined drastically in all First World countries, and in many Third World countries as well—including China, Indonesia, and Bangladesh, with the world's largest, fourth largest, and ninth largest populations respectively. Intrinsic population growth in

Japan and Italy is already below the replacement rate, such that their existing populations (i.e., not counting immigrants) will soon begin shrinking. As for impact per person, the world would not even have to decrease its current consumption rates of timber products or of seafood: those rates could be sustained or even increased, if the world's forests and fisheries were properly managed.

My remaining cause for hope is another consequence of the globalized modern world's interconnectedness. Past societies lacked archaeologists and television. While the Easter Islanders were busy deforesting the highlands of their overpopulated island for agricultural plantations in the 1400s, they had no way of knowing that, thousands of miles to the east and west at the same time, Greenland Norse society and the Khmer Empire were simultaneously in terminal decline, while the Anasazi had collapsed a few centuries earlier, Classic Maya society a few more centuries before that, and Mycenaean Greece 2,000 years before that. Today, though, we turn on our television sets or radios or pick up our newspapers, and we see, hear, or read about what happened in Somalia or Afghanistan a few hours earlier. Our television documentaries and books show us in graphic detail why the Easter Islanders, Classic Maya, and other past societies collapsed. Thus, we have the opportunity to learn from the mistakes of distant peoples and past peoples. That's an opportunity that no past society enjoyed to such a degree. My hope in writing this book has been that enough people will choose to profit from that opportunity to make a difference.

War

by

Brian Orend

Stanford Encyclopedia of Philosophy

First published Fri Feb 4, 2000; substantive revision Thu Jul 28, 2005

<http://plato.stanford.edu/entries/war>

War should be understood as an *actual, intentional* and *widespread* armed conflict between political communities. Thus, fisticuffs between individual persons do not count as a war, nor does a gang fight, nor does a feud on the order of the Hatfields versus the McCoys. War is a phenomenon which occurs *only* between political communities, defined as those entities which either are states or intend to become states (in order to allow for civil war). Classical war is international war, a war between different states, like the two World Wars. But just as frequent is war within a state between rival groups or communities, like the American Civil War. Certain political pressure groups, like terrorist organizations, might also be considered “political communities,” in that they are associations of people with a political purpose and, indeed, many of them aspire to statehood or to influence the development of statehood in certain lands.

What's statehood? Most people follow Max Weber's distinction between nation and state. A nation is a group which thinks of itself as “a people,” usually because they share many things in common, such as ethnicity, language, culture, historical experience, a set of ideals and values, habitat, cuisine, fashion and so on. The state, by contrast, refers much more narrowly to the machinery of government which organizes life in a given territory. Thus, we can distinguish between the American state and the American people, or between the government of France and the French nation. At the same time, you've probably heard the term “nation-state.” Indeed, people often use “nation” and “state” interchangeably but we'll need to keep them conceptually distinct for our purposes. “Nation-state” refers to the relatively recent phenomenon wherein a nation wants its own state, and moves to form one. This started out as a very European trend—an Italian state for the Italian nation, a German state for the German people, etc., but it has spread throughout the world. Note that in some countries, such as America, Australia and Canada, the state actually presides over many nations, and you hear of “multi-national societies.” Most societies with heavy immigration are multi-national. Multi-national countries are sometimes prone to civil wars between the different groups. This has been especially true of central Africa in recent years, as different peoples struggle over control of the one state, or else move to separate themselves from the existing arrangement (itself often having been put in place by distant imperial powers insensitive to local group and ethnic differences).

All these distinctions will come in handy as we proceed. For now, we note how central the issue of statehood is to the essence of warfare. Indeed, it seems that *all warfare is precisely, and ultimately, about governance*. War is a violent way for determining who gets to say what goes on in a given territory, for example, regarding: who gets power, who gets wealth and resources, whose ideals prevail, who is a member and who is not, which laws get made, what gets taught in schools, where the border rests, how much tax is levied, and so on. War is the ultimate means for deciding these issues if a peaceful process or resolution can't be agreed upon.

The mere *threat* of war, and the presence of mutual disdain between political communities, do not suffice as indicators of war. The conflict of arms must be *actual*, and not merely latent, for it to count as war. Further, the actual armed conflict must be both *intentional* and *widespread*: isolated clashes between rogue officers, or border patrols, do not count as actions of war. The onset of war requires a conscious commitment, and a significant mobilization, on the part of the belligerents in question. There's no real war, so to speak, until the fighters *intend* to go to war and until they do so with a heavy *quantum* of force.

Let us here cite, by way of support, the views of the one and only (so-called) “philosopher of war,” Carl von Clausewitz. Clausewitz famously suggested that war is “the continuation of policy by other means.” Surely, as a description, this conception is both powerful and plausible: war is about governance, using violence instead of peaceful measures to resolve policy (which organizes life in a land). This notion fits in

nicely with Clausewitz's own general definition of war as "an act of violence intended to compel our opponent to fulfil our will." War, he says, is like a duel, but on "an extensive scale." As Michael Gelven has written more recently, war is intrinsically vast, communal (or political) and violent. *It is an actual, widespread and deliberate armed conflict between political communities, motivated by a sharp disagreement over governance.* In fact, we might say that Clausewitz was right, but not quite deep enough: it's not just that war is the continuation of policy by other means; it's that war is *about the very thing which creates policy*—i.e., governance itself. War is the intentional use of mass force to resolve disputes over governance. War is, indeed, governance by bludgeon. Ultimately, war is profoundly anthropological: it is about which group of people gets to say what goes on in a given territory.

War is a brutal and ugly enterprise. Yet it remains central to human history and social change. These two facts together might seem paradoxical and inexplicable, or they might reveal deeply disturbing facets of the human character (notably, a drive for dominance over others). What is certainly true, in any event, is that war and its threat continue to be forces in our lives. Recent events graphically demonstrate this proposition, whether we think of the 9-11 attacks, the counter-attack on Afghanistan, the overthrow of Iraq's Saddam Hussein, the Darfur crisis in Sudan, the bombings in Madrid and London, or the on-going "war on terror" more generally. We all had high hopes going into the new millennium in 2000; alas, this new century has already been savagely scarred with warfare.

War's violent nature, and controversial social effects, raise troubling moral questions for any thoughtful person. Is war always wrong? Might there be situations when it can be a justified, or even a smart, thing to do? Will war always be part of human experience, or can we do something to make it disappear? Is war an outcome of unchangeable human nature or, rather, of changeable social practice? Is there a fair and sensible way to wage war, or is it all hopeless, barbaric slaughter? When wars end, how should post-war reconstruction proceed, and who should be in charge? What are our rights, and responsibilities, when our own society makes the move to go to war?

1. The Ethics of War and Peace

Three traditions of thought dominate the ethics of war and peace: *Realism*; *Pacifism*; and *Just War Theory* (and, through just war theory, International Law). Perhaps there are other possible perspectives but it seems that very few theories on the ethics of war succeed in resisting ultimate classification into one of these traditions. They are clearly hegemonic in this regard.

Before discussing the central elements of each tradition, let's declare the basic conceptual differences between "the big three" perspectives. The core, and controversial, proposition of just war theory is that, sometimes, states can have moral justification for resorting to armed force. War is sometimes, but of course not all the time, morally right. The idea here is not that the war in question is merely politically shrewd, or prudent, or bold and daring, but fully moral, just. It is an ethically appropriate use of mass political violence. World War II, on the Allied side, is always trotted out as the definitive example of a just and good war. Realism, by contrast, sports a profound skepticism about the application of moral concepts, such as justice, to the key problems of foreign policy. Power and national security, realists claim, motivate states during wartime and thus moral appeals are strictly wishful thinking. Talk of the morality of warfare is pure bunk: ethics has got nothing to do with the rough-and-tumble world of global politics, where only the strong and cunning survive. A country should tend to its vital interests in security, influence over others, and economic growth—and not to moral ideals. Pacifism does not share realism's moral skepticism. For the pacifist, moral concepts can indeed be applied fruitfully to international affairs. It *does* make sense to ask whether a war is just: that is an important and meaningful issue. But the result of such normative application, in the case of war, is always that war should not be undertaken. Where just war theory is *sometimes* permissive with regard to war, pacifism is *always* prohibitive. For the pacifist, war is always wrong; there's always some better resolution to the problem than fighting. Now let's turn to the elements of each of these three traditions.

2. Just War Theory

Just war theory is probably the most influential perspective on the ethics of war and peace. The just war tradition has enjoyed a long and distinguished pedigree, including such notables as Augustine, Aquinas, Grotius, Suarez, Vattel and Vitoria. Hugo Grotius is probably the most comprehensive and formidable classical member of the tradition; James T. Johnson is the authoritative historian of this tradition; and many recognize Michael Walzer as the dean of contemporary just war theorists. Many credit Augustine with the founding of just war theory but this is incomplete. As Johnson notes, in its origins just war theory is a synthesis of classical Greco-Roman, as well as Christian, values. If we have to “name names”, the founders of just war theory are probably the triad of Aristotle, Cicero and Augustine. Many of the rules developed by the just war tradition have since been codified into contemporary international laws governing armed conflict, such as The United Nations Charter and The Hague and Geneva Conventions. The tradition has thus been doubly influential, dominating both moral and legal discourse surrounding war. It sets the tone, and the parameters, for the great debate.

Just war theory can be meaningfully divided into three parts, which in the literature are referred to, for the sake of convenience, in Latin. These parts are: 1) *jus ad bellum*, which concerns the justice of resorting to war in the first place; 2) *jus in bello*, which concerns the justice of conduct within war, after it has begun; and 3) *jus post bellum*, which concerns the justice of peace agreements and the termination phase of war.

2.1 *Jus ad bellum*

The rules of *jus ad bellum* are addressed, first and foremost, to heads of state. Since political leaders are the ones who inaugurate wars, setting their armed forces in motion, they are to be held accountable to *jus ad bellum* principles. If they fail in that responsibility, then they commit war crimes. In the language of the Nuremberg prosecutors, aggressive leaders who launch unjust wars commit “crimes against peace.” What constitutes a just or unjust resort to armed force is disclosed to us by the rules of *jus ad bellum*. Just war theory contends that, for any resort to war to be justified, a political community, or state, must fulfil each and every one of the following six requirements:

1. Just cause. This is clearly the most important rule; it sets the tone for everything which follows. A state may launch a war only for the right reason. The just causes most frequently mentioned include: self-defence from external attack; the defence of others from such; the protection of innocents from brutal, aggressive regimes; and punishment for a grievous wrongdoing which remains uncorrected. Vitoria suggested that all the just causes be subsumed under the one category of “a wrong received.” Walzer, and most modern just war theorists, speak of the one just cause for resorting to war being the resistance of aggression. Aggression is the use of armed force in violation of someone else’s basic rights.

The basic rights of two kinds of entity are involved here: those of states; and those of their individual citizens. International law affirms that states have many rights, notably those to political sovereignty and territorial integrity. It thus affirms that aggression involves the use of armed forces—armies, navies, air forces, marines, missiles—in violation of these rights. Classic cases would be Nazi Germany into Poland in 1939, and Iraq into Kuwait in 1990, wherein the aggressor used its armed forces to invade the territory of the victim, overthrow its government and establish a new regime in its place. Crucially, the commission of aggression causes the aggressor to forfeit its own state rights, thereby permitting violent resistance. *An aggressor has no right not to be warred against in defence*; indeed, it has the duty to stop its rights-violating aggression.

But why do states have rights? The only respectable answer seems to be that they need these rights to protect their people and to help provide them with the objects of their human rights. As John Locke, and the U.S. Founding Fathers, declared: governments are instituted among people to realize the basic rights of those people. If governments do so, they are legitimate; if not, they have neither right nor reason to exist. This is vital: from the moral point of view, *only legitimate governments have rights*, including those to go to war. We need a theory of legitimate governance to ground just war theory, and Aquinas perhaps saw this

more clearly than any classical member of the tradition. This connection to legitimacy is consistent with the perspective on war offered so far: war, at its heart, is a violent clash over how a territory and its people are to be governed.

Based on international law (see Roth), it seems like there are three basic criteria for a legitimate government. If these conditions are met, the state in question has rights to govern and to be left in peace. They are as follows. First, the state is recognized as legitimate by its own people and by the international community. There is an uncoerced general peace and order within that society, and the state is not shunned as a pariah by the rest of the world. Second, the state avoids violating the rights of other legitimate states. In particular, legitimate governments don't commit aggression against other societies. Finally, legitimate states make every reasonable effort to satisfy the human rights of their own citizens, notably those to life, liberty and subsistence. States failing any of these criteria have no right to govern or to go to war. We can speak of states satisfying these criteria as legitimate, or "minimally just," political communities.

Why do we need to talk about these rights? First, to give state rights moral legitimacy and to avoid fetishizing state rights for their own sake. Second, to describe what is wrong about aggression and why it justifies war in response. Aggression is so serious because it involves the infliction of physical force in violation of the most elemental entitlements people and their communities have: to survive; to be physically secure; to have enough resources to subsist at all; to live in peace; and to choose for themselves their own lives and societies. *Aggression thus attacks the very spine of human civilization itself.* This is what makes it permissible to resist with means as severe as war, provided the other *jus ad bellum* criteria are also met. Third, talk of legitimacy is essential for explaining justice in a civil war, wherein there isn't classical, cross-border aggression between competing countries but, rather, a vicious fight over the one state between rival communities within a formerly united society. The key to discerning morality in such cases revolves around the idea of legitimacy: which, if any, side has minimal justice? Which side is defending—or is seeking to establish—a legitimate political structure in our three-fold sense? That's the side which it is permissible to: a) be part of; or b) if you're an outsider, to support.

How does this conception of just cause impact on the issue of armed humanitarian intervention? This is when a state does not commit cross-border aggression but, for whatever reason, turns savagely against its own people, deploying armed force in a series of massacres against large numbers of its own citizens. Such events happened in Cambodia and Uganda in the 1970s, Rwanda in 1994, Serbia/Kosovo in 1998-9 and in Sudan/Darfur from 2004 to the present. Our definitions allow us to say it's permissible to intervene on behalf of the victims, and to attack with defensive force the rogue regime meting out such death and destruction. Why? There's no logical requirement that aggression can *only* be committed across borders. Aggression is the use of armed force in violation of someone else's basic rights. That "someone else" might be: a) another person (violent crime); b) another state (international or "external" aggression); or c) many other people within one's own community (domestic or "internal" aggression). The commission of aggression, in any of these forms, causes the aggressor to forfeit its rights. The aggressor has no right not to be resisted with defensive force; indeed, the aggressor has the duty to stop and submit itself to punishment. If the aggressor doesn't stop, it is entirely permissible for its victims to resort to force to protect themselves—and for anyone else to do likewise in aid of the victims. Usually, in humanitarian intervention, armed aid from the international community is essential for an effective resistance against the aggression, since domestic populations are at a huge disadvantage, and are massively vulnerable, to the violence of their own state.

Terrorists can commit aggression too. There's nothing to the concept which excludes this: they, too, can deploy armed force in violation of someone else's basic rights. When they do so, they forfeit any right not to suffer the consequences of receiving defensive force in response. Indeed, terrorists almost always commit aggression when they act, since terrorism is precisely the use of random violence—especially killing force—against civilians, with the intent of spreading fear throughout a population, hoping this fear will advance a political objective. On 9/11, the al-Qaeda terrorist group clearly used armed force, both to gain control of the planes and then again when using the planes as missiles against the targets in The Pentagon and The World Trade Center. This use of armed force was in violation of America's state rights to political sovereignty and territorial integrity, and to all those people's human rights to life and liberty.

The terrorist strikes on 9/11 were aggression—defiantly so, deliberately modelled after Pearl Harbor. As such, they justified the responding attack on the Taliban regime in Afghanistan. The Taliban had sponsored and enabled al-Qaeda’s attack, by providing resources, personnel and a safe haven to the terrorist group.

An important issue in just cause is whether, to be justified in going to war, one must wait for the aggression *actually to happen*, or whether in some instances it is permissible to launch a pre-emptive strike against *anticipated* aggression. The tradition is severely split on this issue. Vitoria said you must wait, since it would be absurd to “punish someone for an offense they have yet to commit.” Others, like Walzer, strive to define the exceptional criteria, stressing: the seriousness of the anticipated aggression; the kind and quality of evidence required; the speed with which one must decide; and the issue of fairness and the duty to protect one’s people. If one knows a terrible attack is coming soon, one owes it to one’s people to shift from defense to offense. The best defense, as they say, is a good offense. Why let the aggressor have the upper hand of the first strike? But *that’s the very issue*: can you attack first and not, thereby, yourself become the aggressor? Can *striking first* still be considered an act of *defence* from aggression? International law, for its part, sweepingly forbids pre-emptive strikes unless they are clearly authorized in advance by the UN Security Council. These issues, of course, were highlighted in the run-up to the 2003 U.S.-led pre-emptive strike on Iraq. The U.S. still maintains, in its National Security Strategy, the right to strike first as part of its war on terror. Many other countries find this extremely controversial.

2. Right intention. A state must intend to fight the war only for the sake of its just cause. Having the right reason for launching a war is not enough: the actual motivation behind the resort to war must also be morally appropriate. Ulterior motives, such as a power or land grab, or irrational motives, such as revenge or ethnic hatred, are ruled out. The only right intention allowed is to see the just cause for resorting to war secured and consolidated. If another intention crowds in, moral corruption sets in. International law does not include this rule, probably because of the evidentiary difficulties involved in determining a state’s intent.

3. Proper authority and public declaration. A state may go to war only if the decision has been made by the appropriate authorities, according to the proper process, and made public, notably to its own citizens and to the enemy state(s). The “appropriate authority” is usually specified in that country’s constitution. States failing the requirements of minimal justice lack the legitimacy to go to war.

4. Last Resort. A state may resort to war only if it has exhausted all plausible, peaceful alternatives to resolving the conflict in question, in particular diplomatic negotiation. One wants to make sure something as momentous and serious as war is declared only when it seems the last practical and reasonable shot at effectively resisting aggression.

5. Probability of Success. A state may not resort to war if it can foresee that doing so will have no measurable impact on the situation. The aim here is to block mass violence which is going to be futile. International law does not include this requirement, as it is seen as biased against small, weaker states.

6. Proportionality. A state must, prior to initiating a war, weigh the *universal* goods expected to result from it, such as securing the just cause, against the *universal* evils expected to result, notably casualties. Only if the benefits are proportional to, or “worth”, the costs may the war action proceed. (The universal must be stressed, since often in war states only tally *their own* expected benefits and costs, radically discounting those accruing to the enemy and to any innocent third parties.)

Just war theory insists *all six* criteria must each be fulfilled for a particular declaration of war to be justified: it’s all or no justification, so to speak. Just war theory is thus quite demanding, as of course it should be, given the gravity of its subject matter. It is important to note that the first three of these six rules are what we might call deontological requirements, otherwise known as duty-based requirements or first-principle requirements. For a war to be just, some core duty must be violated: in this case, the duty not to commit aggression. A war in punishment of this violated duty must itself respect further duties: it must be appropriately motivated, and must be publicly declared by (only) the proper authority for doing so. The

next three requirements are consequentialist: given that these first principle requirements have been met, we must also consider the expected consequences of launching a war. Thus, just war theory attempts to provide a common sensical combination of both deontology and consequentialism as applied to the issue of war.

2.2 *Jus in bello*

Jus in bello refers to justice in war, to right conduct in the midst of battle. Responsibility for state adherence to *jus in bello* norms falls primarily on the shoulders of those military commanders, officers and soldiers who formulate and execute the war policy of a particular state. They are to be held responsible for any breach of the principles which follow below. Such accountability may involve being put on trial for war crimes, whether by one's own national military justice system or perhaps by the newly-formed International Criminal Court (created by the 1998 Treaty of Rome).

We need to distinguish between external and internal *jus in bello*. External, or traditional, *jus in bello* concerns the rules a state should observe regarding the enemy and its armed forces. Internal *jus in bello* concerns the rules a state must follow in connection with its own people as it fights war against an external enemy.

There are several rules of external *jus in bello*:

1. Obey all international laws on weapons prohibition. Chemical and biological weapons, in particular, are forbidden by many treaties. Nuclear weapons aren't so clearly prohibited but it seems fair to say a huge taboo attaches to such weapons and any use of them would be greeted with incredible hostility by the international community.

2. Discrimination and Non-Combatant Immunity. Soldiers are only entitled to use their (non-prohibited) weapons to target those who are, in Walzer's words, "engaged in harm." Thus, when they take aim, soldiers must discriminate between the civilian population, which is morally immune from direct and intentional attack, and those legitimate military, political and industrial targets involved in rights-violating harm. While some *collateral* civilian casualties are excusable, it is wrong to take deliberate aim at civilian targets. An example would be saturation bombing of residential areas. (It is worth noting that almost all wars since 1900 have featured larger civilian, than military, casualties. Perhaps this is one reason why this rule is the most frequently and stridently codified rule in all the laws of armed conflict, as international law seeks to protect unarmed civilians as best it can.)

3. Proportionality. Soldiers may only use force proportional to the end they seek. They must restrain their force to that amount appropriate to achieving their aim or target. Weapons of mass destruction, for example, are usually seen as being out of proportion to legitimate military ends.

4. Benevolent quarantine for prisoners of war (POWs). If enemy soldiers surrender and become captives, they cease being lethal threats to basic rights. They are no longer "engaged in harm." Thus it is wrong to target them with death, starvation, rape, torture, medical experimentation, and so on. They are to be provided, as The Geneva Conventions spell out, with benevolent—not malevolent—quarantine away from battle zones and until the war ends, when they should be exchanged for one's own POWs. Do terrorists deserve such protection, too? Great controversy surrounds the detainment and aggressive questioning of terrorist suspects held by the U.S. at jails in Cuba, Iraq and Pakistan in the name of the war on terror.

5. No Means *Mala in Se*. Soldiers may not use weapons or methods which are "evil in themselves." These include: mass rape campaigns; genocide or ethnic cleansing; using poison or treachery (like disguising soldiers to look like the Red Cross); forcing captured soldiers to fight against their own side; and using weapons whose effects cannot be controlled, like biological agents.

6. No reprisals. A reprisal is when country A violates *jus in bello* in war with country B. Country B then retaliates with its own violation of *jus in bello*, seeking to chasten A into obeying the rules. There are strong moral and evidentiary reasons to believe that reprisals don't work, and they instead serve to escalate death and make the destruction of war increasingly indiscriminate. Winning well is the best revenge.

Internal *jus in bello* essentially boils down to the need for a state, even though it's involved in a war, nevertheless *to still respect the human rights of its own citizens as best it can during the crisis*. The following issues arise: is it just to impose conscription, or press censorship? Can one curtail traditional civil liberties, and due process protections, for perceived gains in national security? Should elections be cancelled or postponed? May soldiers disobey orders, e.g. refuse to fight in wars they believe unjust? A comprehensive theory of wartime justice must include consideration of them, and not merely focus on what one may do to the enemy. For some of the worst atrocities in wartime have occurred *within, and not between*, national borders. Some states, historically, have used the cloak of war with foreign powers to engage in massive internal human rights violations, usually against some disfavoured group. Other states, which are otherwise decent, panic amidst the wartime situation and impose emergency legislation which turns out to have been complete overkill, and which they later regret and view as the product of fear rather than reason.

2.3 *Jus post bellum*

Jus post bellum refers to justice during the third and final stage of war: that of war termination. It seeks to regulate the ending of wars, and to ease the transition from war back to peace. There is little international law here—save occupation law and perhaps the human rights treaties—and so we must turn to the moral resources of just war theory. But even here the theory has not dealt with *jus post bellum* to the degree it should. There is a newness, unsettledness and controversy attaching to this important topic. To focus our thoughts, consider the following proposed principles for *jus post bellum*:

1. Proportionality and Publicity. The peace settlement should be measured and reasonable, as well as publicly proclaimed. To make a settlement serve as an instrument of revenge is to make a volatile bed one may be forced to sleep in later. In general, this rules out insistence on unconditional surrender.

2. Rights Vindication. The settlement should secure those basic rights whose violation triggered the justified war. The relevant rights include human rights to life and liberty and community entitlements to territory and sovereignty. This is the main substantive goal of any decent settlement, ensuring that the war will actually have an improving affect. Respect for rights, after all, is a foundation of civilization, whether national or international. Vindicating rights, not vindictive revenge, is the order of the day.

3. Discrimination. Distinction needs to be made between the leaders, the soldiers, and the civilians in the defeated country one is negotiating with. Civilians are entitled to reasonable immunity from punitive post-war measures. This rules out sweeping socio-economic sanctions as part of post-war punishment.

4. Punishment #1. When the defeated country has been a blatant, rights-violating aggressor, proportionate punishment must be meted out. The leaders of the regime, in particular, should face fair and public international trials for war crimes.

5. Punishment #2. Soldiers also commit war crimes. Justice after war requires that such soldiers, *from all sides to the conflict*, likewise be held accountable to investigation and possible trial.

6. Compensation. Financial restitution may be mandated, subject to both proportionality and discrimination. A post-war poll tax on civilians is generally impermissible, and there needs to be enough resources left so that the defeated country can begin its own reconstruction. To beggar thy neighbor is to pick future fights.

7. Rehabilitation. The post-war environment provides a promising opportunity to reform decrepit institutions in an aggressor regime. Such reforms are permissible but they must be proportional to the degree of depravity in the regime. They may involve: demilitarization and disarmament; police and judicial re-training; human rights education; and even deep structural transformation towards a minimally just society governed by a legitimate regime. This is, obviously, the most controversial aspect of *jus post bellum*.

The terms of a just peace should satisfy all these requirements. There needs, in short, to be an *ethical* “exit strategy” from war, and it deserves at least as much thought and effort as the purely military exit strategy so much on the minds of policy planners and commanding officers.

Any serious defection, by any participant, from these principles of just war settlement should be seen as a violation of the rules of just war termination, and so should be punished. At the least, violation of such principles mandates a new round of diplomatic negotiations—even binding international arbitration—between the relevant parties to the dispute. At the very most, such violation may give the aggrieved party a just cause—but *no more than a just cause*—for resuming hostilities. Full recourse to the resumption of hostilities may be made *only if* all the other traditional criteria of *jus ad bellum*—proportionality, last resort, etc.—are satisfied in addition to just cause.

Perhaps a few additional thoughts on coercive regime change should here be added, in light of controversial recent events, especially in Afghanistan and Iraq. Can coercive regime change ever be justified, or is it essentially an act of imperialism? In my view, forcible post-war regime change *can be permissible provided*: 1) the war itself was just and conducted properly; 2) the target regime was illegitimate, thus forfeiting its state rights; 3) the goal of the reconstruction is a minimally just regime; and 4) respect for *jus in bello* and human rights is integral to the transformation process itself. The permission is then granted because the transformation: 1) violates neither state nor human rights; 2) its expected consequences are very desirable, namely, satisfied human rights for the local population and increased international peace and security for everyone; and 3) the post-war moment is especially promising regarding the possibilities for reform. And the transformation will be successful when there’s: 1) a stable new regime; 2) run entirely by locals; which is 3) minimally just. There is extensive historical evidence that this kind of success probably takes from 8 to 12 years to achieve (essentially, a decade). Note that successful, rights-respecting coercive regime change *can* be done, contrary to some pessimistic views; it was actually done in Germany and Japan from 1945-55, and so it is neither conceptually nor empirically impossible. It’s very difficult, to be sure—and, in some cases, it’s not a wise thing to do—but it’s not literally impossible.

A review of the literature suggests something of a 10-point recipe for transforming a defeated aggressive regime into one which is minimally just:

- Adhere diligently to the laws of war during the regime take-down and occupation.
- Purge much of the old regime, and prosecute its war criminals.
- Disarm and demilitarize the society.
- Provide effective military and police security for the whole country.
- Work with a cross-section of locals on a new, rights-respecting constitution which features checks and balances.
- Allow other, non-state associations, or “civil society”, to flourish.
- Forego compensation and sanctions in favour of investing in and re-building the economy.
- If necessary, re-vamp educational curricula to purge past poisonous propaganda and cement new and better values.
- Ensure, in a timely fashion, that the benefits of the new order will be: 1) concrete; and 2) widely, and not narrowly, distributed. The bulk of the population must feel their lives after the regime change are clearly better than their former lives for the change to be sustainable.
- Follow an orderly, not-too-hasty exit strategy when the new regime can stand on its own two feet. Again, this will probably take a decade of intensive effort.

To summarize this whole section, just war theory offers rules to guide decision-makers on the appropriateness of their conduct during the resort to war, conduct during war and the termination phase of the conflict. Its over-all aim is to try and ensure that wars are begun only for a very narrow set of truly defensible reasons, that when wars break out they are fought in a responsibly controlled and targeted manner, and that the parties to the dispute bring their war to an end in a speedy and responsible fashion that respects the requirements of justice.

3. Realism

Realism is most influential amongst political scientists, as well as scholars and practitioners of international relations. While realism is a complex and often sophisticated doctrine, its core propositions express a strong suspicion about applying moral concepts, like justice, to the conduct of international affairs. Realists believe that moral concepts should be employed neither as descriptions of, nor as prescriptions for, state behaviour on the international plane. Realists emphasize power and security issues, the need for a state to maximize its expected self-interest and, above all, their view of the international arena as a kind of anarchy, in which the will to power enjoys primacy.

Referring specifically to war, realists believe that it is an inevitable part of an anarchical world system; that it ought to be resorted to only if it makes sense in terms of national self-interest; and that, once war has begun, a state ought to do whatever it can to win. In other words, “all’s fair in love and war.” During the grim circumstances of war, “anything goes.” So if adhering to the rules of just war theory, or international law, hinders a state during wartime, it should disregard them and stick steadfastly to its fundamental interests in power, security and economic growth. Prominent classical realists include Thucydides, Machiavelli and Hobbes. Modern realists include Hans Morgenthau, George Kennan, Reinhold Niebuhr and Henry Kissinger, as well as so-called neo-realists, such as Kenneth Waltz.

It is important to distinguish between descriptive and prescriptive realism. Descriptive realism is the claim that states, *as a matter of fact*, either do not (for reasons of motivation) or cannot (for reasons of competitive struggle) behave morally, and thus moral discourse surrounding interstate conflict is empty, the product of a category mistake. States are simply not animated in terms of morality and justice: it’s all about power, security and national interest for them. States are not like “big persons”: they are creations of an utterly different kind, and we cannot expect them to live by the same rules and principles we require of individual persons, especially those in peaceful, developed societies. Morality is a luxury states can’t afford, for they inhabit a violent international arena, and they’ve got to be able to get in that game and win, if they are to serve and protect their citizens in an effective way over time. Morality is simply not on the radar screen for states, given their defensive function and the brutal environment in which they subsist.

Walzer offers arguments against this kind of realism, contending that states are in fact responsive to moral concerns, even when they fail to live up to them. States, because they are the creation of individual persons, want to act morally and justly: it could not be otherwise. Walzer goes so far as to say that any state which was motivated by nothing more than the struggle to survive and win power *could not over time sustain the support from its own population*, which demands a deeper sense of community and justice. He also argues that all the pretence regarding “the necessity” of state conduct in terms of pursuing power is exaggerated and rhetorical, ignoring the clear reality of foreign policy *choice* enjoyed by states in the global arena. States are not frequently forced into some kind of dramatic, do-or-die struggle: the choice to go to war is a deliberate one, freely entered into and often hotly debated and agonized over before the decision is made. And this is leaving unspoken the argument regarding the defiant, Machiavellian amorality behind certain kinds of realism, and the moral calibre of the actions it might recommend on this basis. For example, if it’s all about power and winning in the competitive struggle, does that make it alright to unleash weapons of mass destruction? Or to launch a mass rape campaign? Commit genocide and just get rid of those bastards? Just war theory suggests not, and just war theorists like Walzer want to claim that the rest of us agree.

Prescriptive realism, though, need not be rooted in any form of descriptive realism. Prescriptive realism is the claim that a state ought (prudential “ought”) to behave amorally in the international arena. A state should, for prudence’s sake, adhere to an amoral policy of smart self-regard in international affairs. A smart

state will leave its morality at home when considering what to do on the international stage. Why? Because if it's too moral, it will be exploited by other states more ruthless. Nice guys finish last. Or, a moralized and moralizing state will offend other communities, whose communities sport different values. Better to stick to the sober calculus of national interests and leave ethics out of it.

It's important to note that a prescriptive realist might, in the end, actually endorse rules for the regulation of warfare, much like those offered by just war theory. These rules include: "Wars should only be fought in response to aggression"; and "During war, non-combatants should not be directly targeted with lethal violence." Of course, *the reason why* a prescriptive realist might endorse such rules would be very different from the reasons offered by the just war theorist: the latter would talk about abiding moral values whereas the former would refer to useful rules which help establish expectations of behaviour, solve coordination problems and to which prudent bargainers would consent. Just war rules, the prescriptive realist might claim, do not have independent moral purchase on the attention of states. These rules are what Douglas Lackey calls "salient equilibria", stable conventions limiting war's destructiveness which all prudent states can agree on, assuming general compliance. There might even be some room for overlap between this kind of realism and just war theory.

4. Pacifism

It seems best to rely on Jenny Teichman's definition of pacifism as "anti-war-ism." Literally and straightforwardly, a pacifist rejects war in favour of peace. It is not violence in all its forms that the most challenging kind of pacifist objects to; rather, it is the specific kind and degree of violence that war involves which the pacifist objects to. A pacifist objects to killing (not just violence) in general and, in particular, she objects to the mass killing, for political reasons, which is part and parcel of the wartime experience. So, a pacifist rejects war; she believes that there are no moral grounds which can justify resorting to war. War, for the pacifist, is always wrong.

Mention should straight away be made of a very popular just war criticism of pacifism which will not be used here. This criticism is that pacifism amounts to an indefensible "clean hands policy." The pacifist, it is said, refuses to take the brutal measures necessary for the defense of himself and his country, for the sake of maintaining his own inner moral purity. It is contended that the pacifist is thus a kind of free-rider, gathering all the benefits of citizenship while not sharing all its burdens. Another inference drawn is that the pacifist himself constitutes a kind of internal threat to the over-all security of his state.

This "clean hands" argument is easily, and frequently, over-stated. It is important to note that, to the extent to which any moral stance will commend a certain set of actions or intentions deemed morally worthy, and condemn others as being reprehensible, the "clean hands" criticism is so malleable as to apply to nearly any substantive doctrine. Every moral and political theory stipulates that one ought to do what it deems good or just and to avoid what it deems bad or unjust. So this popular just war criticism of pacifism is not strong. The very idea of a selfish pacifist simply does not ring true: many pacifists have, historically, paid a very high price for their pacifism during wartime (through severe ostracism and even jail time) and their pacifism seems less rooted in regard for inner moral purity than it is in regard for constructing a less violent and more humane world order. So, this argument against pacifism fails; but what of others?

Walzer contends that pacifism's idealism is excessively optimistic. In other words, pacifism lacks realism. More precisely, the nonviolent world imagined by the pacifist is not actually attainable, at least for the foreseeable future. Since "ought implies can", the set of "oughts" we are committed to must express a moral outlook on war less utopian in nature. While we are committed to morality in wartime, we are forced to concede that, sometimes in the real world, resorting to war can be morally justified. It's hard to see, e.g., how anything but war could've defeated the Nazis.

Another objection to pacifism is that, by failing to resist international aggression with effective means, it ends up rewarding aggression and failing to protect people who need it. Pacifists reply to this argument by contending that we do not need to resort to war in order to protect people and punish aggression effectively.

In the event of an armed invasion by an aggressor state, an organized and committed campaign of non-violent civil disobedience—perhaps combined with international diplomatic and economic sanctions—would be just as effective as war in expelling the aggressor, with much less destruction of lives and property. After all, the pacifist might say, no invader could possibly maintain its grip on the conquered nation in light of such systematic isolation, non-cooperation and non-violent resistance. How could it work the factories, harvest the fields, or run the stores, when everyone would be striking? How could it maintain the will to keep the country in the face of crippling economic sanctions and diplomatic censure from the international community? And so on.

Though one cannot exactly disprove this pacifist proposition—since it is a counter-factual thesis—there are powerful reasons to agree with John Rawls that such is “an unworldly view” to hold. For, as Walzer points out, the effectiveness of this campaign of civil disobedience relies on the scruples of the invading aggressor. But what if the aggressor is utterly brutal, remorseless? What if, faced with civil disobedience, the invader “cleanses” the area of the native population, and then imports its own people from back home? What if, faced with economic sanctions and diplomatic censure from a neighbouring country, the invader decides to invade it, too? We have some indication from history, particularly that of Nazi Germany, that such pitiless tactics are effective at breaking the will to resist of even very principled people. The defence of our lives and rights may well, against such invaders, require the use of political violence. Under such conditions, Walzer says, adherence to pacifism might even amount to “a disguised form of surrender.”

Pacifists respond to this accusation of “unworldliness” by citing what they believe are real world examples of effective non-violent resistance to aggression. Examples mentioned include Mahatma Ghandi’s campaign to drive the British Imperial regime out of India in the late 1940s and Martin Luther King Jr.’s civil rights crusade in the 1960s on behalf of African-Americans. Walzer replies curtly that there is no evidence that non-violent resistance has ever, of itself, succeeded. This may be rash on his part, though it is clear that Britain’s own exhaustion after WWII, for example, had much to do with the evaporation of its Empire. Walzer’s main counter-argument against these pacifist counter-examples is that *they only illustrate his main point*: that effective non-violent resistance depends upon the scruples of those it is aimed against. It was only because the British and the Americans had some scruples, and were moved by the determined idealism of the non-violent protesters, that they acquiesced to their demands. But aggressors will not always be so moved. A tyrant like Hitler, for example, might interpret non-violent resistance as weakness, deserving contemptuous crushing. “Non-violent defense”, Walzer suggests, “is no defense at all against tyrants or conquerors ready to adopt such measures.”

As sensible as Walzer's remarks might seem, they remain quite narrow, by no means constituting an all-things-considered refutation of pacifism. Generally, there are two kinds of modern secular pacifism to consider: (1) a more consequentialist form of pacifism (or CP), which maintains that the benefits accruing from war can never outweigh the costs of fighting it; and (2) a more deontological form of pacifism (or DP), which contends that the very activity of war is intrinsically wrong, since it violates foremost duties of justice, such as not killing human beings. Most common amongst contemporary secular pacifists, such as Robert Holmes, is a doctrine which attempts to combine both CP and DP. (No discussion will be made here as to religious forms of pacifism. While they have been very influential historically, especially their Christian variants, as theoretical propositions I believe they rest on core premises which are too contentious and exclusionary. But the Christian pacifist literature is a very rich source of information for those interested.)

What arguments might a just war theorist employ to overcome CP and DP? A just war theorist might, for starters, focus on the relationship in CP between consequentialism and the denial of killing. Pacifism in either form places overriding value on respecting human life, notably through its injunction against killing. But this value seems to rest uneasily with consequentialism, for there is nothing inherent to consequentialism which bans killing as such. There is no absolute rule, or side-constraint, that one ought never to kill another person, or that nations ought never to deploy lethal armed force in war. With consequentialism, it’s always a matter of considering the latest costs and benefits, of choosing the best option amongst feasible alternatives. Consequentialism therefore leaves conceptual space open to the claim that under these conditions, at this time and place, and given these alternatives, killing and/or war appears

permissible. After all, what if killing x people (say, soldiers in an aggressive army) appears the best option if we are to save the lives of $x + n$ people (say, fellow citizens who would perish under the brutal heel of an unchecked aggressor)? It is at least conceivable that a quick and decisive resort to war could prevent even greater killing and devastation in the future. Historians speculate, e.g., that an earlier confrontation with Hitler would've prevented World War II from ending up being so widespread and destructive. These are two telling points: CP does not, of itself, ground the *categorical* rejection of killing and war which is the essence of pacifism; and CP is open to counter-examples which question whether consequentialism would reject killing and war at all under certain conditions. Consequentialism might even, in a particular case, go so far as to *recommend* war under certain conditions.

Casting doubt on DP is a complicated procedure. Only a sketch of plausible just war theory arguments can here be offered. The first question to ask is: which foremost duty does DP understand being violated by warfare? If the DP response is the duty not to kill another human being, then contention can be made that this is by no means uncontroversial. Consider the most obvious counter-example: aggressor A attacks B for no defensible reason, posing a serious threat to B 's life. Some would suggest, in good faith, that B is not duty-bound not to kill A if such seems necessary to stop A 's aggression. Indeed, they would argue that B may kill A in legitimate self-defence. The DP pacifist, however, might reply that extending B moral permission to kill A , even in self-defence, violates the human rights of A . He might contend that just war theory merely compounds the wrongness of the situation by paradoxically permitting lethal force to stop lethal force. There's a clever phrase nowadays: an eye for an eye leaves us both blind.

One just war theory rejoinder to this DP contention is this: B does no wrong whatsoever—violates no human rights—by responding to A 's aggression with lethal force if required. Why does B do nothing wrong? First, it is A who is *responsible* for forcing B to choose between her own life and rights and those of A . We can hardly blame B for choosing her own. For if she does not choose her own, she loses an enormous amount, perhaps everything. And it is patently *unreasonable* to expect creatures like us to suffer catastrophic loss by default. Consider also the issue of *fairness*: if B is not allowed to use lethal force, if necessary, against A in the event of A 's aggression, then B loses everything while A loses nothing. Indeed, A gains whatever object he desired in violating or killing B . Such is an unfair reward of awful behaviour. Finally, B 's having rights at all provides her with an *implicit entitlement* to use those means necessary to secure her rights, including the use of force in the face of a serious physical threat. These powerful considerations of responsibility, reasonableness, fairness and implicit entitlement come together in support of the just war claims that: B may respond with needed lethal force to A 's initial aggression; B does no wrong in doing so; it would be wrong to prohibit B 's doing so; and that A bears all of the blame for the situation. It is A who should stop, not B who should succumb.

DP pacifists are not, at this point, out of options. Holmes, for example, suggests that the foremost duty of justice violated by war is not the duty not to kill aggressors, but rather the duty not to kill innocent, non-aggressive human beings. To be innocent here means to have done nothing which would justify being harmed or killed; in particular, it means not constituting a serious threat to the lives and rights of other people. It is this sense of innocence that just war theory invokes when it claims that civilians should not be directly attacked during wartime. Even if civilians support the war effort politically, or even in terms of their personal attitudes towards the war, they clearly do not pose serious threats to others. Only armed forces, and the political-industrial-technological complexes which guide them, constitute serious threats against which threatened communities may respond in kind. Civilian populations, just war theory surmises, are morally off-limits as targets. Holmes contends that this just war (and international law) rule of non-combatant immunity can never be satisfied. For all possible wars in this world—given the nature of military technology and tactics, the heat of battle, and the limits of human knowledge and self-discipline—involve the killing of innocents, thus defined. We know this to be true from history and have no good reason for expecting otherwise in the future. But the killing of innocents, Holmes says, is always unjust. So no war can ever be fought justly, regardless of the nature of the goal sought after, such as national defence from an aggressor's attack. The very activities needed to fight wars are intrinsically corrupt, and cannot be redeemed by the supposed justice of the ends they are aimed at. How is a just war theorist to respond to this DP challenge?

Some respond by casting doubt on the concept of innocence in wartime. But a just war theorist subscribing to the rule of non-combatant immunity will neither want, nor logically be at liberty, to argue in this fashion. It is hard to see, for example, how infants could be anything other than innocent during a war, and as such entitled not to be made the object of direct and intentional attack. It is only those who, in Walzer's phrase, are "involved in harming us"—i.e. those who pose serious threats to our lives and rights—that we can justly target in a direct and intentional fashion during wartime.

The more appropriate just war response invokes, alongside Walzer, the doctrine of double effect (or DDE). The DDE, invented by Aquinas, is a complex idea. In spite of its apparent technicality, though, the DDE is closely related to our ordinary ways of thinking about moral life. The DDE assumes the following scenario: agent *X* is considering performing an action *T*, which *X* foresees will produce both good/moral/just effects *J* and bad/immoral/unjust effects *U*. The DDE permits *X* to perform *T* only if: 1) *T* is otherwise permissible; 2) *X* only intends *J* and not *U*; 3) *U* is not a means to *J*; and 4) the goodness of *J* is worth, or is proportionately greater than, the badness of *U*. Assume now that *X* is a country and *T* is war. The government of *X*, contemplating war in response to an attack by aggressor country *Y*, foresees that, should it embark on war to defend itself, civilian casualties will result, probably in both *X* and *Y*. The DDE stipulates that *X* may launch into this defensive (and thus otherwise permissible) war only if: 1) *X* does not intend the resulting civilian casualties but rather aims only at defending itself and its people; 2) such casualties are not themselves the means whereby *X*'s end is achieved; and 3) the importance of *X*'s defending itself and its people from *Y*'s aggression is proportionately greater than the badness of the resulting civilian casualties. The DDE, in making these claims, refers to common shared principles regarding the moral importance of intent, of appealing to better expected consequences, and insisting that bad not be done so that good may follow from it.

Just war theorists claim that civilians are not entitled to absolute immunity from attack during wartime. Civilians are owed neither more nor less than what Walzer calls "due care" from the belligerent governments that they not be made casualties of the war action in question. "Due care" involves fighting only in certain ways, applying limited force to specific targets. Essentially, "due care" means fighting in adherence with *jus in bello*. But does this just war claim simply beg the question against the latest DP principle? DPs insist on *absolute* immunity for civilians, which in our world would result in banning warfare, whereas just war theorists, acknowledging the threat, seem to dodge it by re-defining the immunity to which civilians are entitled, demoting it to mere "due care." Despite appearances, it is *not* question-begging *but* principled disagreement which roots the difference. Just war theorists will argue that civilians cannot be entitled to absolute immunity because that would outlaw all warfare. But outlawing all warfare would ignore *both* the responsibility for interstate aggression *and* the implicit entitlement of a state to use necessary means (including armed force) to secure the lives and rights of its citizens from serious and standard threats to them. In the real world, it is neither reasonable nor fair to require a political community not to avail itself of the most effective means available for resisting an aggressive invasion which threatens the lives and rights of its citizens. It is simply not reasonable to require a state to stand down while an aggressor—be it state or terrorist—wreaks havoc, murder and mayhem upon its people.

This is not a complete defeat for DP, merely a suggestion of how such defeat might be sought. DP probably constitutes the most formidable moral challenge to just war theory (whereas prescriptive realism constitutes the most formidable prudential challenge to just war theory). Suffice it for our purposes to say that the DDE is the just war principle most frequently employed to defeat the DP pacifist's assertion that it is always wrong to kill innocent human beings. Just war theorists prefer to substitute, for this DP claim, the following proposition: what is always wrong, both in peace and war, is to kill innocent human beings *intentionally and deliberately*. Unintended, collateral civilian casualties can be excused during the prosecution of an otherwise just war, wherein the end is the repulsion of aggression and the means are aimed at legitimate military targets.

5. Conclusion

This [encyclopedia] entry provides a sample of the rich and controversial argumentation surrounding philosophical discourse on war. This discourse is dominated by three major traditions of thought: just war

theory (and its international law subsidiary); realism; and pacifism. The interaction between these three traditions structures the contemporary discussion of wartime issues, at the same time as it fuels fascinating debate about them. While just war theory occupies an especially large and influential space within the discourse, its realist and pacifist alternatives endure as provocative challenges to the philosophical mainstream which it represents.

6. Guide to the Literature

I discuss all these issues and more, with extensive reference to cases, in my forthcoming book, *The Morality of War* (Broadview, 2006).

All the works cited in this entry, plus relevant other works, are listed below. It may be helpful to first locate and emphasize some of the major and most influential sources.

For scholarship on the history and development of just war theory, consult the works of James T. Johnson. Hugo Grotius is often cited as the most formidable classical just war theorist (though I'd rank Vitoria up there myself). A translation of his works can be found in J. Scott's edition of *Classics of International Law*. The major contemporary statement of just war theory remains Michael Walzer's *Just and Unjust Wars*. For other comprehensive contemporary statements, see the works of: Paul Christopher; J.B. Elshtain; Michael Ignatieff; Doug Lackey; Brian Orend; and Richard Regan. Works critical of just war theory can be found in the pacifist and realist tracts below.

Other important articles on particular aspects of just war theory include: on *jus ad bellum*, D. Luban, "Just War and Human Rights"; on *jus in bello*, T. Nagel's "War and Massacre" and R. Fullinwider's "War and Innocence"; and on *jus post bellum*, Kant's "Perpetual Peace" (in his *Political Writings*) and B. Orend's "Justice After War".

Hans Morgenthau's *Politics Among Nations* remains an often-cited defense of realism, as does G. Kennan's *Realities of American Foreign Policy*. Henry Kissinger's *Diplomacy* provides the same outlook in perhaps more accessible form. Two of the most focused and effective criticisms of the realist approach to war occur at: Chapter 1 of Walzer's *Just and Unjust Wars*; and Chapters 1–3 of R. Holmes' *On War and Morality*.

The three best contemporary, secular works defending pacifism are: R. Holmes, *On War and Morality*; J. Teichman, *Pacifism and the Just War*; and R. Norman, *Ethics, Killing and War*. Two renowned critical essays on pacifism, both reprinted in R. Wasserstrom, ed. *War and Morality*, are G.E.M. Anscombe's "War and Murder" and Jan Narveson's "Pacifism: A Philosophical Analysis".

One prominent writer on the philosophy of war who resists easy classification into any of these categories is Carl von Clausewitz. Clausewitz wrote *On War*, one of the most influential general sources, cited by soldiers and statesmen as often as by philosophers or international lawyers. M. Gelven's *War and Existence* is an interesting contemporary piece on the meaning and experience of war, with a Clausewitzian flavor to it.

In terms of international law, I strongly recommend the web-sites below. For hard copy sources, see especially: W. Reisman and C. Antoniou, eds. *The Laws of War: A Comprehensive Collection of Primary Documents Governing Armed Conflict* and A. Roberts and R. Guelff, eds., *Documents on The Laws of War*.

[An extensive bibliography – 4 pages at 9 point font – follows at this point of the original article. I have excised it to save space. PR]

Computers, Ethics, and Collective Violence

by
Craig Summers and Eric Markusen

Controversy Corner from *The Journal of Systems and Software*, vol. 17, pp. 91–103, 1992

This article extends the emerging debate and discussion over ethical dimensions of computer science from issues such as software piracy, viruses, and unauthorized systems entry to the realm of collective violence. We view collective violence as actions by large numbers of people that contributes to large-scale destruction. Several ways in which computer professionals may contribute to actual or potential violence are briefly discussed. Then, to understand how well-meaning computer professionals can do work of the highest technical quality, but which is routinized and isolated from its social effects, we discuss three types of psychosocial mechanisms: (1) psychological-level aspects of one's own role; (2) bureaucratic factors routinizing individual involvement, and (3) specific factors in scientific and technological work affecting perceived responsibility. To understand why these mechanisms occur, the importance of importance of perceived short-term economic needs for day-to-day living are considered against values and ethics. A predictive model of temporal and social "traps" is outlined that explains when individuals may contribute to harmful projects regardless of social values and human welfare. Finally, we explore how codes of ethics, education about ethics, and other policy initiatives can help professionals do work that avoids harmful risks and consequences and produces benefits individually and collectively.

INTRODUCTION

Professions in contemporary society can be characterized by four defining features: they possess specialized knowledge; they are important to society; they enjoy a high degree of autonomy and self-regulation; and they are guided by an ideology of public service [1]. The latter two features involve ethics, defined here as moral guidelines for behavior. Thus, most professions have codes of ethics to which all members in good standing are expected to adhere.

However, simply having codes of ethics does not guarantee ethical behavior. As society and technology changes, new situations arise which create new ethical dilemmas. Also, if students and practitioners of a profession are not carefully instructed about ethical issues and concerns relevant to their profession, it is unlikely that they will be guided by them.

Ethics are every bit as relevant to the profession of computer science as they are to other contemporary professions. There has been widespread and influential dissemination of computer technology in recent years, although this profession is still relatively young (e.g., personal computers are less than 15 years old). Examination of ethical issues that relate to computer professionals¹—as embodied in this special issue of *The Journal of Systems and Software*—is therefore both welcome and necessary. Practices such as illegal duplication of software, insertion of harmful viruses, and unauthorized entry and retrieval of private files all need careful exposure and analysis in terms of ethical principles.

This article, however, examines a rather different ethical dimension that is nonetheless relevant to computer scientists. Rather than focus on ethical issues such as viruses, abuse of passwords, privacy, and copyrights, we are concerned with the possibility that computer professionals may lend their expertise to activities and projects that involve harm to other human beings on a large scale. We are, in short, concerned with the relations among computers, ethics, and collective violence. By "collective violence" we mean large-scale destruction to which many people have contributed.

This article has five primary objectives, which are examined in the sections that follow. First, we will briefly address the problem of collective violence during the twentieth century. Second, we hope to persuade readers that they should be concerned with the problem of collective violence. Third and fourth, we will summarize relevant literature from psychology and sociology to explain how and why normal individuals—including professionals—contribute to collective violence. Finally, we suggest how

¹ "Professional" is used here in a broad sense, referring to occupations including programmers, systems analysts, engineers, technicians, and computer scientists.

professional codes of ethics and education about ethics can help professionals and professions avoid unethical behavior and involvement in harmful enterprises.

COLLECTIVE VIOLENCE TURING THE TWENTIETH CENTURY

Anyone who reads the newspaper or watches the news on television is painfully aware of the prevalence of collective violence throughout the world. In this section, we discuss a number of relationships between professionals and collective violence.

First, collective violence can occur in a wide variety of forms. Warfare, which can take place between nations or groups of nations (international war) as well as between groups within a nation (civil war), is perhaps the most widely recognized and thoroughly studied form of collective violence. Genocide, a term invented only in 1944, refers to the deliberate destruction of groups of human beings because of their racial, ethnic, religious, or political identity. When governments permit and enforce official discrimination and violation of human rights—for example, apartheid in South Africa and torture and “disappearances” in Argentina—large numbers of people suffer and some lose their lives. Likewise, certain corporate practices, such as exploitation of the environment or tolerance of dangerous workplace conditions, can hurt many people. Finally, the nuclear arms race, even though it has been justified as a deterrent, poses the ever-present threat of collective violence on an unimaginable scale.

Second, some scholars have argued that the scale of collective violence is greater during this century than at any other period in history [2]. One analyst of genocidal violence estimates that more than 100,000,000 people have been killed by governments during the twentieth century [3]. Another scholar counted 22 wars underway in 1987—more than in any other single year in human history [4]. Military historians and weapons experts argue that the intensity and lethality of war in the present century greatly exceeds anything in history [5, 6]. Projections of the possible results of a nuclear war have estimated that more than one billion people could be killed [7] and the planetary ecosystems catastrophically damaged [8]. The unprecedented levels of violence probably do not reflect any increase in aggressiveness or brutality among human beings, but rather their possession of more effective technologies for killing [9].

A third aspect of professionals and collective violence is that most of the individuals who contribute to collective violence are psychologically normal and motivated by idealistic concerns. Studies of the Holocaust, for example, have found that the vast majority of Nazi perpetrators were “...normal people according to currently accepted definitions by the mental healthy profession” [10, p. 148]. This finding has been corroborated by numerous other scholars [1].

Finally, professions and professionals make crucial contributions to most forms of collective violence. Again using the Holocaust as an illustration, there is strong consensus among scholars that educated professionals played indispensable roles in rationalizing and implementing the extermination of the Jews [11]. In his study of German doctors in the Holocaust, Robert Lifton [12] found that these health care professionals made crucial contributions to the killing process, even peering through the peepholes in the gas chamber doors to determine when the victims were dead.

WHY COMPUTER PROFESSIONALS SHOULD BE CONCERNED ABOUT COLLECTIVE VIOLENCE

If psychologically normal professionals could be implicated in violence as repugnant and brutal as the Holocaust, it is also conceivable that other professionals could make equally destructive contributions now, particularly if the effects are less apparent. Therefore, the primary reason that computer professionals should be concerned about collective violence is as potential contributors.

One area of potential abuse of information technology is in intelligence—spying on individual citizens and other computer systems. In 1988, Canadian newspapers obtained a report by Atomic Energy of Canada on its computerized data base tracking the actions of environmental groups [13]. The report also outlined plans for obtaining unauthorized access to other data bases. At around the same time, break-ins occurred at the offices of a Member of Parliament and a number of environmental groups [14-16]: “Computerized records were taken but valuable computer equipment ignored. ... ‘They took seven entire computer systems and left 25 wires dangling,’ said the network’s director [17].

There are many questionable uses of computers in this one government-related example. The work done by computer professionals in South Africa has even more direct consequences for human welfare. As this is

being written, ordinary people are working conscientiously at keyboards in the banking systems, the governments, universities, and software companies, all upholding the Apartheid regime. These are ordinary, well-educated people, who go home at night to their families. They are not individually malicious, but are still co-opted into maintaining a society where other human beings are systematically starved, dehumanized, and deprived of education, health care, and other basic human rights. Recent legislative changes may improve this situation, but so far the injustice has continued.

Computer technology may also adversely effect human welfare through military weapons use. One of the first computer professionals to recognize this was Norbert Wiener, the developer of cybernetics [18–20]. A substantial portion of government research (in North America at least) is through military agencies [4, 21, 22]. This involves a broad cross-section of scientists and researchers who have little or no control over how their published work is subsequently developed or used.

The greatest threat of computers in the military is in nuclear weapons systems. A war fought with nuclear weapons would constitute a human and environmental disaster. Such a war would not be possible without computers and computer professionals. Computer professionals contribute to the preparations for nuclear war in at least four ways: 1) computers and the professionals who operate them are essential components of the early warning and command and control systems for nuclear weapons. Malfunctions in these systems may be catastrophic [23, 24], yet in an 18-month period in 1979 – 1980 alone, the U.S. Senate Armed Services Committee reported 151 “serious” false alarms, and 3,703 others [25]; 2) computer professionals help devise and use computer simulations of nuclear war—so-called “war games” [26]. While computer game simulations are designed to alert officials to the uncertainties and complexities involved in the actual use of nuclear weapons, some analysts have expressed concern that this makes preparations for nuclear war routine [26]; 3) computer professionals may obtain scientific results with eventual applications to nuclear weapons. Scientists conduct basic research without knowing how it will be used; and 4) the most direct way in which computer scientists “up the stakes” for global destruction is in the actual design and development of nuclear weapons and missile guidance systems.

Therefore, computer professionals can do work of the highest technical quality, yet be isolated from the potential human costs. Even those computer professionals who have no direct involvement with these or other forms of collective violence should nevertheless be concerned about the problem, since they and their families are potential victims.

HOW DESTRUCTIVE PROFESSIONAL WORK IS JUSTIFIED

It is disturbing and regrettable to have to consider violent images and atrocities in relation to our everyday, comfortable lives. But perhaps recognizing the problems, and that the corporations and government agencies we work for have vested interests independent of human needs, is the first step in differentiating economic practicalities from values and human welfare.

In the preceding section, we showed how apparently legitimate work routines can threaten human welfare in the most inhumane ways. Therefore, it is logical to ask how well-meaning individuals perceive their role in the profession. Psychological and social mechanisms related to this are listed in Table 1. This is not necessarily intended to be the definitive taxonomy or to cover every possible example, but it should provide a useful summary of processes that may be new to the computer professional. These have been defined from the few [27–29] case studies, autobiographies [30], ethnographies [31] and related theoretical works [3–35].

We have attempted to list mechanisms which are applicable in many different situations. These have been classified as 1) general psychological processes, 2) processes specific to work in large bureaucracies and organizations, and 3) mechanisms that allow scientific and technological work independent of social values.

Table 1. Mechanisms That Could Maintain Conflict Between Job Actions and Personal and Social Values in Work with Computer Technology

Psychological mechanisms	Dissociation Rationalization
Organizational factors	Compartmentalization Hierarchical authority structure Amoral rationality
Facilitating factors in science and technology	Technological curiosity Distancing effects

Psychological Mechanisms

The mind is capable of playing subtle tricks on us. We do not always take the most rational alternative, or pay equal attention to equally important information. Therefore, we are susceptible to the following psychological mechanisms in many different types of dilemmas.

Dissociation. This involves a separation of different parts of conscious knowledge. The effect is to continue thinking and cognitive functioning by isolating incapacitating feelings and emotional responses [29]. It prevents full awareness of disquieting or unsettling information. Lifton and Markusen [29] state that this may ultimately involve “doubling” of one’s personality, as if separate roles or personalities develop for more and less humane behavior. It may be invoked when a role, at work begins to contradict one’s personal role [36].² As an illustration, Del Tredici [37] recorded the following dialogue with the spouse of a nuclear plant worker:

‘He was just real happy about being hired at Rocky Flats. We were a young couple, expecting a family, and the benefits were very good. The pay was great—you get what they call “hot pay” for working with radiation, so that’s why he wanted the process operator’s job...’

Did Don ever talk to you about the fact the he was making bombs?

‘He never did go into that’ ([37], pp. 173–174).

Several other authors have also described dissociation [1, 31, 38]. A similar procedure is often used in everyday life, e.g., when conscious attention is not used in an activity such as driving, changing gears, or locking a door. We can then devote complete attention to something else, such as an ongoing conversation (although we may later find ourselves wondering whether we actually locked that door).

“Psychic numbing” is a type of dissociation. Lifton [39] documented this in nuclear survivors in Hiroshima. He argues that in the nuclear age, it functions to mask the threat of instant extinction in our daily lives. Ironically, it operates in perpetrators as well as victims, and may allow either to shut out recognition of brutality.

Rationalization. This involves after-the-fact explanations of actions. Festinger developed a theory explaining how a post hoc shift in attitudes results from “cognitive dissonance” [40]. When we become aware that our actions contradict our values, we may rearrange our values after the fact to reduce inconsistency. When we are drawn into taking risks, we may adjust our beliefs about the likelihood of negative outcomes. This style of justification for one’s actions is typified by commonly-heard explanations commonly-heard explanations for why a particular project was accepted: “Better I do this than someone else”; “If I don’t do this, someone else will.”

² It should also be recognized that many individuals would not report any conflict between their personal values and job actions. We are interested in cases, however, where the individual has a vested interest in carrying out organizational goals independent of social values. The psychological mechanisms outlined show how conflict between vested work interests and values can then be obscured.

Bureaucratic Factors

Most computer scientists work within bureaucracies, often as specialists on sections of large projects. People who work in large organizations are susceptible to the following ways of separating work and values.

Compartmentalization. A diffusion of responsibility tends to occur naturally with complex technology, since technological work relies on numerous different specialists [35]. Therefore, most individuals have only small parts in the ultimate product, for which they do not feel responsible. (There are also situations in which a compartmentalized product is benign, but could be developed in future for either beneficial or harmful applications.) Lempert [27] reports interviews with four engineering students with summer jobs at Lawrence Livermore (nuclear weapons) Labs: “All four seemed to agree that in only a few months one could not possibly make a large enough contribution to feel one had personally helped to develop new nuclear arms” ([27], p. 63). This type of perception then leads to logic of the following sort: “I only ____, I don’t actually use them. ” One may fill in the blank with any application: “*write viruses,*” “*assemble the weapons,*” etc.

Although the division of labor in a large project may contribute to knowledge compartmentalization, it may also be the case that the “big picture” is purposely withheld. Diffusion of responsibility is explicit in cases of military compartmentalization for security reasons [30]. This was true of the thousands of people who moved to the Hanford nuclear reservation for a “top secret” project in the 1940s [31]. Soviet scientist and dissident Andrei Sakharov also noted this in the case of Soviet military research: “I was thankful that I was not told everything, despite my high-level security clearance” ([41], p. 268). However, in military or civilian work, compartmentalization and diffusion of responsibility lead to situations in which no one seems to actually have responsibility, as illustrated by three examples of work that is heavily reliant on computer technology:

It’s not like I’m designing the weapons. The guys who design them are in physics. *An engineer at Lawrence Livermore (nuclear weapons) Labs* ([27], p. 63).

Savannah River is the only facility that is producing weapons-grade plutonium to the defense programs. It is also the sole source of tritium. But we don’t have anything to do here with the actual fabrication of weapons. *James Gaver, Public Relations Officer for the U.S. Department of Energy, Savannah River Plant, North Carolina* ([37], p. 141).

Sandia’s role in the U.S. nuclear weapons program extends from applied research through development of new weapons and evaluation of their reliability throughout their stockpile lifetimes. We do not manufacture or assemble weapons components ... Sandia does not produce weapons and components. *Sandia National (nuclear weapons) Labs* ([42], p. 5).

A hierarchical authority structure. In a classic study of obedience, Milgram [32] told individuals in an experiment to administer electric shocks to people making mistakes on a learning test. He found that individuals would follow orders from a stranger to what they thought were life-threatening extents (see update and social applications in Kelman and Hamilton [33]). Although computer professionals in most contemporary jobs do not receive explicit orders (except in the military), there can still be penalties for not following procedures and instructions from superiors: these include implicit sanctions such as loss of status, or the possibility of being passed over for promotion [30]. The hierarchical authority structure is usually quite clear in most organizations.

It is sometimes argued that technicians and computer professionals should leave decisions about ethics and values to government leaders. Individual employees are not elected, and not authorized to make autonomous decisions affecting policy [43, 44]. However, this does not recognize the expertise of those directly involved in a particular project. This logic leads to what Johnson calls the “guns for hire” doctrine [45]. This view suggests that computer professionals should let society regulate what is acceptable through government representatives. Noting that the government cannot always be trusted to provide objective information, however, Sussman [46] states that our “leaders’ deliberate avoidance of true debate, the contempt they show the public during political campaigning, their use and refinement of propaganda

techniques, the attentiveness of so many of them to moneyed interests and not to the people generally, are all major causes of resentment and distrust” ([46], p. 49).

Amoral rationality. This is a preoccupation with procedural and technical aspects of work, while ignoring its moral, human, and social implications. The focus is on how to best do a job, with little attention to broader values and social effects. Responsibility for the work is perceived to be limited to technical aspects. In the Nazi death camps, amoral rationality allowed health professionals to serve as professional killers. Lifton reports that “an S.S. doctor said to me, ‘Ethics was not a word used in Auschwitz. Doctors and others spoke only about how to do things most efficiently’ ” [12, p. 294]. Albert Speer, Minister of Armaments and War in the Third Reich and a primary director of slave labor, directly addressed this in a 1944 note to Hitler: “The task that I am to perform is unpolitical. I have felt very good about my work so long as both I and my work were evaluated purely on the basis of my professional performance” ([38], p. 3; [47]). Wooten refers to this as a system of amoral functionalism, “one essentially devoid of morals and ethics in its decision-making process and one concerned only with *how* things get done and not *whether* they should get done” ([48], p. 21; emphasis in original). Computer science can be similarly promoted as highly technical, but independent of value considerations.

Once more fundamental social considerations are recognized, it becomes apparent that these questions must be addressed first. As the inventor of the hydrogen bomb in the Soviet Union, Sakharov notes that

Our reports, and the conferences where we discussed a strategic thermonuclear strike on a potential enemy, transformed the unthinkable and monstrous into a subject for detailed investigation and calculation. It became a fact of life—still hypothetical, but already seen as something possible. I could not stop thinking about this, and I came to realize that the technical, military, and economic problems are secondary; the fundamental issues are political and ethical” ([41], p. 268).

It will be argued in the final section that this way of thinking is reflected in codes of professional ethics and in educational curricula on science and technology.

Facilitating Factors in Science and Technology

These are processes encountered in professions based on science and technology. Again, they are distorting mechanisms that separate individual value judgements from the collective effects of work.

Technological curiosity. Regardless of the overall consequences, intelligent computer systems can be inherently interesting and can distract the worker from thoughts about the ethical implications of his or her work. Chalk describes a “primitive fascination” [20] with new technology (also see [27]). Since any type of basic research has by its nature no direct application, this must be a primary motivation for work on many scientific projects. Lifton and Markusen [29] discuss this general “passion for problem solving” in the work of nuclear physicists. Hayes [49] argues that work has changed as it has become more technology based; this may be due in part to this curiosity. “What mattered was the product’s capacity to provide more interesting work—a capacity that usually dovetailed with the corporate concern for profitability.” However, “among computer professionals, work was so self-referential, so thoroughly personalized, that it no longer required a public rationale in order to yield meaning” ([49], p. 32).

Distancing effects of technology. By operating as an intermediate processor in some situations, computers make eventual effects seem more distant. Just as pilots dropping bombs are removed from the human suffering that results, computers can remove the human initiator even more from personal involvement. This can occur in time, with contributions to a project or product to be implemented at a later date. A situation more unique to the computer industry, though, is where the human operator is present at the same point in time, but simply removed from the decision-making process: a preplanned procedure is carried through with automated control. (Note that bureaucracies also serve to distance policy makers from front-line effects, and front-line workers from responsibility for policies.)

WHY DESTRUCTIVE PROFESSIONAL WORK OCCURS: A PREDICTIVE MODEL

Taken together, these mechanisms can result in a situation where many highly-trained people work on projects that ultimately have very large human costs. Use of mechanisms such as these could be reinforced by socialization and professional training [30]. Recruitment, selection, and promotion may all depend on one's ability to go along with routines unquestioningly. The atmosphere in many settings may not allow open discussion of the effects of a project on society and on human welfare, and may emphasize distinct roles and hierarchies (e.g., with the use of uniforms or titles).

These mechanisms are factors affecting or in response to decisions we make. However, it is not the mechanisms per se that cause contributions to collective violence. For example, although obedience to a higher authority is often cited as a cause of irresponsible individual behavior [32, 33], we make autonomous decisions before following orders. We are not reflexively and automatically obedient to any higher authority (although we may decide that it is in our interest to be obedient). As another example, dissociation can not fundamentally explain behavior in dilemmas at work. We dissociate as a result of an earlier decision or an event. It is not dissociation that causes computer professionals to work on weapons of mass destruction; rather, they may do so because of practical employment needs, but then dissociate knowledge of destructive effects. To better explain these underlying causes, we will now present a predictive model. It explains why we contribute to large-scale risks that are not in our own or society's long-term interests, and therefore why mechanisms such as psychic numbing, rationalization, and obedience are needed.

It seems fundamental to the human condition that although we espouse certain values, individual actions ultimately come down to economic practicalities. For example: "Marie is a mother of two living in a small village in Vichy France in 1941 under Nazi control. Everyone is hustling for a position in the new regime, a pass for curfew, a bit of meat; resistance is not an option ..." [50]. The demands of daily living [51] were a priority for survival, and still figure prominently in many cases. But even when extreme affluence is attained, the focus on self-interest in the short term does not change. We can see the same process in the following biographical note on a defense electronics executive:

RAYTHEON. Thomas L. Phillips, Lexington, Mass. 617- 862-6600. SALES: \$8.8 bil. PROFITS: \$529 mil. Career path—engineering/technical; tenure—42 years, CEO 22 years. Compensation: 1989 salary & bonus, \$1,215,000; ownership, 136,000 shares. Not fretting about defense cuts, thanks to his electronics, commercial businesses, now 40% of sales ... One soft target: \$40 billion Milstar communications satellite—for use after nuclear war. Scheduled to retire at yearend to enjoy New Hampshire lakefront home [52].

Of course, wealth is not unethical in and of itself. But certainly when profiting from nuclear war, it is reasonable to wonder how justifications, vested interests, and psychological mechanisms are related. Obviously, day-to-day practicalities for this business executive do not mean actual survival, as they did for the oppressed mother in Nazi-occupied France. In both cases, though, there are immediate, tangible incentives for individuals to contribute to a system in which maximizing their own interests adds to the risk of harm for others later on.

The incentives for decisions that we are faced with can be defined in terms of a number of interacting parameters, such as the value of different alternatives, the probability associated with each alternative, and the type of each alternative [53]. In computer work, one might have to decide between

1. developing a profitable computer project with a 10% chance of eventual misuse or failure, or
2. not developing this project, therefore creating no chance of misuse or failure but possibly incurring negative consequences for one's job.

Note that the two alternatives differ in both probability (0% vs. 10%) and value (profit vs. negative consequences). The value can be conceptualized as coming in positive (reinforcing) or negative (punishing) forms. Either type can elicit behavior, although positive incentives are much more desirable. For example, a programmer would obviously rather work for intellectual or monetary rewards, than because he or she was forced to under threat of penalty (e.g., by an oppressive government, or simply because of monetary losses).

Parameters such as the value or magnitude of rewards and punishments tend to be relative, rather than absolute. For example, the difference we perceive between \$20 and \$30 is likely to be seen as more valuable than the difference between \$1,020 and \$1,030 (also a difference of \$10). The interesting thing for

dilemmas faced by computer professionals, though, is not a choice based on the perceived value of a single dimension. In alternatives where two parameters interact, each parameter has to be weighed, and trade-offs evaluated. Therefore, the computer professional may be faced with choosing between a profitable but low-probability project, for example, or one which offers less profit but a better chance of success.

Another important parameter in the subjective value of different alternatives is time delay. A basic principle of learning theory is that as the delay of a reward increases, its value decreases. Just as the subjective value of an additional \$10 varies according to whether it is in the context of \$20 or \$1,020, \$10 received now is likely to be seen as preferable to \$10 received tomorrow. This in turn has more value than a promise of \$10 or more in five weeks. Interestingly, we can obtain the relative importance of magnitude and time delay by asking how much money *would* be equally valuable: “Would you take \$12 tomorrow instead of \$10 now?” “Would you take \$30 in five weeks instead of \$10 now?” Regardless of the actual value in dollars, the psychological value is thus a nonlinear function of time ([54]).

Magnitude and time delay trade off in a predictable manner, although some irrational decisions are produced that do not maximize benefits, as will be discussed below. Rachlin notes the disproportional increase in value of some jobs initially because of this: “In the army ... you get an enlistment (or reenlistment) bonus so that the delay between signing up and your first pay check is very short” ([53], p. 142). Even advertisements for military service stick to payoffs that are both in one’s self-interest and immediate: “travel ... summer employment ... interesting people ... earn extra money ... build on your career ... part-time adventure” [55]. Recruiting has historically appealed to broad patriotic and nationalist values, but these are apparently not as marketable as early pay checks and the promise of more and earlier money, friends, adventure, and jobs. This situation is not unlike that of many computer professionals, for whom a fundamental motivation for many work decisions is economic: the need for a job that satisfies day-to-day needs [51].

A specific model, based on “social traps” [56, 57] relates incentives for individuals in their jobs to larger collective effects. As is true of all traps, a social trap presents an enticing opportunity, or bait. Like a more tangible trap, a social trap is a situation in which one choice that seems beneficial carries with it other negative consequences. Baron [58] emphasizes that this model is fundamental to dilemmas in many social situations:

Because so many situations can be analyzed as social dilemmas, much of the philosophy and psychology of morality is contained in this problem ... If everybody lies, we will not be able to depend on each other for information, and we will all lose. Likewise ... cheating on one’s taxes (making the government spend more money on enforcement), building up arms stocks in the context of an arms race, accepting bribes, polluting the environment, and having too many children are all examples ([58], pp. 399–400).

Two different types of traps can be defined, both of which are based on conflicting alternatives. Strictly speaking, “social” traps, or social dilemmas, apply only to a choice between self-interest and broader social or group interests (e.g., [59]). This model has been formally tested in laboratory simulations of conflict and cooperation between individuals and between countries [60]. However, there has been practically no attempt to collect empirical data or quantitatively model choices between self- and group interests in real life individual dilemmas, whether political, occupational or ethical.

“Temporal” traps could also be defined, for conflict between an immediate, short-term incentive, and a later one. The significance of these choices is that one has to wait to obtain the preferable alternative. Experiments with children on delay of gratification have identified cultural and personality variables affecting self-control [61], although the process of weighing different alternatives in decisions is more directly relevant in the present context. Quantitative models have been developed in numerous studies on animal learning defining tradeoffs between parameters such as the magnitude and delay of rewards [62–64]. Nevertheless, until now there have been very few attempts to apply these to the dilemmas that people face.

Table 2 shows how decisions using types of incentives and time delays can either maximize gains or lead to disastrous outcomes. Structuring social traps in this way allows predictions of how and when irrational decisions will be made. The table shows the preferred choices in four different sets of circumstances. For the computer professional, positive benefits shown on Table 2 related to the ethical context being discussed here might include (roughly in increasing order of importance): receiving praise for work well done, getting a raise, obtaining a well-paying job, making a positive contribution to the employer, producing a computer component that contributes to international stability and prosperity, or other contributions to human welfare. Similarly, negative consequences include: receiving a poor evaluation, losing one’s job, not being able to

support oneself or one's family, contributing to a harmful weapons system, or participating in collective violence. It should be noted, however, that for the purposes of the model, the importance of specific costs and rewards will be perceived individually; those specified here serve only as possibilities. Table 2 shows types of choices that computer professionals and others have to make and offers predictions about circumstances that may lead to irrational decisions. Four different sets of choices are shown; we will use choice C, a decision between a small cost immediately or a much larger cost in the future, for illustration here. From the above list of possible costs relevant to a computer professional, this decision may be between negative job effects now (e.g., poor evaluation, unemployment) and, say, the development of a weapon of mass destruction. This model would predict the final decision by measuring the psychological value of each alternative and scaling these as a function of time delay to obtain a total subjective value for each alternative based on the trade-off of time and value. The alternative with the greater subjective value is then chosen. In C the effect of time means that the predicted preference is not the one with the most benefit (least cost). Thus, negative job effects such as unemployment may be given more weight than contributing to future collective violence.

Table 2. Effects of Time Course and Value of Alternatives on Decision Making

Which would you rather have: a small but immediate reward or a larger delayed one?		
Now	Later	Net outcome
A. Small benefit (choice preferred)	Large consequence	Rational
B. Small benefit (choice preferred)	Large benefit	Mistake
C. Small consequence	Large consequence (choice preferred)	Mistake
D. Small consequence	Large benefit (choice preferred)	Rational

The psychological literature reviewed suggests that the computer professional tends to maximize short-term gains. Therefore the preferred choice is the rational one in only half of the cases (A and D). In situations configured as social or temporal traps, this means later suffering the consequences.

For individuals in single-industry towns, the practicality of having to avoid the consequences of unemployment may be much more salient than the possibility of producing a weapon that fuels the arms race [65–67]. Moreover, the weightings that we subjectively give to immediate, local needs over global consequences at some point in the future can be rationalized or overlooked with many of the psychological mechanisms discussed earlier. From interviews with computer professionals, physicists, and engineers working on nuclear weapons, Lempert [27] has noted the motivation that short-term economic needs provides: “in a tight job market, a young man or woman with a newly-earned degree might abandon a primary academic interest for a tempting salary” ([27], p. 62).

It should be clear that some of our decision preferences may be short sighted, and lead us into traps in which there are much larger consequences to suffer. It is also important to emphasize, however, that this model of social-temporal traps does not specify that individuals always choose the short term. Rather, decisions involve weighing the parameters of each alternative and evaluating trade-offs. With other things being equal, the short-term incentive will have greater perceived value.

Looking at decision making in terms of social and temporal traps is useful for explaining work behavior at all levels of organizational hierarchies. How does the data entry operator perceive and weigh conflicting responsibilities or interests? The model is equally applicable to the executive policy maker.

Although many of the problems of sustainability that we face at the end of the 20th century relate to institutions, organizations, industry, and so on, ultimately these are all made up of individual people. In affirming the importance of individuals and the collective effects of their work, Baron [58] has noted that

the problems caused by the existence of social dilemmas are among the most important that human beings have to solve. If we could learn ways to cooperate, wars would disappear and prosperity would prevail ... more cooperation

would solve many other human problems, from conflicts among roommates and family members to problems of protecting the world environment” ([58], pp. 403–404).

PRACTICAL APPLICATIONS TO ETHICAL DECISION MAKING

The psychological model and collective effects outlined here suggest that the wheels of the technological machine may be powered more by short-term economic interests and psychological, organizational, and technical mechanisms than by actual scientific or social needs (to say nothing of moral and ethical concerns). This can lead to devastating human costs on a world-wide scale. As Bandura [35] notes,

Given the many psychological devices for disengagement of moral control, societies cannot rely solely on individuals, however honorable their standards, to provide safeguards against inhumanities. To function humanely, societies must establish effective social safeguards against moral disengagement practices that foster exploitive and destructive conduct ([35], p. 27).

In view of this process, then, what practical alternatives are there to facilitate the choice of the right overall decision, rather than simply the one with immediate rewards?

Professional Codes of Ethics

Professional codes of ethics are one method through which short-term self interest could be balanced with broader alternatives. These codes exist in hundreds of professional societies [68, 69], as well as in some universities [70] and university departments [71]. In computing, codes exist for professional associations such as the Institute of Electrical and Electronics Engineers (IEEE), the Association for Computing Machinery (ACM), and the Canadian Information Processing Society. A number of codes have also existed in related areas, dating to before the advent of computers, e.g., the Code of Principles of Professional Conduct of the American Institute of Electrical Engineers (1912).³

Codes such as these have the potential to objectively structure an ethical dilemma for rationally evaluating possible alternatives, with ethical implications for each alternative clearly laid out. The basic purpose for all codes of ethics is to ensure that work has moral integrity and is for the public good. For example, the IEEE code is very germane to collective professional violence in specifying that its members “accept responsibility in making engineering decisions consistent with the safety, health and welfare of the public and to disclose promptly factors that might endanger the public or the environment.” The ACM code similarly makes reference to “the advancement of human welfare.” Codes are often oriented to protect consumers and society from conflicts of interest when the professional is in a position of power.

Two critical reviews of the actual effectiveness of codes of ethics have been carried out [68, 72]. Unfortunately, in codes of ethics generally, ideals such as “honorable” and “the public good” tend to be undermined by being open to multiple interpretations [68,72]. Codes may be particularly amenable to serving government aggression if they do not clearly differentiate human welfare from national welfare, and clearly define how “the public good” relates to these.

Many professional codes of ethics apply only to individual abuses, with no consideration of institutionalized destructiveness [68]. That is, they prohibit unethical behavior by one individual, but do not address unethical policies, professional practices, or committee actions. Codes also typically emphasize procedures and technical issues [68] (in their own form of selective attention and amoral rationality). For instance, the Rensselaer Polytechnical Institute code of ethics for computing is not unusual in focusing on procedural violations such as:

6. Bypassing accounting mechanisms;
7. Violating copyright or licensing agreements ...
8. Deliberately wasting computer resources (e.g., printing blank pages or unnecessary copies).

³ Copies of these and other professional codes of ethics can be obtained from C. S. or from archives such as Center for the Study of Ethics in the Professions, Life Sciences Building, Illinois Institute of Technology, 3101 S. Dearborn Street, Chicago, Illinois, 60616-3793.

Fundamental ethical questions are not routinely discussed, but should precede consideration of these types of how-to's and procedural do's and don'ts. No professional codes of ethics contain statements on ethical justifications for weapons development or professional involvement in wars or killing, for example [68]. Ethical considerations for the computer professional typically deal with what you do after you sit down at a terminal. However, an initial consideration should be why one is sitting down at the terminal in the first place.

Effective codes of ethics at both the fundamental and procedural levels serve two complementary purposes. First, they protect consumers and subgroups of society from institutionalized destructiveness by encouraging professional activities that are in the interests of human welfare. Second, the various types of ethical guidelines can protect the individual professional who receives instructions to carry out questionable or unconscionable institutional goals [20]. Employees in this situation can "pass the buck," deferring responsibility for their inability to serve the company or government agency to an objective, often (and preferably) international code of ethics.

Codes of ethics can similarly facilitate whistleblowing, making clear, as the 1990 IEEE code does, that this type of criticism may be in the public interest. The IEEE code encourages "disclosing promptly factors that might endanger the public," although no actual protection is mentioned in the code itself. Johnson criticizes a statement on whistleblowing in the ACM code for not ensuring that action is taken [72].

It is notable that although codes of ethics are oriented towards human welfare, a computer-intensive organization involved in developing nuclear weapons has no formal, written ethics policies. Los Alamos National Labs has no code of ethics or even ethical guidelines (R. Glasser, Los Alamos National Labs, personal communication). This has the double-pronged implication that society and (international) human welfare are not protected from the technology developed in the weapons labs, but also that employees have limited recourse when directed to carry out any unethical or unconscionable project. Of course, it may not be useful for nuclear weapons labs to have codes, when engineering and computing associations do already. But certainly many corporations employing computer professionals and engineers have their own specific codes; even Martin Marietta Energy Systems, which is extensively involved in nuclear weapons research and development, has such a code [68]. This is particularly necessary in work where an employee does not need to be accredited by or have membership in a professional association having a code. Having several overlapping codes from professional associations, companies, and possibly the government should not be a problem, particularly if basic principles such as human welfare can be prioritized in case conflicting guidelines are encountered. Similar procedures must already be followed within any detailed, prescriptive code when two guidelines conflict.

Other Policy Implications

Education is another area in which a narrow, technical focus may be established, similar to the psychological process of amoral rationality. Some professional associations in computing do have statements on the inclusion of ethics in computer science curricula. However, as in other areas in science, there may be an implicit assumption in many textbooks and lecture halls that the process of advancing knowledge through research and development is value free. One way to put computer science in a broader social context would be for textbooks to mention that some of the research they review has been funded for military purposes, or that there are ethical questions surrounding a technical issue being presented. This would provide a more complete education. Ethical questions should not be compartmentalized in specialty textbooks if we want to avoid the psychological and bureaucratic mechanisms discussed earlier.

Government policy initiatives also provide a direct way of mediating vested short-term interests in a particular type of work. For a brief period in 1990 there was talk of a "peace dividend" as a result of the end of the Cold War and the political restructuring in Eastern Europe. Massive military expenditures would no longer be needed, and could be redirected to immediate human needs: better schools, road repairs, funding for research, foreign aid, lowered taxes. Social trap models allow governments to anticipate resistance when self interests conflict with demilitarization. To begin implementing a policy of economic conversion, incentives must be provided to meet short-term economic needs, both at individual and organizational levels.

There are also several specific implications from the experimental research on incentives. The salience of long-term goals and benefits can be heightened in several ways. Simple periodic reminders may be effective. This is the effect of the Surgeon General's warning against the immediate rewards of smoking.

Also, making an early commitment to alternatives delayed in time is effective [73]. Short-term alternatives may seem more attractive. With an early commitment to one choice, however, both alternatives are long term, and the time delay has less effect than the actual values of the possible outcomes.

SUMMARY AND CONCLUSIONS

Organizations such as governments, companies, and the military involve many professionals, but can have goals independent of human needs. Because of the role computer technology now plays in any large project, computer professionals may face ethical decisions between organizational interests and social values. Unfortunately, if there are vested job interests, the reliance on higher authority, regular routines, and technological curiosity may support amoral rationality: do a good job technically, but leave responsibility to the larger organization. Because of this process, professionals have been participants in collective violence.

Social and temporal traps provide a useful framework for evaluating the role of individuals in collective violence. These models look at the value and timing (delay) of the alternatives in a decision. Lawful predictions can then be made for both rational and shortsighted behavior. This approach has the advantage of applying to individuals at all levels of organizational hierarchies, and in many different situations.

Finally, in response to the conflicting interests that may arise for computer professionals, there are several approaches that may help to structure and prioritize the alternatives. Professional codes of ethics, education, and government policies may all facilitate choices that provide benefits individually *and* collectively.

REFERENCES

1. E. Markusen, Professions, professionals and genocide, in *Genocide: A Critical Bibliographic Review* (I. W. Charney, ed.), Facts on File Publishing, New York, 1991.
2. P. R. Ehrlich and A. H. Ehrlich, *Extinction: The Causes and Consequences of the Disappearance of Species*, Ballantine Books, New York, 1985.
3. R. Rummel, *Lethal Politics: Soviet Genocide and Mass Murder Since 1917*, Transaction Books, New Brunswick, New Jersey, 1990.
4. R. Sivard, *World Military and Social Expenditures 1987-88*, World Priorities, Leesburg, Virginia, 1988.
5. G. Dyer, *War*, Crown, New York, 1985.
6. R. O'Connell, *Of Arms and Men*, Oxford University Press, New York, 1989.
7. World Health Organization, *Effects of Nuclear War on Health and Health Services*, World Health Organization, Geneva, 1984.
8. C. Sagan and R. Turco, *A Path Where No Man Thought: Nuclear Winter and the End of the Arms Race*, Random House, New York, 1990.
9. E. Markusen, Genocide and total war, in *Genocide and the Modern Age* (I. Wallimann and M. Dobkowski, eds.), Greenwood Press, New York, 1987.
10. I. W. Charney, Genocide and mass destruction: doing harm to others as a missing dimension in psychopathology, *Psychiatry* 49,144-157 (1986).
11. R. L. Rubenstein and J. Roth, *Approaches to Auschwitz*, John Knox Press, Atlanta, Georgia, 1987.
12. R. J. Lifton, Medical killing in Auschwitz, *Psychiatry* 45, 283-297 (1982).
13. I. Mumford, Memorandum (restricted commercial), *Atomic Energy of Canada Limited*, January 12, 1988.
14. R. Ludlow, N-industry spying, activists allege, *The Vancouver Sun*, Vancouver, British Columbia, Canada, July 19, 1988, p. A7.
15. Office break-ins might be tied to MP's report, *The Ottawa Citizen*, Ottawa, Ontario, Canada, March 15, 1989, p. A4.
16. P. O'Neil, Fulton's fears break-ins linked called paranoia, *The Vancouver Sun*, Vancouver, British Columbia, Canada, March 15, 1989, p. C8.
17. R. Cleroux, Burglars seek records, not cash from MP, environmental groups, *The Globe & Mail*, March 17, 1989, p. A11.
18. N. Wiener, A scientist rebels, *The Atlantic* 179, 46 (January 1947).
19. T. Winograd, CPSR president Winograd presents Norbert Wiener award to Parnas, *Computer Professionals for Social Responsibility Newsletter* 6, 10-12 (1988).
20. R. Chalk, Drawing the line, an examination of conscientious objection in science, in *Ethical Issues Associated with Scientific and Technological Research for the Military* (C. Mitcham and P. Siekeitz, eds.), *Ann. NY Acad. Sci.* 577, 61-74 (1989).
21. American Psychological Society, The importance of the citizen scientist in national science policy, *APS Observer*, 4, 10-23 (July 1991).
22. D. D. Noble, Mental materiel: the militarization of learning and intelligence in U.S. education, in *Cyborg Worlds: The Military Information Society*

- (L. Levidow and K. Robins, eds.), Free Association Books, London, 1989.
23. D. Ford, *The Button: The Pentagon's Strategic Command and Control System*, Simon & Schuster, New York, 1985.
 24. S. Gregory, Command and Control, in *The Greenpeace Book of the Nuclear Age: The Hidden History, The Human Cost* (J. May, ed.), McClelland and Stewart, Toronto, 1989.
 25. D. P. Barash, *Introduction to Peace Studies*, Wadsworth, Belmont, California, 1991.
 26. T. B. Allen, *War Games*, McGraw-Hill, New York, 1987.
 27. R. Lempert, Will young scientists build bombs? *Bull. Atomic Sci.* 37, 61-64 (1981).
 28. S. Kull, *Minds at War: Nuclear Reality and the Inner Conflicts of Defense Policymakers*, Basic Books, New York, 1988.
 29. R. J. Lifton and E. Markusen, *The Genocidal Mentality: Nazi Holocaust and Nuclear Threat*, Basic Books, New York, 1990.
 30. H. T. Nash, The bureaucratization of homicide, in *Protest and Survive* (E. P. Thompson and D. Smith, eds.), Monthly Review Press, New York, 1981.
 31. P. Loeb, *Nuclear Culture: Living and Working in the World's Largest Atomic Complex*, New Society Publishers, Philadelphia, Pennsylvania, 1986.
 32. S. Milgram, Some conditions of obedience and disobedience, *Hum. Relat.* 18, 57-76 (1965).
 33. H. C. Kelman, and V. L. Hamilton, *Crimes of Obedience: Toward a Social Psychology of Authority and Responsibility*, Yale University Press, New Haven, Connecticut, 1989.
 34. E. Staub, *The Roots of Evil: The Origins of Genocide and Other Group Violence*, Cambridge University, New York, 1989.
 35. A. Bandura, Selective activation and disengagement of moral control, *J. Social Issues* 46,27-46 (1990).
 36. J. S. Coleman, *The Asymmetric Society*, Syracuse University Press, Syracuse, New York, 1982.
 37. R. Del Tredici, *At Work in the Fields of the Bomb*, Harper & Row, New York, 1987.
 38. J. H. Barton, Editorial, *IEEE Potentials* 5, 3 (1986).
 39. R. J. Lifton, *Death in Life: Survivors of Hiroshima*, Random House, New York, 1967 (also 1991, University of North Carolina Press).
 40. L. Festinger, *A Theory of Cognitive Dissonance*, Stanford University Press, Stanford, California, 1957.
 41. A. Sakharov, *Memoirs* (R. Lourie, trans.), Alfred A. Knopf, New York, 1990.
 42. *Sandia National Laboratories, A Report for A. T. & T. Shareholders*, A.T.&T, New York, 1990.
 43. K. W. Kemp, Conducting scientific research for the military as a civic duty, in *Ethical Issues Associated with Scientific and Technological Research for the Military*. (C. Mitcham and P. Siekevitz, eds.), *Ann. NY Acad. Sci.* 577, 115-121 (1989).
 44. K. Kipnis, Engineers who kill: professional ethics and the paramountcy of public safety, *Bus. Profess. Ethics J.* 1, 77-91 (1981).
 45. D. G. Johnson, The social/professional responsibility of engineers, in *Ethical Issues Associated with Scientific and Technological Research for the Military* (C. Mitcham and P. Siekevitz, eds.), *Ann. NY Acad. Sci.* 577, 106-114 (1989).
 46. B. Sussman, *What Americans Really Think, and Why Our Politicians Pay no Attention*, Pantheon Books, New York, 1988.
 47. A. Speer, *Trial of the Major War Criminals*, International Tribunal, Nuremberg, Germany, 1947. (Quoted from [38].) vol. XVI.
 48. L. M. Wooten, Albert Speer: how to manage an atrocity, *J. Hum. Psychol.* 21, 21-38 (1981).
 49. D. Hayes, *Behind the Silicon Curtain: The Seductions of Work in a Lonely Era*, South End Press, Boston, 1989.
 50. *Story of Women* (film, 1989). Advertisement.
 51. R. K. Gilbert, The dynamics of inaction: psychological factors inhibiting arms control activism, *Am. Psychol.* 43, 755-764 (1989).
 52. The corporate elite: the chief executives of the 1000 most valuable publicly held U.S. companies, *Business Week*, 55-274 (19 October 1990).
 53. H. Rachlin, *Judgement, Decision and Choice: A Cognitive / Behavioural Synthesis*, W. H. Freeman and Co., New York, 1989.
 54. B. S. Gorman, A. E. Wessman, G. R. Schmeidler, and S. Thayer, Linear representation of temporal location and Stevens' law, *Mem. Cognit.* 1, 169-171 (1973).
 55. Canadian Armed Forces, The reserve: part-time adventure (advertisement), *The Tribune-Post*, Sackville, New Brunswick, Canada, Oct. 17, 1990, p. 17.
 56. J. Platt, Social traps, *Am. Psychol.* 28, 641-651 (1973).
 57. G. Cross and M. J. Guyer, *Social Traps*, University of Michigan Press, Ann Arbor, Michigan, 1980.
 58. J. Baron, *Thinking and Deciding*, Cambridge University Press, New York, 1988.
 59. G. R. Hardin, The tragedy of the commons, *Science* 162, 1243-1248 (1968).
 60. A. Rappaport, Experiments with N-person social traps I, *J. Conflict Res.* 32, 457-472 (1989).
 61. W. Mischel, Preference for delayed reinforcement: an experimental study of a cultural observation, *J. Abnormal Soc. Psychol.* 56, 57-61 (1958).
 62. M. L. Commons, R. J. Herrnstein, and H. Rachlin, eds. *Quantitative Analyses of Behavior, Vol. II: Matching and Maximizing Accounts*, Ballinger Publishing, Cambridge, Massachusetts, 1982.
 63. H. Rachlin, R. Battalio, J. Kagel, and L. Green, Maximization theory in behavioral psychology, *Behav. Brain Sci.* 4, 371-417 (1981).

64. R. J. Herrnstein, Rational choice theory: necessary but not sufficient, *Am. Psychol.* 45, 356-367 (1990).
65. P. Sanger, *Blind Faith: The Nuclear Industry in One Small Town*, McGraw-Hill Ryerson Ltd., Toronto, 1981.
66. C. Giangrande, *The Nuclear North: The People, The Regions and The Arms Race*, Anansi, Toronto, 1983.
67. A. G. Mojtabai, *Blessed Assurance: At Home with The Bomb in Amarillo, Texas*, Houghton-Mifflin, Boston, 1986.
68. C. Summers, An international and historical comparison of codes of professional ethics, *Annual Convention of the Canadian Psychological Association*, Calgary, Alberta, Canada, 1991.
69. R. A. Gorlin, ed., *Codes of Professional Responsibility*, 2nd ed. The Bureau of National Affairs, Washington, DC, 1990.
70. *UCLA Faculty Handbook, Supplement A: Policies and Procedures*, University of California at Los Angeles Publications Office, Los Angeles, 1987.
71. Queen's University Department of Psychology, *Ethics Clearance for Research on Human Subjects*, Queen's University, Kingston, Ontario, Canada, 1989.
72. D. G. Johnson, *Computer Ethics*, Prentice-Hall, Englewood Cliffs, New Jersey, 1985.
73. B. Schwartz and H. Lacey, *Behaviorism, Science and Human Nature*, W. W. Norton & Co., New York, 1982.

Eisenhower's Farewell Address to the Nation

January 17, 1961

Good evening, my fellow Americans: First, I should like to express my gratitude to the radio and television networks for the opportunity they have given me over the years to bring reports and messages to our nation. My special thanks go to them for the opportunity of addressing you this evening.

Three days from now, after a half century of service of our country, I shall lay down the responsibilities of office as, in traditional and solemn ceremony, the authority of the Presidency is vested in my successor.

This evening I come to you with a message of leave-taking and farewell, and to share a few final thoughts with you, my countrymen.

Like every other citizen, I wish the new President, and all who will labor with him, Godspeed. I pray that the coming years will be blessed with peace and prosperity for all.

Our people expect their President and the Congress to find essential agreement on questions of great moment, the wise resolution of which will better shape the future of the nation.

My own relations with Congress, which began on a remote and tenuous basis when, long ago, a member of the Senate appointed me to West Point, have since ranged to the intimate during the war and immediate post-war period, and finally to the mutually interdependent during these past eight years.

In this final relationship, the Congress and the Administration have, on most vital issues, cooperated well, to serve the nation well rather than mere partisanship, and so have assured that the business of the nation should go forward. So my official relationship with Congress ends in a feeling on my part, of gratitude that we have been able to do so much together.

We now stand ten years past the midpoint of a century that has witnessed four major wars among great nations. Three of these involved our own country. Despite these holocausts America is today the strongest, the most influential and most productive nation in the world. Understandably proud of this pre-eminence, we yet realize that America's leadership and prestige depend, not merely upon our unmatched material progress, riches and military strength, but on how we use our power in the interests of world peace and human betterment.

Throughout America's adventure in free government, such basic purposes have been to keep the peace; to foster progress in human achievement, and to enhance liberty, dignity and integrity among peoples and among nations.

To strive for less would be unworthy of a free and religious people.

Any failure traceable to arrogance or our lack of comprehension or readiness to sacrifice would inflict upon us a grievous hurt, both at home and abroad.

Progress toward these noble goals is persistently threatened by the conflict now engulfing the world. It commands our whole attention, absorbs our very beings. We face a hostile ideology global in scope, atheistic in character, ruthless in purpose, and insidious in method. Unhappily the danger it poses promises



to be of indefinite duration. To meet it successfully, there is called for, not so much the emotional and transitory sacrifices of crisis, but rather those which enable us to carry forward steadily, surely, and without complaint the burdens of a prolonged and complex struggle – with liberty the stake. Only thus shall we remain, despite every provocation, on our chartered course toward permanent peace and human betterment.

Crises there will continue to be. In meeting them, whether foreign or domestic, great or small, there is a recurring temptation to feel that some spectacular and costly action could become the miraculous solution to all current difficulties. A huge increase in the newer elements of our defenses; development of unrealistic programs to cure every ill in agriculture; a dramatic expansion in basic and applied research – these and many other possibilities, each possibly promising in itself, may be suggested as the only way to the road we wish to travel.

But each proposal must be weighed in light of a broader consideration; the need to maintain balance in and among national programs – balance between the private and the public economy, balance between the cost and hoped for advantages – balance between the clearly necessary and the comfortably desirable; balance between our essential requirements as a nation and the duties imposed by the nation upon the individual; balance between the actions of the moment and the national welfare of the future. Good judgment seeks balance and progress; lack of it eventually finds imbalance and frustration.

The record of many decades stands as proof that our people and their Government have, in the main, understood these truths and have responded to them well in the face of threat and stress.

But threats, new in kind or degree, constantly arise.

Of these, I mention two only.

A vital element in keeping the peace is our military establishment. Our arms must be mighty, ready for instant action, so that no potential aggressor may be tempted to risk his own destruction.

Our military organization today bears little relation to that known by any of my predecessors in peacetime, or indeed by the fighting men of World War II or Korea.

Until the latest of our world conflicts, the United States had no armaments industry. American makers of plowshares could, with time and as required, make swords as well. But now we can no longer risk emergency improvisation of national defense; we have been compelled to create a permanent armaments industry of vast proportions. Added to this, three and a half million men and women are directly engaged in the defense establishment. We annually spend on military security more than the net income of all United States corporations.

This conjunction of an immense military establishment and a large arms industry is new in the American experience. The total influence – economic, political, even spiritual – is felt in every city, every Statehouse, every office of the Federal government. We recognize the imperative need for this development. Yet we must not fail to comprehend its grave implications. Our toil, resources and livelihood are all involved; so is the very structure of our society.

In the councils of government, we must guard against the acquisition of unwarranted influence, whether sought or unsought, by the military-industrial complex. The potential for the disastrous rise of misplaced power exists and will persist.

American makers of plowshares could, with time and as required, make swords as well. But now we can no longer risk emergency improvisation of national defense; we have been compelled to create a permanent armaments industry of vast proportions.

We must never let the weight of this combination endanger our liberties or democratic processes. We should take nothing for granted. Only an alert and knowledgeable citizenry can compel the proper meshing of the huge industrial and military machinery of defense with our peaceful methods and goals, so that security and liberty may prosper together.

Akin to, and largely responsible for the sweeping changes in our industrial-military posture, has been the technological revolution during recent decades.

In this revolution, research has become central, it also becomes more formalized, complex, and costly. A steadily increasing share is conducted for, by, or at the direction of, the Federal government.

Today, the solitary inventor, tinkering in his shop, has been overshadowed by task forces of scientists in laboratories and testing fields. In the same fashion, the free university, historically the fountainhead of free ideas and scientific discovery, has experienced a revolution in the conduct of research. Partly because of the huge costs involved, a government contract becomes virtually a substitute for intellectual curiosity. For every old blackboard there are now hundreds of new electronic computers.

The prospect of domination of the nation's scholars by Federal employment, project allocations, and the power of money is ever present – and is gravely to be regarded.

Yet, in holding scientific research and discovery in respect, as we should, we must also be alert to the equal and opposite danger that public policy could itself become the captive of a scientific-technological elite.

The prospect of domination of the nation's scholars by Federal employment, project allocations, and the power of money is ever present – and is gravely to be regarded.

It is the task of statesmanship to mold, to balance, and to integrate these and other forces, new and old, within the principles of our democratic system – ever aiming toward the supreme goals of our free society.

Another factor in maintaining balance involves the element of time. As we peer into society's future, we – you and I, and our government – must avoid the impulse to live only for today, plundering for, for our own ease and convenience, the precious resources of tomorrow.

We cannot mortgage the material assets of our grandchildren without asking the loss also of their political and spiritual heritage. We want democracy to survive for all generations to come, not to become the insolvent phantom of tomorrow.

Down the long lane of the history yet to be written America knows that this world of ours, ever growing smaller, must avoid becoming a community of dreadful fear and hate, and be, instead, a proud confederation of mutual trust and respect.

Such a confederation must be one of equals. The weakest must come to the conference table with the same confidence as do we, protected as we are by our moral, economic, and military strength. That table, though scarred by many past frustrations, cannot be abandoned for the certain agony of the battlefield.

Disarmament, with mutual honor and confidence, is a continuing imperative. Together we must learn how to compose differences, not with arms, but with intellect and decent purpose. Because this need is so sharp and apparent I confess that I lay down my official responsibilities in this field with a definite sense of disappointment. As one who has witnessed the horror and the lingering sadness of war – as one who knows that another war could utterly destroy this civilization which has been so slowly and painfully built over thousands of years – I wish I could say tonight that a lasting peace is in sight.

Happily, I can say that war has been avoided. Steady progress toward our ultimate goal has been made. But, so much remains to be done. As a private citizen, I shall never cease to do what little I can to help the world advance along that road.

So – in this my last good night to you as your President – I thank you for the many opportunities you have given me for public service in war and peace. I trust that in that service you find some things worthy; as for the rest of it, I know you will find ways to improve performance in the future.

You and I – my fellow citizens – need to be strong in our faith that all nations, under God, will reach the goal of peace with justice. May we be ever unswerving in devotion to principle, confident but humble with power, diligent in pursuit of the Nations' great goals.

To all the peoples of the world, I once more give expression to America's prayerful and continuing aspiration:

We pray that peoples of all faiths, all races, all nations, may have their great human needs satisfied; that those now denied opportunity shall come to enjoy it to the full; that all who yearn for freedom may experience its spiritual blessings; that those who have freedom will understand, also, its heavy responsibilities; that all who are insensitive to the needs of others will learn charity; that the scourges of poverty, disease and ignorance will be made to disappear from the earth, and that, in the goodness of time, all peoples will come to live together in a peace guaranteed by the binding force of mutual respect and love.

Now, on Friday noon, I am to become a private citizen. I am proud to do so. I look forward to it.

Thank you, and good night.

A longer and more scholarly version of this essay appears as “The Second Enclosure Movement and the Construction of the Public Doman.” It appears in *Law and Contemporary Problems*, 66, pp. 33–74, 2003 and is available at URL www.law.duke.edu/pd/papers/boyle.pdf. See www.law.duke.edu/boylesite/ for a short biography of Prof. Boyle.

Fencing off Ideas Enclosure & the Disappearance of the Public Domain

by

James Boyle

Dadealus, 131(2), pp. 13–25, 2002

The law locks up the man or woman
Who steals the goose from off the common
But leaves the greater villain loose
Who steals the common from off the goose.

The law demands that we atone
When we take things we do not own
But leaves the lords and ladies fine
Who take things that are yours and mine.

The poor and wretched don't escape
If they conspire the law to break;
This must be so but they endure
Those who conspire to make the law.

The law locks up the man or woman
Who steals the goose from off the common
And geese will still a common lack
Till they go and steal it back.

This poem is one of the pithiest condemnations of the English enclosure movement, the process of fencing off common land and turning it into private property. (Although we refer to it as “the enclosure movement,” it was actually a series of enclosures that started in the fifteenth century and went on, with differing means, ends, and varieties of state involvement, until the nineteenth.) The poem manages in a few lines to criticize double standards, expose the artificial and controversial nature of property rights, and take a slap at the legitimacy of state power. And it does it all with humor, without jargon, and in rhyming couplets. ...

The enclosure movement continues to draw our attention. It offers irresistible ironies about the two-edged sword of “respect for property” and lessons about the role of the state in making controversial, policy-laden decisions to define property rights in ways that subsequently come to seem both natural and neutral.

Following in the footsteps of Thomas More, critics have long argued that the enclosure movement imposed devastating costs on one segment of society. Some of these costs were brutally and relentlessly “material”—for example, the conversion of crofters and freeholders into peons, seasonal wage-laborers, or simply, as More argued in *Utopia*, beggars and thieves. But other harms were harder to classify: the loss of a form of life, and the relentless power of market logic to migrate to new areas, disrupting traditional social relationships, views of the self, and even the relationship of human beings to the environment.

A great many economic historians have begged to differ. As they see the matter, the critics of enclosure have fallen prey to the worst kind of sentimentality, romanticizing a form of life that was neither comfortable nor noble, and certainly not very egalitarian.

From an economist's point of view, the key fact about the enclosure movement is that it worked: this new property regime allowed an unparalleled expansion of productive possibilities. By transferring inefficiently managed common land into the hands of a single owner, enclosure averted one aptly named "tragedy" of the commons: overuse. It also created incentives for large-scale investment, allowed control over exploitation, and in general ensured that the resource would be used efficiently. Unless the feudal lord knew that the fruits of his labor would be his alone, he would not have invested in drainage schemes, the purchase of sheep, or the rotation of crops in order to increase the yield of his acreage.

Strong private-property rights helped to avoid the tragedies of both overuse and underinvestment. As a result of the enclosure movement, fewer Englishmen starved: more grain was grown, and more sheep were raised. If the price of this social gain was a greater concentration of economic power in fewer hands and despoliation of the environment, so be it. Those who weep about the terrible effects of private property should realize that it literally saved lives. Or so say the economic historians.

This is a debate of more than antiquarian interest, for we are in the midst of a new kind of enclosure movement, this one aimed at exploiting a new and intangible kind of commons—call it a "commons of the mind." (1) Once again, things that were formerly thought to be uncommodifiable, essentially common, or outside the market altogether are being turned into private possessions under a new kind of property regime. But this time the property in question is intangible, existing in databases, business methods, and gene sequences.

Take the human genome as an example. The opponents of "enclosure" have claimed that the genome "belongs to everyone," that it is literally "the common heritage of humankind." They say that the code of life ought not and perhaps in some sense cannot be owned by an individual or a corporation. When patents have been granted for stem cells and gene sequences, critics have mused darkly about the way in which the state is simply handing over monopoly power to private parties, potentially thwarting future research and innovation. The new monopolists have names like Geron, Celera, and Human Genome Sciences, and their holdings are in the form of patent portfolios rather than oil wells or steel plants.

Alongside these reports about the beneficiaries of the new property scheme run news stories about those who were not so fortunate, the commoners of the genetic enclosure. Law students across America now read *Moore v. Regents*, a California Supreme Court case deciding that poor Mr. Moore had no property right to a cell line derived from his spleen. In this case, the court decided that giving property rights to "sources" would make it more difficult for scientists to share cell lines with fellow researchers—reading the decision, one can almost picture the Styrofoam coolers criss-crossing the country by Federal Express in an orgy of altruistic flesh swapping. Yet this fear of the pernicious effects of property rights did not last for long. In another portion of the opinion the court speaks approvingly of the patent granted to the doctors whose inventive genius created a billion-dollar cell line from Mr. Moore's "naturally occurring raw material." Like the commoners, Mr. Moore finds that his naturalistic and traditional property claims are portrayed as impediments to innovation. Like the beneficiaries of enclosure, the doctors are granted a property right to encourage efficient development of a wasted resource.

Of course, like the first enclosure movement, this new one has its defenders. To the question "should there be patents over human genes?" the answer will be "private property saves lives." Only by extending the reach of property rights can the state guarantee the investment of time, ingenuity, and capital necessary to produce new drugs and gene therapies. Private-property rights are a necessary incentive to research; economists need only worry about how to allocate these rights most efficiently. Or so say the advocates of private-property rights.

The genome is not the only area to have been partially "enclosed" in the past decade. In recent years, intellectual property rights have been dramatically expanded in many different fields of human endeavor—from business method patents to the Digital Millennium Copyright Act, from trademark antidilution rulings to the European Database Protection Directive.

In 1918, the American jurist Louis Brandeis confidently claimed that “[t]he general rule of law is, that the noblest of human productions—knowledge, truths ascertained, conceptions, and ideas—become, after voluntary communication to others, free as the air to common use.” At the time that Brandeis made that remark, intellectual property rights were the exception rather than the rule; it was widely agreed that ideas and facts must always remain in the public domain. But that old consensus is now under attack. Long-standing limits on the reach of intellectual property—the antierosion walls around the public domain—are being eaten away each year.

The annual process of updating my syllabus for a basic intellectual property course provides a nice snapshot of what is going on. I can wax nostalgic looking back to a five-year-old text, with its confident list of the subject matter that intellectual property rights couldn’t cover, the privileges that circumscribed the rights that did exist, the length of time before a work fell into the public domain. In each case, the old limits have recently been changed or challenged.

Patents are increasingly stretched out to cover “ideas” that twenty years ago all scholars would have agreed were unpatentable: the so-called business method patents, which cover such “inventions” as auctions or accounting methods, are an obvious example. Most troubling of all are the attempts to introduce intellectual property rights over mere compilations of facts. If Anglo-American intellectual property law had an article of faith, it was that unoriginal compilations of facts would remain in the public domain. This was “no mere accident of a statutory scheme,” as the Supreme Court once put it: protecting the raw material of science and speech is as important to the next generation of innovation as the intellectual property rights themselves. The system would offer a limited monopoly for an invention or an original expression of ideas, but the monopoly was to be tightly confined to the layer of invention or expression. The facts below, or the ideas above, would remain free for all to build upon. Even the stuff that could be protected by intellectual property—the drug or the poem, say—was supposed to pass into the public domain after a certain number of years. As Jefferson and Macaulay both observed, intellectual property rights were necessary evils. They should be strictly limited in both time and extent.

Today, these traditional assumptions about intellectual property law are under attack. Some of the challenges are subtle. In patent law, stretched interpretations of novelty and nonobviousness allow intellectual property rights to move closer and closer to the underlying data layer; gene sequence patents come very close to being rights over a particular discovered arrangement of data—C’s, G’s, A’s, and T’s. Other challenges are overt; the European Database Directive does (and the various proposed database bills in the United States would) create proprietary rights over compilations of facts, often without even the carefully framed exceptions of the copyright scheme, such as the usefully protean category of “fair use.”

The older strategy of intellectual property law was a “braided” one: thread a thin layer of intellectual property rights around a commons of material from which future creators would draw. Even that thin layer of intellectual property rights was limited so as to allow access to the material when the private-property owner might charge too much, or just refuse; fair use allows for parody, commentary, and criticism, and also for “decompilation” of computer programs so that Microsoft Word’s competitors can reverse-engineer its features in order to make sure that their program can convert Word files. (Those who prefer topographical metaphors might imagine a quilted pattern of public and private land, with legal rules specifying that certain areas—beaches, say—can never be privately owned, and accompanying rules giving public right of way through private land if there is a danger that access to the commons might otherwise be blocked.)

From the inception of intellectual property law in the eighteenth century until quite recently, protection of the public domain—the intangible commons—was one fundamental goal of the law in most nations. In the new vision of intellectual property, however, property rights should be established everywhere; more is better. Expanding patentable and copyrightable subject matter, lengthening the copyright term, giving legal protection to “digital barbed wire,” even if it is used in part to prevent fair use: each of these can be understood as a vote of no confidence in the productive powers of the commons. We seem to be shifting from Brandeis’s assumption that the “noblest of human productions are free as the air to common use” to the assumption that any human production left open to free use is inefficient, if not tragic.

So far I have argued that there are profound similarities between the first enclosure movement and our contemporary expansion of intellectual property. Today, as in the fifteenth century, proponents and opponents of enclosure are locked in battle, hurling at each other incommensurable claims about innovation, efficiency, traditional values, the boundaries of the market, the saving of lives, the loss of familiar liberties. Once again, opposition to enclosure is portrayed as economically illiterate; the beneficiaries of enclosure tell us that an expansion of property rights is needed in order to fuel progress. Indeed, the post-Cold War “Washington Consensus” is invoked to claim that the lesson of history itself is that the only way one gets growth and efficiency is through markets; property rights, surely, are the sine qua non of markets.

But if there are similarities between the two enclosure movements, there are also crucial differences. The digitized and networked “commons of the mind,” circa 2002, differs greatly from the grassy and isolated common plots of land that dotted England circa 1400. (2) Some of the key differences should lead us to question whether stronger intellectual property rights are really either necessary or desirable.

For example, consider the well-known fact that a digital text, unlike a plot of land, can be used by countless people simultaneously without mutual interference or destruction of the shared resource. Unlike an earthly commons, the commons of the mind is generally what economists call “nonrival.” Many uses of land are mutually exclusive. If I am using the field for grazing, it may interfere with your plans to use it for growing crops. By contrast, a gene sequence, an MP3 file, or an image may be used by multiple parties; my use does not interfere with yours. To simplify a complicated analysis, this means that the depredations through overuse that affect fields and fisheries are generally not a problem with intellectual property. (The exceptions to this statement turn out to be fascinating; in the interest of brevity I will ignore them entirely.)

Thus, one cause of tragedy on the earthly commons generally does not arise on the commons of the mind. Overuse is normally not a problem. But what about incentives to create the intellectual resources in the first place?

Here intellectual property, especially in our digitized age, seems at first glance to pose a unique problem. It has long been relatively easy for pirates to produce unauthorized copies of poems, novels, treatises, and musical compositions. In the language of the economists, it has long been difficult, and in some cases virtually impossible, to stop one unit of an intellectual good from satisfying an infinite number of users at zero marginal cost. A familiar conclusion seems irresistible: without an ability to protect their creations against theft, creators will be unable to earn an adequate living. There will be inadequate incentives to create. Thus the law must step in and create a monopoly called an intellectual property right.

This is the standard argument in favor of intellectual property rights, but it has recently acquired a historical dimension, a teleology of expansion over time. After all, in our digitized age, it is easier than ever before for pirates to copy not just a book, but a film, a photograph, a recorded piece of music, a drug formula, a computer program - the list goes on. Surely the historical lowering of copying and transmission costs implies a corresponding need to increase the strength of intellectual property rights.

Imagine a line. At one end sits a monk, painstakingly transcribing Aristotle’s *Poetics*. In the middle lies the Gutenberg printing press. Three-quarters of the way along the line is a photocopying machine. At the end lies the Internet. At each stage, copying costs are lowered: Aristotle’s text becomes ever more freely and widely accessible; indeed, the complete text is currently available in both Greek and English to anyone with access to the Internet.

Among some analysts, the assumption seems to be that the strength of intellectual property rights must correspond inversely to the cost of copying. The argument goes something like this: To deal with the monk-copyist, we need no intellectual property right; physical control of the manuscript is enough. To deal with the Gutenberg press, we need the Statute of Anne. But to deal with the Internet, we need the Digital Millennium Copyright Act, the No Electronic Theft Act, the Sonny Bono Copyright Term Extension Act, and perhaps even the Collections of Information Anti-Piracy Act. Why? As copying costs approach zero, intellectual property rights must approach perfect control. And if a greater proportion of product value and

GNP is now in the form of information, then obviously we have an independent reason to need strengthened protection. A five-dollar padlock would do for a garden shed, but not for a vault.

Like any attractive but misleading argument, this one has some truth. The Internet does lower the cost of copying and facilitates illicit copying. The same technology also lowers the costs of production, distribution, and advertising—and dramatically increases the size of the potential market.

Is the “net” result, then, a loss to rights-holders such that we need to increase protection in order to maintain a constant level of incentives? The answer is not self-evident.

A large, leaky market may actually produce more revenue than a small, tightly controlled market. What’s more, the same technologies that allow for cheap copying also allow for swift and encyclopedic search engines—the best detection device for illicit copying ever invented. It would be impossible to say, on the basis of the evidence we have, that owners of protected content are better or worse off as a result of the Internet.

My intuition—as well as our historical experience with prior “dangerous” technologies such as the VCR—points strongly to the possibility that copyright holders are better off. In any case, there simply isn’t enough evidence, either to support my intuition or to support the conclusion that as copy costs decline intellectual property rights must be strengthened. Furthermore, given the known static and dynamic costs of monopolies, and the constitutional injunction to encourage the progress of science and useful arts, the burden should be on those requesting expanded intellectual property rights to prove their value.

Another argument commonly offered in defense of granting new intellectual property rights stresses the increasing importance of products that use, embody, or process information in today’s global economy. Perhaps the commons of the mind requires enclosure because it is now such a vital sector of economic activity. The importance of agriculture to the economy was certainly one of the arguments for the first enclosure movement. (Lovers of Patrick O’Brian’s novels may remember Maturin’s stolid silence in the face of an admiral’s increasingly vehement insistence that enclosure was essential to produce the corn necessary to fight the Napoleonic war.)

Here we come to another big difference between the commons of the mind and the earthly commons. As has frequently been pointed out (by Jessica Litman, Pamela Samuelson, and Richard A. Posner, among others), information products are frequently made out of fragments of other information products; one person’s information output is someone else’s information input. These inputs may be snippets of code, discoveries, prior research, images, genres of work, cultural references, databases of single nucleotide polymorphisms—all can function as raw material for future innovation. And every potential increase of protection over such products also raises the costs of, or reduces access to, the raw material to create new products.

The right balance is difficult to strike. One Nobel Prize-winning economist has claimed that it is actually impossible to produce an “informationally efficient” market. (3) Whether or not it is impossible in theory, it is surely a difficult problem in practice. In other words, even if enclosure of the arable commons always produced gains (itself a subject of debate), enclosure of the information commons clearly has some potential to harm intellectual innovation. More property rights, even though they supposedly offer greater incentives, do not necessarily ensure greater intellectual productivity. Sometimes just the opposite may be true. (4)

My arguments so far have taken as a given the various problems to which modern intellectual property laws have been a response. I have discussed the extent to which the logic of enclosure works for the commons of the mind as well as it did for the arable commons, taking into account the effects of an information society and a global Internet. Remember that when I speak of enclosure, I am talking about increases in the level of rights: protecting new subject matter for longer periods of time, criminalizing certain technologies, making

it illegal to cut through digital fences even if they have the effect of foreclosing previously lawful uses, and so on.

What I have not yet done is ask whether the brute fact of the Internet actually unsettles old assumptions and forces us to reconsider the need for incentives—at least in certain areas. But this is a question that cannot be evaded.

For anyone interested in the way that computer networks may embody a new mode of collaborative production, an exemplary case to study is the open-source software movement. (5) This software is released under a series of licenses, the most important being the General Public License, or GPL. The GPL specifies that anyone may copy the software, provided the license remains attached and the “source code” for the software always remains available. (6) Users may add to or modify the code, may build on it and incorporate it into their own work, but if they do so then the new program created is also covered by the GPL. Some people refer to this as the “viral” nature of the license; others find the term offensive. The point, however, is that the open quality of the creative enterprise spreads; it is not simply a donation of a program or a work to the public domain, but a continual accretion in which all gain the benefits of the program on pain of agreeing to give their own additions and innovations back to the communal project.

The open-source software movement has produced software that either rivals or exceeds the productive capacities of conventional proprietary software. Its adoption on the enterprise level is impressive, as are the various technical encomia to its strengths.

But the most remarkable aspect of the open-source software movement is harder to see. It functions as a new kind of social system: many of those who contribute to the movement by writing a part of the software do so as volunteers, without direct remuneration. Here, it seems, we have a classic public good—code that can be copied freely, and sold or redistributed without paying the creator or creators.

Skeptics, of course, wonder if this mode of production can be sustained. There seem to be inadequate incentives to ensure continued productivity and innovation. *E pur si muove*, as Galileo is reputed to have said in the face of Cardinal Bellarmine’s certainties—“And yet it moves.”

Still, there is no consensus about why the system works. Perhaps the open-source software movement is actually a contemporary form of potlatch, in which one gains prestige by the extravagance of the resources one “wastes.” Perhaps it is simply a smart way for a young programmer to build a resume that will eventually pay off in a conventional job. Or perhaps the movement is driven by what Karl Marx considered an innate aspect of our “species-being”: namely, the urge to create, which drives human beings to labor out of love rather than material need.

Like Yochai Benkler and Eben Moglen, I believe that such speculation is interesting but irrelevant. (7) My own explanation for why the system works is this:

Assume a random distribution of incentive structures in different people, a global network. Assume also that the costs of transmission, information sharing, and copying approach zero. Assume finally a modular creation process. With these assumptions, it just doesn’t matter why unpaid code writers do what they do; what matters is that a certain number of people will do what the unpaid code writers do. One may do it for love of the species, another in the hope of a better job, a third for the joy of solving puzzles, and so on. Each person also has his or her own “reserve price,” the point at which he or she says “now I will turn off *Survivor* and go and create something.” But on a global network, there are a lot of people, and with numbers that big, and information-overhead that small, even relatively hard projects will attract a sufficient number of motivated and skilled people to sustain the creative process. For the whole structure to work without large-scale centralized coordination, the creation process has to be modular, with “units” of different size and complexity, each requiring slightly different expertise, all of which can be added together to make a grand whole. I can work on the sendmail program, you on the search algorithms. More likely, lots of people try to solve the sendmail and search algorithm problems, and their products are judged by the

community and the best ones adopted. Under these conditions—an ad hoc mode of production that curiously combines anarchism and entrepreneurialism, Kropotkin and Adam Smith—we will get innovation and productivity, without having to rely on the proprietary model.

What's more (and this is a truly fascinating twist), when the production process does need more centralized coordination, some governance that guides how the modular bits are most productively associated, it is at least theoretically possible that we can come up with the control system in exactly the same way; distributed production is potentially recursive. Governance processes, too, can be assembled through distributed methods on a global network, by people with widely varying motivations, skills, and reserve prices.

Again, skeptics will have their doubts. One organization theorist I know dismisses the possibility of anarchic coordination as “governance by food fight.” Anyone who has ever been on an organizational listserv, or been part of a global production process run by people who are long on brains and short on social skills, knows how accurate that description is. *E pur si muove*.

But, in the language of computer programmers, does the open-source software movement “scale”? Can we generalize anything from this limited example? How many types of production, innovation, and research fit into the model I have just described? After all, for lots of types of innovation and invention one needs hardware, capital investment, large-scale real-world data collection, stuff—in all its facticity and infinite recalcitrance. Maybe the open-source model has solved the individual incentives problem, but that's not the only problem. And how many types of innovation or cultural production are as modular as software?

My own guess is that this method of production is far more common than we realize. “Even before the Internet,” as some of my students have taken to saying portentously, science, law, education, and musical genres all developed in ways that are markedly similar to the model I have described. “The marketplace of ideas,” the continuous roiling development in thought and norm that our political culture spawns, is itself an idea that owes much more to the distributed, nonproprietary model than it does to the special case of commodified innovation that we regulate through intellectual property law. It's not that copyright and patent haven't helped power the rise of modern civilization; it's just that it would be wrong to see them as the only engine of innovation. Indeed, the mottoes of free software development have their counterparts in the theory of democracy and the open society. The open-source movement describes its advantage over closed and secretive systems concisely: “given enough eyeballs, all bugs are shallow.” Karl Popper would have cheered.

Furthermore, I suspect that the increasing migration of the sciences toward data-rich, processing-rich models will make it likely that a greater amount of innovation and discovery could follow the distributed, nonproprietary model of intellectual production. Bio-informatics and computational biology, the open-source genomics project at www.ensembl.org, the possibility of distributed data scrutiny by lay volunteers that NASA used on the Mars landing data—all of these offer intriguing glances of a possible future. And finally, of course, the Internet is one big experiment in distributed cultural production.

My own utopia would include modes of nonproprietary intellectual production flourishing alongside a scaled-down but still powerful intellectual property regime. Of course, my utopia hinges on a hunch about the future. Still, there is some possibility (I might say hope) that we could have a world in which much more intellectual production is free—“free” meaning that it is not subject to centralized control, and “free” meaning that its products are available without payment. Insofar as this is at least a possible future, then surely we should think twice before foreclosing it.

Yet foreclosing this possibility is precisely what lawmakers and government regulators in America are now doing. The point about the dramatic recent expansion of intellectual property rights—in database protection bills and directives that extend intellectual property rights to the layer of facts, in the efflorescence of software patents, in the validation of shrink-wrap licenses that bind third parties, in the Digital Millennium Copyright Act's anticircumvention provisions—is not merely that they hamper the nonproprietary mode of

intellectual production unfairly and without justification. The point is rather that they run the risk of ruling it out altogether. (8)

We have come full circle. As I have shown, we are in many ways in the midst of a second enclosure movement. The opponents and proponents of enclosure are currently locked in battle, each appealing to conflicting and sometimes incommensurable claims about efficiency, innovation, justice, and the limits of the market.

But should there be a second enclosure movement? Do we know that property rights in this sphere will yield the same surge of productive energy that they did when applied to arable land?

I think the answer is a resounding No. We are rushing to fence in ever-larger stretches of the commons of the mind without convincing economic evidence that enclosure will help either productivity or innovation—and with very good reason to believe it may actually hurt them. (9)

As I have argued elsewhere, this process should bother people across the ideological spectrum, from civil libertarians to free marketeers. Researchers and scientists should be particularly worried by what is happening. Up to now, the American system of science, for all its flaws, has worked astoundingly well; changing some of its fundamental premises, such as by moving property rights into the data layer, is not something to be done lightly.

The dangers are particularly acute at the moment for three reasons. First, under the conditions that currently obtain in our digitized commons of the mind, the creation of new intellectual property rights tends, in a vicious circle, to create still further demands for new intellectual property rights. The argument is a little too complicated to lay out here. (10) But in essence the position is this: once a new intellectual property right has been created over some informational good, the only way to ensure efficient allocation of that good is to give the rights holder the ability to charge every user the exact maximum each consumer is willing to pay, so that the market can be perfectly segregated by price. In order to protect their ability to set prices for digital intellectual property goods, whose marginal cost to produce and distribute in fact approaches zero, the rights holders will inevitably argue that they need even more changes of the rules in their favor: relaxed privacy standards, so they can know more about consumers' price points; enforceable shrink-wrap or click-wrap contracts, so that consumers can be held to the term of a particular license, no matter how restrictive; and changes in antitrust rules, to allow for a variety of practices that are currently illegal, such as resale price maintenance and various forms of "tying." Rights holders will also claim that they need technical changes with legal backing: for example, the creation of personalized digital objects surrounded by state-sanctioned digital fences, objects that are tied to particular users and particular computers, so that reading my e-book on your machine is either technically impossible, a crime, or a tort—or possibly all three. My conclusion: extending ever-stronger intellectual property rights is a very slippery slope.

Second, the broader the scope of intellectual property rights, the more the characteristics of the Internet that have made it so attractive to civil libertarians—its distributed, anonymous character, its resistance to control or filtering by public or private entities, its global nature—start to seem like vices rather than virtues. The process of trying to make the Net safe for price discrimination has already begun. Yet as Lawrence Lessig has argued, this is a fundamental political choice that ought to be made deliberately and publicly, not as a side effect of an economically dubious digital enclosure movement. Because of some threats, such as terrorism, we might choose to live in a pervasively monitored electronic environment in which identity and geography, and thus regulability, have been reintroduced. (In my own view, the price is not worth paying.) But to do so on the basis of some bad microeconomic arguments about the needs of the entertainment industry and in the absence of good empirical evidence and to foreclose some of the most interesting new productive possibilities in the process -- well, that would be really sad.

Third, the arguments in favor of the new enclosure movement depend heavily on the intellectually complacent, analytically unsound assumptions of "neoliberal orthodoxy," the "Washington consensus."

Convinced that property is good, and that creating more property rights is better, neoliberals are primed to hand out patents on gene sequences and stem cell lines and copyrights on compilations of facts. It would be ironic, to say the least, to let such neoliberal convictions determine the fate of the information commons, the one area where the pros and cons of a property regime need to be most delicately balanced, and also an area where the possible consequences for the public good ought to be vigorously and openly debated.

What is to be done, then? I cannot lay out a full answer here, but I would suggest two broad strategies. First, we ought to insist on considerably better empirical and economic evidence before signing on to the proposals of the second enclosure movement. There are a few serious comparative and historical studies of the economics of innovation, but we need a lot more. Indeed, there should be an annual audit of our intellectual property system, perhaps by the General Accounting Office. What are the costs—static and dynamic—and the benefits of our current intellectual property regime? After all, this is one of the largest industry subsidies given by government (through its granting of patents and copyrights); it deserves the same searching scrutiny that we apply to the recipients of other state subsidies. I am a firm believer in intellectual property rights; properly balanced and judiciously applied, such rights promise us a wonderfully decentralized system for the promotion of innovation. But this is a rational belief in particular rules based on empirical evidence, not an unquestioning faith that any increase in intellectual property rights is automatically good.

Second, we need to make clear the current dangers to the public domain, in the same way that environmental activists in the 1950s and 1960s made visible not only particular environmental threats but the very existence of “the environment” itself. The environmental movement gained much of its political power by pointing out that there were structural reasons why lawmakers were likely to make bad environmental decisions: a legal system based on a particular notion of what “private property” entailed, and a technological tendency to treat the world as a simple, linear set of causes and effects, ignoring the complex interrelationship among natural systems. In both of these conceptual systems, the environment actually disappeared; there was no place for it in the analysis. Small surprise, then, that lawmakers were not able to protect it properly.

We should press a similar argument—as I have done here—in the case of the public domain. (11) We should exploit the power of a concept like the public domain both to clarify and to reshape perceptions of self-interest. The idea that there is a public domain—a commons of the mind—can help a coalition to be built around a refrained conception of common interest. In the narrowest sense, that common interest might be the realization, spurred by greater attention to intellectual interrelationships, that the freest possible circulation of ideas and facts is important to anyone whose well-being significantly depends on intellectual innovation and productivity—that is to say, every citizen of the world.

The poem with which I began this essay contained some advice: And geese will still a common lack / Till they go and steal it back.

I can’t match the terseness or the rhyme. But if we blithely assume that the second enclosure movement will have the same benign effects as the first, we may look like very silly geese indeed.

(1.) The analogy to the enclosure movement has been too succulent to resist. To my knowledge, Ben Kaplan, Pamela Samuleson, Yochai Benkler, David Lange, Christopher May, and Keith Aoki have all employed the trope, as I have myself on previous occasions. For a particularly thoughtful and careful development of the parallelism see Hannibal Travis, “Pirates of the Information Infrastructure: Blackstonian Copyright and the First Amendment,” *Berkeley Tech. Law Journal* 15 (2) (Spring 2000): 777.

(2.) The differences are particularly strong in the arguments over “desert”—are these property rights deserved, or are they simply violations of the public trust, privatizations of the commons? For example, some would say that we never had the same traditional claims over the genetic commons that the victims of the first enclosure movement had over theirs; this is more like newly discovered frontier land, or perhaps even privately drained marshland, than it is like well-

known common land that all have traditionally used. In this case, the enclosers can claim (though their claims are disputed) that they discovered or perhaps simply made usable the territory they seek to own. The opponents of gene patenting, on the other hand, turn more frequently than the farmers of the eighteenth century to religious and ethical arguments about the sanctity of life and the incompatibility of property with living systems. These arguments, or the appeals to free speech that dominate debates over digital intellectual property, have no precise analogue in debates over hunting or pasturage, although, again, there are common themes. For example, we are already seeing nostalgic laments of the loss of the immemorial rights of Internet users. At the same time, the old language of property law is turned to this more evanescent subject matter; a favorite article title is “The Ancient Doctrine of Trespass to Websites” (I. Trotter Hardy, “The Ancient Doctrine of Trespass to Web Sites,” *Journal of Online Law* [Oct. 1996]: art. 7).

(3.) Sanford J. Grossman and Joseph E. Stiglitz, “On the Impossibility of Informationally Efficient Markets,” *American Economic Review* 70 (1980): 393.

(4.) For a more technical account, see James Boyle, “Cruel, Mean, or Lavish? Economic Analysis, Price Discrimination and Digital Intellectual Property,” *Vanderbilt Law Review* 536 (2000): 2007, .

(5.) Glyn Moody, *The Rebel Code: The Inside Story of Linux and the Open Source Revolution* (Cambridge, Mass.: Perseus Pub., 2001); Peter Wayner, *Free for All: How Linux and the Free Software Movement Undercut the High-tech Titans* (New York: HarperBusiness, 2000). See also Eben Moglen, “Anarchism Triumphant: Free Software and the Death of Copyright,” in the online journal *First Monday* (1999).

(6.) Proprietary, or “binary only,” software is generally released only once the source code has been compiled into machine-readable object code format, a form that is impenetrable to the user. Even if you were a master programmer, and if the provisions of the Copyright Act, the appropriate licenses, and the DMCA did not forbid you from doing so, you would be unable to modify commercial proprietary software so as to customize it for your needs, remove a bug, or add a feature. Open-source programmers say disdainfully that it is like buying a car with “the hood welded shut.”

(7.) See Yochai Benkler, “Coase’s Penguin, or, Linux and the Nature of the Firm,” October 2001, unpublished draft. For a seminal statement relying on the innate human love of creativity as the motivation, see Moglen, “Anarchism Triumphant.” “[I]ncentives” is merely a metaphor, and as a metaphor to describe human creative activity it’s pretty crummy. I have said this before, but the better metaphor arose on the day Michael Faraday first noticed what happened when he wrapped a coil of wire around a magnet and spun the magnet. Current flows in such a wire, but we don’t ask what the incentive is for the electrons to leave home. We say that the current results from an emergent property of the system, which we call induction. The question we ask is ‘what’s the resistance of the wire?’ So Moglen’s Metaphorical Corollary to Faraday’s Law says that if you wrap the Internet around every person on the planet and spin the planet, software flows in the network. It’s an emergent property of connected human minds that they create things for one another’s pleasure and to conquer their uneasy sense of being too alone. The only question to ask is, what’s the resistance of the network? Moglen’s Metaphorical Corollary to Ohm’s Law states that the resistance of the network is directly proportional to the field strength of the ‘intellectual property’ system. So the right answer to the econodwarf is, resist the resistance.”

(8.) This point has been ably made by, inter alia, Pamela Samuelson, Jessica Litman, Jerry Reichman, Larry Lessig, and Yochai Benkler. Each has a slightly different focus and emphasis on the problem, but each has pointed out the impediments now being erected to distributed, nonproprietary solutions. See also Boyle, “Cruel, Mean, or Lavish?”

(9.) Some of the legislation involved is also constitutionally dubious, under the First Amendment and Art 1 sec. 8 cl. 8 of the Constitution, but that is a point for another paper.

(10.) The full version is given in Boyle, “Cruel, Mean, or Lavish?”

(11.) An expanded version of this argument can be found in “A Politics of Intellectual Property: Environmentalism for the Net,” *Duke Law Journal* 47 (1) (1997): 87.

This is a verbatim copy of www.gnu.org/gnu/manifesto.html, converted into pdf for inclusion in a course reader.

The GNU Manifesto

by

Richard Stallman

The GNU Manifesto (which appears below) was written by Richard Stallman at the beginning of the GNU Project, to ask for participation and support. For the first few years, it was updated in minor ways to account for developments, but now it seems best to leave it unchanged as most people have seen it.

Since that time, we have learned about certain common misunderstandings that different wording could help avoid. Footnotes added in 1993 help clarify these points.

For up-to-date information about the available GNU software, please see the information available on our web server, in particular our list of software. For how to contribute, see <http://www.gnu.org/help>.

What's GNU? Gnu's Not Unix!

GNU, which stands for Gnu's Not Unix, is the name for the complete Unix-compatible software system which I am writing so that I can give it away free to everyone who can use it. (1) Several other volunteers are helping me. Contributions of time, money, programs and equipment are greatly needed.

So far we have an Emacs text editor with Lisp for writing editor commands, a source level debugger, a yacc-compatible parser generator, a linker, and around 35 utilities. A shell (command interpreter) is nearly completed. A new portable optimizing C compiler has compiled itself and may be released this year. An initial kernel exists but many more features are needed to emulate Unix. When the kernel and compiler are finished, it will be possible to distribute a GNU system suitable for program development. We will use TeX as our text formatter, but an nroff is being worked on. We will use the free, portable X window system as well. After this we will add a portable Common Lisp, an Empire game, a spreadsheet, and hundreds of other things, plus on-line documentation. We hope to supply, eventually, everything useful that normally comes with a Unix system, and more.

GNU will be able to run Unix programs, but will not be identical to Unix. We will make all improvements that are convenient, based on our experience with other operating systems. In particular, we plan to have longer file names, file version numbers, a crashproof file system, file name completion perhaps, terminal-independent display support, and perhaps eventually a Lisp-based window system through which several Lisp programs and ordinary Unix programs can share a screen. Both C and Lisp will be available as system programming languages. We will try to support UUCP, MIT Chaosnet, and Internet protocols for communication.

GNU is aimed initially at machines in the 68000/16000 class with virtual memory, because they are the easiest machines to make it run on. The extra effort to make it run on smaller machines will be left to someone who wants to use it on them.

To avoid horrible confusion, please pronounce the 'G' in the word 'GNU' when it is the name of this project.

Why I Must Write GNU

I consider that the golden rule requires that if I like a program I must share it with other people who like it. Software sellers want to divide the users and conquer them, making each user agree not to share with others. I refuse to break solidarity with other users in this way. I cannot in good conscience sign a nondisclosure agreement or a software license agreement. For years I worked within the Artificial Intelligence Lab to resist such tendencies and other inhospitalities, but eventually they had gone too far: I could not remain in an institution where such things are done for me against my will.

So that I can continue to use computers without dishonor, I have decided to put together a sufficient body of free software so that I will be able to get along without any software that is not free. I have resigned from the AI lab to deny MIT any legal excuse to prevent me from giving GNU away.

Why GNU Will Be Compatible with Unix

Unix is not my ideal system, but it is not too bad. The essential features of Unix seem to be good ones, and I think I can fill in what Unix lacks without spoiling them. And a system compatible with Unix would be convenient for many other people to adopt.

How GNU Will Be Available

GNU is not in the public domain. Everyone will be permitted to modify and redistribute GNU, but no distributor will be allowed to restrict its further redistribution. That is to say, proprietary modifications will not be allowed. I want to make sure that all versions of GNU remain free.

Why Many Other Programmers Want to Help

I have found many other programmers who are excited about GNU and want to help.

Many programmers are unhappy about the commercialization of system software. It may enable them to make more money, but it requires them to feel in conflict with other programmers in general rather than feel as comrades. The fundamental act of friendship among programmers is the sharing of programs; marketing arrangements now typically used essentially forbid programmers to treat others as friends. The purchaser of software must choose between friendship and obeying the law. Naturally, many decide that friendship is more important. But those who believe in law often do not feel at ease with either choice. They become cynical and think that programming is just a way of making money.

By working on and using GNU rather than proprietary programs, we can be hospitable to everyone and obey the law. In addition, GNU serves as an example to inspire and a banner to rally others to join us in sharing. This can give us a feeling of harmony which is impossible if we use software that is not free. For about half the programmers I talk to, this is an important happiness that money cannot replace.

How You Can Contribute

(Nowadays, for software tasks to work on, see [the GNU task list](#). For other ways to contribute, see <http://www.gnu.org/help>.)

I am asking computer manufacturers for donations of machines and money. I'm asking individuals for donations of programs and work.

One consequence you can expect if you donate machines is that GNU will run on them at an early date. The machines should be complete, ready to use systems, approved for use in a residential area, and not in need of sophisticated cooling or power.

I have found very many programmers eager to contribute part-time work for GNU. For most projects, such part-time distributed work would be very hard to coordinate; the independently-written parts would not work together. But for the particular task of replacing Unix, this problem is absent. A complete Unix system contains hundreds of utility programs, each of which is documented separately. Most interface specifications are fixed by Unix compatibility. If each contributor can write a compatible replacement for a single Unix utility, and make it work properly in place of the original on a Unix system, then these utilities will work right when put together. Even allowing for Murphy to create a few unexpected problems, assembling these components will be a feasible task. (The kernel will require closer communication and will be worked on by a small, tight group.)

If I get donations of money, I may be able to hire a few people full or part time. The salary won't be high by programmers' standards, but I'm looking for people for whom building community spirit is as important as making money. I view this as a way of enabling dedicated people to devote their full energies to working on GNU by sparing them the need to make a living in another way.

Why All Computer Users Will Benefit

Once GNU is written, everyone will be able to obtain good system software free, just like air.[\(2\)](#)

This means much more than just saving everyone the price of a Unix license. It means that much wasteful duplication of system programming effort will be avoided. This effort can go instead into advancing the state of the art.

Complete system sources will be available to everyone. As a result, a user who needs changes in the system will always be free to make them himself, or hire any available programmer or company to make them for him. Users will no longer be at the mercy of one programmer or company which owns the sources and is in sole position to make changes.

Schools will be able to provide a much more educational environment by encouraging all students to study and improve the system code. Harvard's computer lab used to have the policy that no program could be installed on the system if its sources were not on public display, and upheld it by actually refusing to install certain programs. I was very much inspired by this.

Finally, the overhead of considering who owns the system software and what one is or is not entitled to do with it will be lifted.

Arrangements to make people pay for using a program, including licensing of copies, always incur a tremendous cost to society through the cumbersome mechanisms necessary to figure out how much (that is, which programs) a person must pay for. And only a police state can force everyone to obey them. Consider a space station where air must be manufactured at great cost: charging each breather per liter of air may be fair, but wearing the metered gas mask all day and all night is intolerable even if everyone can afford to pay the air bill. And the TV cameras everywhere to see if you ever take the mask off are outrageous. It's better to support the air plant with a head tax and chuck the masks.

Copying all or parts of a program is as natural to a programmer as breathing, and as productive. It ought to be as free.

Some Easily Rebutted Objections to GNU's Goals

“Nobody will use it if it is free, because that means they can't rely on any support.”

“You have to charge for the program to pay for providing the support.”

If people would rather pay for GNU plus service than get GNU free without service, a company to provide just service to people who have obtained GNU free ought to be profitable.[\(3\)](#)

We must distinguish between support in the form of real programming work and mere handholding. The former is something one cannot rely on from a software vendor. If your problem is not shared by enough people, the vendor will tell you to get lost.

If your business needs to be able to rely on support, the only way is to have all the necessary sources and tools. Then you can hire any available person to fix your problem; you are not at the mercy of any individual. With Unix, the price of sources puts this out of consideration for most businesses. With GNU this will be easy. It is still possible for there to be no available competent person, but this problem cannot be blamed on distribution arrangements. GNU does not eliminate all the world's problems, only some of them.

Meanwhile, the users who know nothing about computers need handholding: doing things for them which they could easily do themselves but don't know how.

Such services could be provided by companies that sell just hand-holding and repair service. If it is true that users would rather spend money and get a product with service, they will also be willing to buy the service having got the product free. The service companies will compete in quality and price; users will not be tied to any particular one. Meanwhile, those of us who don't need the service should be able to use the program without paying for the service.

“You cannot reach many people without advertising, and you must charge for the program to support that.”

“It's no use advertising a program people can get free.”

There are various forms of free or very cheap publicity that can be used to inform numbers of computer users about something like GNU. But it may be true that one can reach more microcomputer users with advertising. If this is really so, a business which advertises the service of copying and mailing GNU for a fee ought to be successful enough to pay for its advertising and more. This way, only the users who benefit from the advertising pay for it.

On the other hand, if many people get GNU from their friends, and such companies don't succeed, this will show that advertising was not really necessary to spread GNU. Why is it that free market advocates don't want to let the free market decide this?[\(4\)](#)

“My company needs a proprietary operating system to get a competitive edge.”

GNU will remove operating system software from the realm of competition. You will not be able to get an edge in this area, but neither will your competitors be able to get an edge over you. You and they will compete in other areas, while benefiting mutually in this one. If your business is selling an operating system, you will not like GNU, but that's tough on you. If your business is something else, GNU can save you from being pushed into the expensive business of selling operating systems.

I would like to see GNU development supported by gifts from many manufacturers and users, reducing the cost to each.[\(5\)](#)

“Don't programmers deserve a reward for their creativity?”

If anything deserves a reward, it is social contribution. Creativity can be a social contribution, but only in so far as society is free to use the results. If programmers deserve to be rewarded for creating innovative programs, by the same token they deserve to be punished if they restrict the use of these programs.

“Shouldn't a programmer be able to ask for a reward for his creativity?”

There is nothing wrong with wanting pay for work, or seeking to maximize one's income, as long as one does not use means that are destructive. But the means customary in the field of software today are based on destruction.

Extracting money from users of a program by restricting their use of it is destructive because the restrictions reduce the amount and the ways that the program can be used. This reduces the amount of wealth that humanity derives from the program. When there is a deliberate choice to restrict, the harmful consequences are deliberate destruction.

The reason a good citizen does not use such destructive means to become wealthier is that, if everyone did so, we would all become poorer from the mutual destructiveness. This is Kantian ethics; or, the Golden Rule. Since I do not like the consequences that result if everyone hoards information, I am required to consider it wrong for one to do so. Specifically, the desire to be rewarded for one's creativity does not justify depriving the world in general of all or part of that creativity.

“Won't programmers starve?”

I could answer that nobody is forced to be a programmer. Most of us cannot manage to get any money for standing on the street and making faces. But we are not, as a result, condemned to spend our lives standing on the street making faces, and starving. We do something else.

But that is the wrong answer because it accepts the questioner's implicit assumption: that without ownership of software, programmers cannot possibly be paid a cent. Supposedly it is all or nothing.

The real reason programmers will not starve is that it will still be possible for them to get paid for programming; just not paid as much as now.

Restricting copying is not the only basis for business in software. It is the most common basis because it brings in the most money. If it were prohibited, or rejected by the customer, software business would move to other bases of organization which are now used less often. There are always numerous ways to organize any kind of business.

Probably programming will not be as lucrative on the new basis as it is now. But that is not an argument against the change. It is not considered an injustice that sales clerks make the salaries that they now do. If programmers made the same, that would not be an injustice either. (In practice they would still make considerably more than that.)

“Don't people have a right to control how their creativity is used?”

“Control over the use of one's ideas” really constitutes control over other people's lives; and it is usually used to make their lives more difficult.

People who have studied the issue of intellectual property rights(6) carefully (such as lawyers) say that there is no intrinsic right to intellectual property. The kinds of supposed intellectual property rights that the government recognizes were created by specific acts of legislation for specific purposes.

For example, the patent system was established to encourage inventors to disclose the details of their inventions. Its purpose was to help society rather than to help inventors. At the time, the life span of 17 years for a patent was short compared with the rate of advance of the state of the art. Since patents are an issue only among manufacturers, for whom the cost and effort of a license agreement are small compared with setting up production, the patents often do not do much harm. They do not obstruct most individuals who use patented products.

The idea of copyright did not exist in ancient times, when authors frequently copied other authors at length in works of non-fiction. This practice was useful, and is the only way many authors' works have survived even in part. The copyright system was created expressly for the purpose of encouraging authorship. In the domain for which it was invented—books, which could be copied economically only on a printing press—it did little harm, and did not obstruct most of the individuals who read the books.

All intellectual property rights are just licenses granted by society because it was thought, rightly or wrongly, that society as a whole would benefit by granting them. But in any particular situation, we have to ask: are we really better off granting such license? What kind of act are we licensing a person to do?

The case of programs today is very different from that of books a hundred years ago. The fact that the easiest way to copy a program is from one neighbor to another, the fact that a program has both source code and object code which are distinct, and the fact that a program is used rather than read and enjoyed, combine to create a situation in which a person who enforces a copyright is harming society as a whole both materially and spiritually; in which a person should not do so regardless of whether the law enables him to.

“Competition makes things get done better.”

The paradigm of competition is a race: by rewarding the winner, we encourage everyone to run faster. When capitalism really works this way, it does a good job; but its defenders are wrong in assuming it always works this way. If the runners forget why the reward is offered and become intent on winning, no matter how, they may find other strategies—such as, attacking other runners. If the runners get into a fist fight, they will all finish late.

Proprietary and secret software is the moral equivalent of runners in a fist fight. Sad to say, the only referee we've got does not seem to object to fights; he just regulates them (“For every ten yards you run, you can fire one shot”). He really ought to break them up, and penalize runners for even trying to fight.

“Won't everyone stop programming without a monetary incentive?”

Actually, many people will program with absolutely no monetary incentive. Programming has an irresistible fascination for some people, usually the people who are best at it. There is no shortage of professional musicians who keep at it even though they have no hope of making a living that way.

But really this question, though commonly asked, is not appropriate to the situation. Pay for programmers will not disappear, only become less. So the right question is, will anyone program with a reduced monetary incentive? My experience shows that they will.

For more than ten years, many of the world's best programmers worked at the Artificial Intelligence Lab for far less money than they could have had anywhere else. They got many kinds of non-monetary rewards: fame and appreciation, for example. And creativity is also fun, a reward in itself.

Then most of them left when offered a chance to do the same interesting work for a lot of money.

What the facts show is that people will program for reasons other than riches; but if given a chance to make a lot of money as well, they will come to expect and demand it. Low-paying organizations do poorly in competition with high-paying ones, but they do not have to do badly if the high-paying ones are banned.

“We need the programmers desperately. If they demand that we stop helping our neighbors, we have to obey.”

You’re never so desperate that you have to obey this sort of demand. Remember: millions for defense, but not a cent for tribute!

“Programmers need to make a living somehow.”

In the short run, this is true. However, there are plenty of ways that programmers could make a living without selling the right to use a program. This way is customary now because it brings programmers and businessmen the most money, not because it is the only way to make a living. It is easy to find other ways if you want to find them. Here are a number of examples.

A manufacturer introducing a new computer will pay for the porting of operating systems onto the new hardware.

The sale of teaching, hand-holding and maintenance services could also employ programmers.

People with new ideas could distribute programs as freeware⁽⁷⁾, asking for donations from satisfied users, or selling hand-holding services. I have met people who are already working this way successfully.

Users with related needs can form users’ groups, and pay dues. A group would contract with programming companies to write programs that the group’s members would like to use.

All sorts of development can be funded with a Software Tax:

Suppose everyone who buys a computer has to pay x percent of the price as a software tax. The government gives this to an agency like the NSF to spend on software development.

But if the computer buyer makes a donation to software development himself, he can take a credit against the tax. He can donate to the project of his own choosing—often, chosen because he hopes to use the results when it is done. He can take a credit for any amount of donation up to the total tax he had to pay.

The total tax rate could be decided by a vote of the payers of the tax, weighted according to the amount they will be taxed on.

The consequences:

- The computer-using community supports software development.
- This community decides what level of support is needed.
- Users who care which projects their share is spent on can choose this for themselves.

In the long run, making programs free is a step toward the post-scarcity world, where nobody will have to work very hard just to make a living. People will be free to devote themselves to activities that are fun, such as programming, after spending the necessary ten hours a week on required tasks such as legislation,

family counseling, robot repair and asteroid prospecting. There will be no need to be able to make a living from programming.

We have already greatly reduced the amount of work that the whole society must do for its actual productivity, but only a little of this has translated itself into leisure for workers because much nonproductive activity is required to accompany productive activity. The main causes of this are bureaucracy and isometric struggles against competition. Free software will greatly reduce these drains in the area of software production. We must do this, in order for technical gains in productivity to translate into less work for us.

Footnotes

1. The wording here was careless. The intention was that nobody would have to pay for **permission** to use the GNU system. But the words don't make this clear, and people often interpret them as saying that copies of GNU should always be distributed at little or no charge. That was never the intent; later on, the manifesto mentions the possibility of companies providing the service of distribution for a profit. Subsequently I have learned to distinguish carefully between "free" in the sense of freedom and "free" in the sense of price. Free software is software that users have the freedom to distribute and change. Some users may obtain copies at no charge, while others pay to obtain copies—and if the funds help support improving the software, so much the better. The important thing is that everyone who has a copy has the freedom to cooperate with others in using it.
2. This is another place I failed to distinguish carefully between the two different meanings of "free". The statement as it stands is not false—you can get copies of GNU software at no charge, from your friends or over the net. But it does suggest the wrong idea.
3. Several such companies now exist.
4. The Free Software Foundation for 10 years raised most of its funds from a distribution service, although it is a charity rather than a company. You can [order things from the FSF](#).
5. A group of computer companies pooled funds around 1991 to support maintenance of the GNU C Compiler.
6. In the 80s I had not yet realized how confusing it was to speak of "the issue" of "intellectual property". That term is obviously biased; more subtle is the fact that it lumps together various disparate laws which raise very different issues. Nowadays I urge people to reject the term "intellectual property" entirely, lest it lead others to suppose that those laws form one coherent issue. The way to be clear is to discuss patents, copyrights, and trademarks separately. See [further explanation](#) of how this term spreads confusion and bias.
7. Subsequently we have learned to distinguish between "free software" and "freeware". The term "freeware" means software you are free to redistribute, but usually you are not free to study and change the source code, so most of it is not free software. See [the Confusing Words and Phrases page](#) for more explanation.

Please send FSF & GNU inquiries to gnu@gnu.org. There are also [other ways to contact](#) the FSF. Please send broken links and other corrections or suggestions to webmasters@gnu.org.

Please see the [Translations README](#) for information on coordinating and submitting translations of this article.

Copyright © 1985, 1993, 2003, 2005, 2007 Free Software Foundation, Inc., 51 Franklin St, Fifth Floor, Boston, MA 02110-1301, USA. Permission is granted to anyone to make or distribute verbatim copies of this document, in any medium, provided that the copyright notice and permission notice are preserved, and that the distributor grants the recipient permission for further redistribution as permitted by this notice. Modified versions may not be made.

Updated: 2008/04/07 14:21:35

The Darknet and the Future of Content Distribution

Peter Biddle, Paul England, Marcus Peinado, and Bryan Willman

Microsoft Corporation¹

Abstract

We investigate the *darknet* – a collection of networks and technologies used to share digital content. The darknet is not a separate physical network but an application and protocol layer riding on existing networks. Examples of darknets are peer-to-peer file sharing, CD and DVD copying, and key or password sharing on email and newsgroups. The last few years have seen vast increases in the darknet’s aggregate bandwidth, reliability, usability, size of shared library, and availability of search engines. In this paper we categorize and analyze existing and future darknets, from both the technical and legal perspectives. We speculate that there will be short-term impediments to the effectiveness of the darknet as a distribution mechanism, but ultimately the darknet-genie will not be put back into the bottle. In view of this hypothesis, we examine the relevance of content protection and content distribution architectures.

1 Introduction

People have always copied things. In the past, most items of value were physical objects. Patent law and economies of scale meant that small scale copying of physical objects was usually uneconomic, and large-scale copying (if it infringed) was stoppable using policemen and courts. Today, things of value are increasingly less tangible: often they are just bits and bytes or can be accurately represented as bits and bytes. The widespread deployment of packet-switched networks and the huge advances in computers and codec-technologies has made it feasible (and indeed attractive) to deliver such digital works over the Internet. This presents great opportunities and great challenges. The opportunity is low-cost delivery of personalized, desirable high-quality content. The challenge is that such content can be distributed illegally. Copyright law governs the legality of copying and distribution of such valuable data, but copyright protection is increasingly strained in a world of programmable computers and high-speed networks.

For example, consider the staggering burst of creativity by authors of computer programs that are designed to share audio files. This was first popularized by Napster, but today several popular applications and services offer similar capabilities. CD-writers have become mainstream, and DVD-writers may well follow suit. Hence, even in the absence of network connectivity, the opportunity for low-cost, large-scale file sharing exists.

1.1 The Darknet

Throughout this paper, we will call the shared items (e.g. software programs, songs, movies, books, etc.) *objects*. The persons who copy objects will be called *users* of the darknet, and the computers used to share objects will be called *hosts*.

The idea of the darknet is based upon three assumptions:

¹ Statements in this paper represent the opinions of the authors and not necessarily the position of Microsoft Corporation.

1. Any widely distributed object will be available to a fraction of users in a form that permits copying.
2. Users will copy objects if it is possible and interesting to do so.
3. Users are connected by high-bandwidth channels.

The *darknet* is the distribution network that emerges from the injection of objects according to assumption 1 and the distribution of those objects according to assumptions 2 and 3.

One implication of the first assumption is that any content protection system will leak popular or interesting content into the darknet, because some fraction of users--possibly experts--will overcome any copy prevention mechanism or because the object will enter the darknet before copy protection occurs.

The term “widely distributed” is intended to capture the notion of mass market distribution of objects to thousands or millions of practically anonymous users. This is in contrast to the protection of military, industrial, or personal secrets, which are typically not widely distributed and are not the focus of this paper.

Like other networks, the darknet can be modeled as a directed graph with labeled edges. The graph has one vertex for each user/host. For any pair of vertices (u,v), there is a directed edge from u to v if objects can be copied from u to v. The edge labels can be used to model relevant information about the physical network and may include information such as bandwidth, delay, availability, etc. The vertices are characterized by their object library, object requests made to other vertices, and object requests satisfied.

To operate effectively, the darknet has a small number of technological and infrastructure requirements, which are similar to those of legal content distribution networks. These infrastructure requirements are:

1. facilities for injecting new objects into the darknet (input)
2. a distribution network that carries copies of objects to users (transmission)
3. ubiquitous rendering devices, which allow users to consume objects (output)
4. a search mechanism to enable users to find objects (database)
5. storage that allows the darknet to retain objects for extended periods of time. Functionally, this is mostly a caching mechanism that reduces the load and exposure of nodes that inject objects.

The dramatic rise in the efficiency of the darknet can be traced back to the general technological improvements in these infrastructure areas. At the same time, most attempts to fight the darknet can be viewed as efforts to deprive it of one or more of the infrastructure items. Legal action has traditionally targeted search engines and, to a lesser extent, the distribution network. As we will describe later in the paper, this has been partially successful. The drive for legislation on mandatory watermarking aims to deprive the darknet of rendering devices. We will argue that watermarking approaches are technically flawed and unlikely to have any material impact on the darknet. Finally, most content protection systems are meant to prevent or delay the injection of new objects into the darknet. Based on our first assumption, no such system constitutes an impenetrable barrier, and we will discuss the merits of some popular systems.

We see no technical impediments to the darknet becoming increasingly efficient (measured by aggregate library size and available bandwidth). However, the darknet, in all its transport-layer embodiments, is under legal attack. In this paper, we speculate on the technical and legal future of the darknet, concentrating particularly, but not exclusively, on peer-to-peer networks.

The rest of this paper is structured as follows. Section 2 analyzes different manifestations of the darknet with respect to their robustness to attacks on the infrastructure requirements described above and speculates on the future development of the darknet. Section 3 describes content protection mechanisms, their probable effect on the darknet, and the impact of the darknet upon

them. In sections 4 and 5, we speculate on the scenarios in which the darknet will be effective, and how businesses may need to behave to compete effectively with it.

2 The Evolution of the Darknet

We classify the different manifestations of the darknet that have come into existence in recent years with respect to the five infrastructure requirements described and analyze weaknesses and points of attack.

As a system, the darknet is subject to a variety of attacks. Legal action continues to be the most powerful challenge to the darknet. However, the darknet is also subject to a variety of other common threats (e.g. viruses, spamming) that, in the past, have led to minor disruptions of the darknet, but could be considerably more damaging.

In this section we consider the potential impact of legal developments on the darknet. Most of our analysis focuses on system robustness, rather than on detailed legal questions. We regard legal questions only with respect to their possible effect: the failure of certain nodes or links (vertices and edges of the graph defined above). In this sense, we are investigating a well known problem in distributed systems.

2.1 Early Small-Worlds Networks

Prior to the mid 1990s, copying was organized around groups of friends and acquaintances. The copied objects were music on cassette tapes and computer programs. The rendering devices were widely-available tape players and the computers of the time – see Fig. 1. Content injection was trivial, since most objects were either not copy protected or, if they were equipped with copy protection mechanisms, the mechanisms were easily defeated. The distribution network was a “sneaker net” of floppy disks and tapes (storage), which were handed in person between members of a group or were sent by postal mail. The bandwidth of this network – albeit small by today’s standards – was sufficient for the objects of the time. The main limitation of the sneaker net with its mechanical transport layer was latency. It could take days or weeks to obtain a copy of an object. Another serious limitation of these networks was the lack of a sophisticated search engine.

There were limited attempts to prosecute individuals who were trying to sell copyrighted objects they had obtained from the darknet (commercial piracy). However, the darknet as a whole was never under significant legal threat. Reasons may have included its limited commercial impact and the protection from legal surveillance afforded by sharing amongst friends.

The sizes of object libraries available on such networks are strongly influenced by the interconnections between the networks. For example, schoolchildren may copy content from their “family network” to their “school network” and thereby increase the size of the darknet object library available to each. Such networks have been studied extensively and are classified as “interconnected small-worlds networks.” [24] There are several popular examples of the characteristics of such systems. For example, most people have a social group of a few score of people. Each of these people has a group of friends that partly overlap with their friends’ friends, and also introduces more people. It is estimated that, on average, each person is connected to every other person in the world by a chain of about six people from which arises the term “six degrees of separation”.

These findings are remarkably broadly applicable (e.g. [20],[3]). The chains are on average so short because certain super-peers have many links. In our example, some people are gregarious and have lots of friends from different social or geographical circles..

We suspect that these findings have implications for sharing on darknets, and we will return to this point when we discuss the darknets of the future later in this paper.

The small-worlds darknet continues to exist. However, a number of technological advances have given rise to new forms of the darknet that have superseded the small-worlds for some object types (e.g. audio).

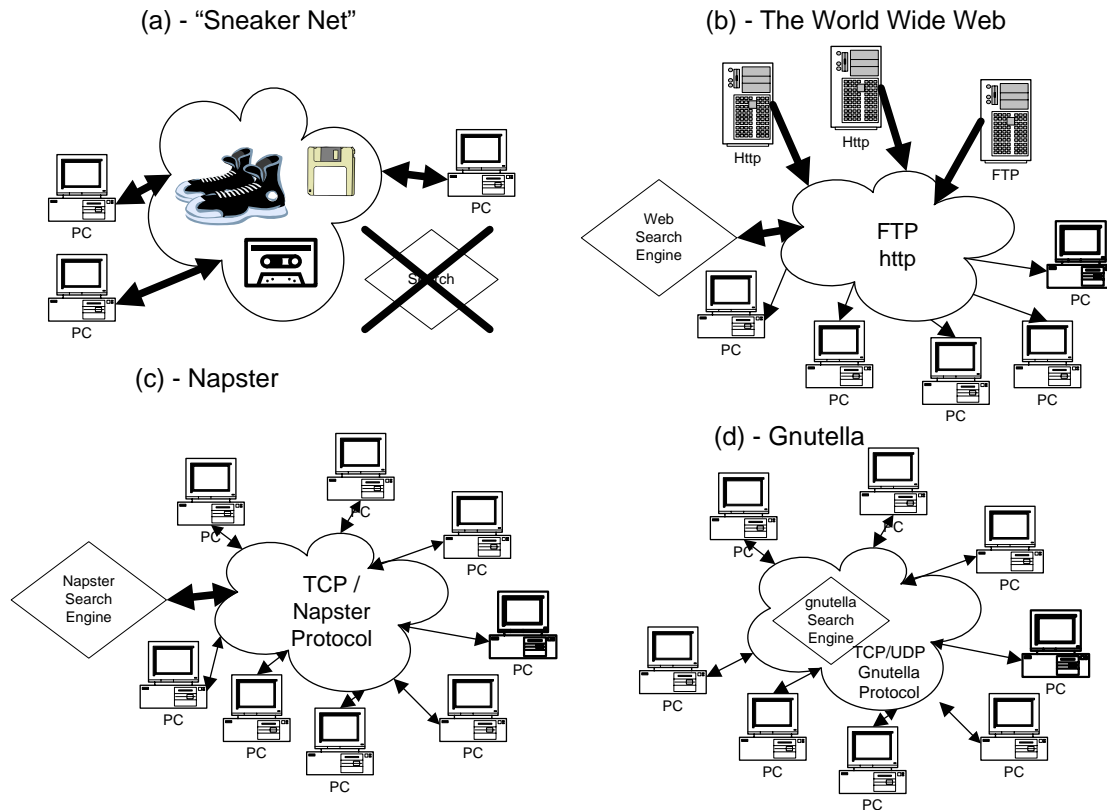


Figure 1: Historical evolution of the Darknet. We highlight the location of the search engine (if present) and the effective bandwidth (thicker lines represent higher bandwidth). Network latencies are not shown, but are much longer for the sneaker net than for the IP-based networks.

2.2 Central Internet Servers

By 1998, a new form of the darknet began to emerge from technological advances in several areas. The internet had become mainstream, and as such its protocols and infrastructure could now be relied upon by anyone seeking to connect users with a centralized service or with each other. The continuing fall in the price of storage together with advances in compression technology had also crossed the threshold at which storing large numbers of audio files was no longer an obstacle to mainstream users. Additionally, the power of computers had crossed the point at which they could be used as rendering devices for multimedia content. Finally, "CD ripping" became a trivial method for content injection.

The first embodiments of this new darknet were central internet servers with large collections of MP3 audio files. A fundamental change that came with these servers was the use of a new distribution network: The internet displaced the sneaker net – at least for audio content. This solved several problems of the old darknet. First, latency was reduced drastically.

Secondly, and more importantly, discovery of objects became much easier because of simple and powerful search mechanisms – most importantly the general-purpose world-wide-web search engine. The local view of the small world was replaced by a global view of the entire collection accessible by all users. The main characteristic of this form of the darknet was centralized storage and search – a simple architecture that mirrored mainstream internet servers.

Centralized or quasi-centralized distribution and service networks make sense for legal online commerce. Bandwidth and infrastructure costs tend to be low, and having customers visit a commerce site means the merchant can display adverts, collect profiles, and bill efficiently. Additionally, management, auditing, and accountability are much easier in a centralized model.

However, centralized schemes work poorly for *illegal* object distribution because large, central servers are large single points of failure: If the distributor is breaking the law, it is relatively easy to force him to stop. Early MP3 Web and FTP sites were commonly “hosted” by universities, corporations, and ISPs. Copyright-holders or their representatives sent “cease and desist” letters to these web-site operators and web-owners citing copyright infringement and in a few cases followed up with legal action [15]. The threats of legal action were successful attacks on those centralized networks, and MP3 web and FTP sites disappeared from the mainstream shortly after they appeared.

2.3 Peer-to-Peer Networks

The realization that centralized networks are not robust to attack (be it legal or technical) has spurred much of the innovation in peer-to-peer networking and file sharing technologies. In this section, we examine architectures that have evolved. Early systems were flawed because critical components remained centralized (Napster) or because of inefficiencies and lack of scalability of the protocol (gnutella) [17]. It should be noted that the problem of object location in a massively distributed, rapidly changing, heterogeneous system was new at the time peer-to-peer systems emerged. Efficient and highly scalable protocols have been proposed since then [9],[23].

2.3.1. Napster

Napster was the service that ignited peer-to-peer file sharing in 1999 [14]. There should be little doubt that a major portion of the massive (for the time) traffic on Napster was of copyrighted objects being transferred in a peer-to-peer model in violation of copyright law. Napster succeeded where central servers had failed by relying on the distributed storage of objects not under the control of Napster. This moved the injection, storage, network distribution, and consumption of objects to users.

However, Napster retained a centralized database² with a searchable index on the file name. The centralized database itself became a legal target [15]. Napster was first enjoined to deny certain queries (e.g. “Metallica”) and then to police its network for all copyrighted content. As the size of the darknet indexed by Napster shrank, so did the number of users. This illustrates a general characteristic of darknets: there is positive feedback between the size of the object library and aggregate bandwidth and the appeal of the network for its users.

2.3.2. Gnutella

The next technology that sparked public interest in peer-to-peer file sharing was Gnutella. In addition to distributed object storage, Gnutella uses a fully distributed database described more fully in [13]. Gnutella does not rely upon any centralized server or service – a peer just needs the IP address of one or a few participating peers to (in principle) reach any host on the Gnutella darknet. Second, Gnutella is not really “run” by anyone: it is an open protocol and anyone can write a Gnutella client application. Finally, Gnutella and its descendants go beyond sharing audio and have substantial non-infringing uses. This changes its legal standing markedly and puts it in a similar category to email. That is, email has substantial non-infringing use, and so email itself is not under legal threat even though it may be used to transfer copyrighted material unlawfully.

² Napster used a farm of weakly coupled databases with clients attaching to just one of the server hosts.

2.4 Robustness of Fully Distributed Darknets

Fully distributed peer-to-peer systems do not present the single points of failure that led to the demise of central MP3 servers and Napster. It is natural to ask how robust these systems are and what form potential attacks could take. We observe the following weaknesses in Gnutella-like systems:

- Free riding
- Lack of anonymity

2.4.1 Free Riding

Peer-to-peer systems are often thought of as fully decentralized networks with copies of objects uniformly distributed among the hosts. While this is possible in principle, in practice, it is not the case. Recent measurements of libraries shared by gnutella peers indicate that the majority of content is provided by a tiny fraction of the hosts [1]. In effect, although gnutella *appears* to be a peer-to-peer network of cooperating hosts, in actual fact it has evolved to effectively be another largely centralized system – see Fig. 2. *Free riding* (i.e. downloading objects without sharing them) by many gnutella users appears to be main cause of this development. Widespread free riding removes much of the power of network dynamics and may reduce a peer-to-peer network into a simple unidirectional distribution system from a small number of sources to a large number of destinations. Of course, if this is the case, then the vulnerabilities that we observed in centralized systems (e.g. FTP-servers) are present again. Free riding and the emergence of super-peers have several causes:

Peer-to-peer file sharing assumes that a significant fraction of users adhere to the somewhat post-capitalist idea of sacrificing their own resources for the “common good” of the network. Most free-riders do not seem to adopt this idea. For example, with 56 kbps modems still being the network connection for most users, allowing uploads constitutes a tangible bandwidth sacrifice. One approach is to make collaboration mandatory. For example, Freenet [6] clients are required to contribute some disk space. However, enforcing such requirements without a central infrastructure is difficult.

Existing infrastructure is another reason for the existence of super-peers. There are vast differences in the resources available to different types of hosts. For example, a T3 connection provides the combined bandwidth of about one thousand 56 kbps telephone connections.

2.4.2 Lack of Anonymity

Users of gnutella who share objects they have stored are not anonymous. Current peer-to-peer networks permit the server endpoints to be determined, and if a peer-client can determine the IP address and affiliation of a peer, then so can a lawyer or government agency. This means that users who share copyrighted objects face some threat of legal action. This appears to be yet another explanation for free riding.

There are some possible technological workarounds to the absence of endpoint anonymity. We could imagine anonymizing routers, overseas routers, object fragmentation, or some other means to complicate the effort required by law-enforcement to determine the original source of the copyrighted bits. For example, Freenet tries to hide the identity of the hosts storing any given object by means of a variety of heuristics, including routing the object through intermediate hosts and providing mechanisms for easy migration of objects to other hosts. Similarly, Mnemosyne [10] tries to organize object storage, such that individual hosts may not know what objects are stored on them. It is conjectured in [10] that this may amount to common-carrier status for the host. A detailed analysis of the legal or technical robustness of these systems is beyond the scope of this paper.

2.4.3 Attacks

In light of these weaknesses, attacks on gnutella-style darknets focus on their object storage and search infrastructures. Because of the prevalence of super-peers, the gnutella darknet depends on a relatively small set of powerful hosts, and these hosts are promising targets for attackers.

Darknet hosts owned by corporations are typically easily removed. Often, these hosts are set up by individual employees without the knowledge of corporate management. Generally corporations respect intellectual property laws. This together with their reluctance to become targets of lawsuits, and their centralized network of hierarchical management makes it relatively easy to remove darknet hosts in the corporate domain.

While the structures at universities are typically less hierarchical and strict than those of corporations, ultimately, similar rules apply. If the .com and .edu T1 and T3 lines were pulled from under a darknet, the usefulness of the network would suffer drastically.

This would leave DSL, ISDN, and cable-modem users as the high-bandwidth servers of objects. We believe limiting hosts to this class would present a far less effective piracy network today from the perspective of acquisition because of the relative rarity of high-bandwidth consumer connections, and hence users would abandon this darknet. However, consumer broadband is becoming more popular, so in the long run it is probable that there will be adequate consumer bandwidth to support an effective consumer darknet.

The obvious next legal escalation is to bring direct or indirect (through the affiliation) challenges against users who share large libraries of copyrighted material. This is already happening and the legal threats or actions appear to be successful [7]. This requires the collaboration of ISPs in identifying their customers, which appears to be forthcoming due to requirements that the carrier must take to avoid liability³ and, in some cases, because of corporate ties between ISPs and content providers. Once again, free riding makes this attack strategy far more tractable.

It is hard to predict further legal escalation, but we note that the DMCA (digital millennium copyright act) is a far-reaching (although not fully tested) example of a law that is potentially quite powerful. We believe it probable that there will be a few more rounds of technical innovations to sidestep existing laws, followed by new laws, or new interpretations of old laws, in the next few years.

³ The Church of Scientology has been aggressive in pursuing ISPs that host its copyright material on newsgroups. The suit that appeared most likely to result in a clear finding, filed against Netcom, was settled out of court. Hence it is still not clear whether an ISP has a responsibility to police the users of its network.

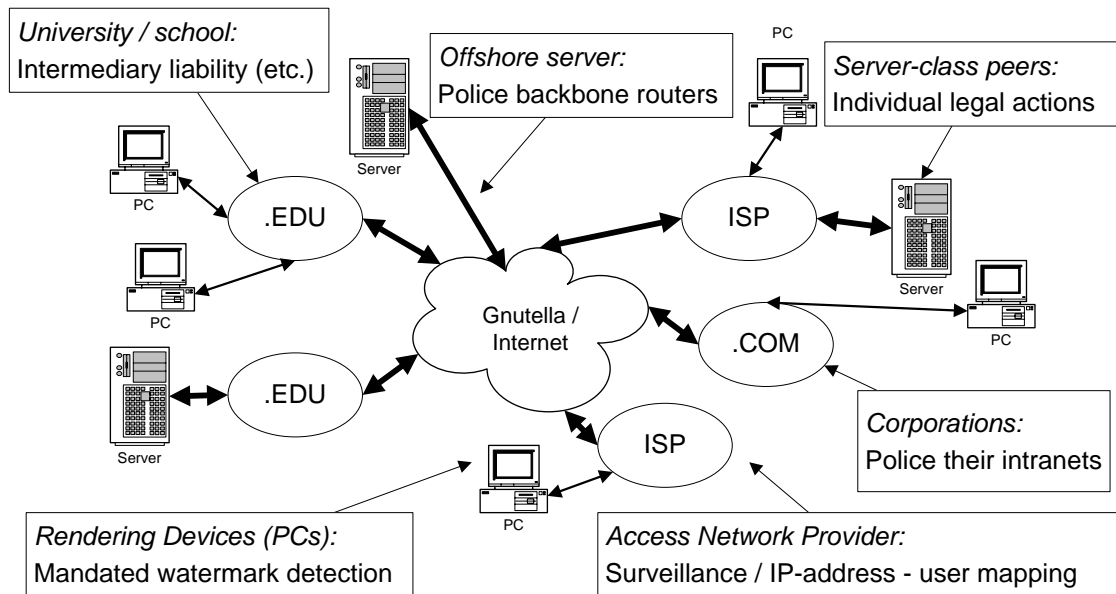


Figure 2: Policing the darknet. Gnutella-style networks appear hard to police because they are highly distributed, and there are thousands or millions of peers. Looking more closely there are several potential vulnerabilities.

2.4.4 Conclusions

All attacks we have identified exploit the lack of endpoint anonymity and are aided by the effects of free riding. We have seen effective legal measures on all peer-to-peer technologies that are used to provide effectively global access to copyrighted material. Centralized web servers were effectively closed down. Napster was effectively closed down. Gnutella and Kazaa are under threat because of free rider weaknesses and lack of endpoint anonymity.

Lack of endpoint anonymity is a direct result of the globally accessible global object database, and it is the existence of the global database that most distinguishes the newer darknets from the earlier small worlds. At this point, it is hard to judge whether the darknet will be able to retain this global database in the long term, but it seems clear that legal setbacks to global-index peer-to-peer will continue to be severe.

However, should Gnutella-style systems become unviable as darknets, systems, such as Freenet or Mnemosyne might take their place. Peer-to-peer networking and file sharing does seem to be entering into the mainstream – both for illegal and legal uses. If we couple this with the rapid build-out of consumer broadband, the dropping price of storage, and the fact that personal computers are effectively establishing themselves as centers of home-entertainment, we suspect that peer-to-peer functionality will remain popular and become more widespread.




Figure 3: Interconnected small-worlds darknets. Threats of surveillance close global darknets. Darknets form around social groups, but use high-bandwidth, low latency communications (the internet) and are supported by search engines. Custom applications, Instant-Messenger-style applications or simple shared file-systems host the darknet. People's social groups overlap so objects available in one darknet diffuse to others: in the terminology used in this paper, each peer that is a member of more than one darknet is an introduction host for objects obtained from other darknets.

2.5 Small Worlds Networks Revisited

In this section we try to predict the evolution of the darknet should global peer-to-peer networks be effectively stopped by legal means. The globally accessible global database is the only infrastructure component of the darknet that can be disabled in this way. The other enabling technologies of the darknet (injection, distribution networks, rendering devices, storage) will not only remain available, but rapidly increase in power, based on general technological advances and the possible incorporation of cryptography. We stress that the networks described in this section (in most cases) provide poorer services than global network, and would only arise in the absence of a global database.

In the absence of a global database, small-worlds networks could again become the prevalent form of the darknet. However, these small-worlds will be more powerful than they were in the past. With the widespread availability of cheap CD and DVD readers and writers as well as large hard disks, the bandwidth of the sneaker net has increased dramatically, the cost of object storage has become negligible and object injection tools have become ubiquitous. Furthermore, the internet is available as a distribution mechanism that is adequate for audio for most users, and is becoming increasingly adequate for video and computer programs. In light of strong cryptography, it is hard to imagine how sharing could be observed and prosecuted as long as users do not share with strangers.

In concrete terms, students in dorms will establish darknets to share content in their social group. These darknets may be based on simple file sharing, DVD-copying, or may use special application programs or servers: for example, a chat or instant-messenger client enhanced to share content with members of your buddy-list. Each student will be a member of other darknets: for example, their family, various special interest groups, friends from high-school, and colleagues in part-time jobs (Fig. 3). If there are a few active super-peers - users that locate and share objects with zeal - then we can anticipate that content will rapidly diffuse between darknets, and relatively small darknets arranged around social groups will approach the aggregate libraries that are provided by the global darknets of today. Since the legal exposure of such sharing is quite limited, we believe that sharing amongst socially oriented groups will increase unabated.

Small-worlds networks suffer somewhat from the lack of a global database; each user can only see the objects stored by his small world neighbors. This raises a number of interesting questions about the network structure and object flow:

- What graph structure will the network have? For example, will it be connected? What will be the average distance between two nodes?

- Given a graph structure, how will objects propagate through the graph? In particular, what fraction of objects will be available at a given node? How long does it take for objects to propagate (diffuse) through the network?

Questions of this type have been studied in different contexts in a variety of fields (mathematics, computer science, economics, and physics). A number of empirical studies seek to establish structural properties of different types of small world networks, such as social networks [20] and the world-wide web [3]. These works conclude that the diameter of the examined networks is small, and observe further structural properties, such as a power law of the degree distribution [5]. A number of authors seek to model these networks by means of random graphs, in order to perform more detailed mathematical analysis on the models [2],[8],[21],[22] and, in particular, study the possibility of efficient search under different random graph distributions [18],[19]. We will present a quantitative study of the structure and dynamics of small-worlds networks in an upcoming paper, but to summarize, small-worlds darknets can be extremely efficient for popular titles: very few peers are needed to satisfy requests for top-20 books, songs, movies or computer programs. If darknets are interconnected, we expect the effective introduction rate to be large. Finally, if darknet clients are enhanced to actively seek out new popular content, as opposed to the user-demand based schemes of today, small-worlds darknets will be very efficient.

3 Introducing Content into the Darknet

Our analysis and intuition have led us to believe that efficient darknets – in global or small-worlds form -- will remain a fact of life. In this section we examine rights-management technologies that are being deployed to limit the introduction rate or decrease the rate of diffusion of content into the darknet.

3.1 Conditional Access Systems

A conditional-access system is a simple form of rights-management system in which subscribers are given access to objects based (typically) on a service contract. Digital rights management systems often perform the same function, but typically impose restrictions on the use of objects after unlocking.

Conditional access systems such as cable, satellite TV, and satellite radio offer little or no protection against objects being introduced into the darknet from subscribing hosts. A conditional-access system customer has no access to channels or titles to which they are not entitled, and has essentially free use of channels that he has subscribed or paid for. This means that an investment of ~\$100 (at time of writing) on an analog video-capture card is sufficient to obtain and share TV programs and movies. Some CA systems provide post-unlock protections but they are generally cheap and easy to circumvent.

Thus, conditional access systems provide a widely deployed, high-bandwidth source of video material for the darknet. In practice, the large size and low cost of CA-provided video content will limit the exploitation of the darknet for distributing video in the near-term.

The same can *not* be said of the use of the darknet to distribute conditional-access system broadcast keys. At some level, each head-end (satellite or cable TV head-end) uses an encryption key that must be made available to each customer (it is a broadcast), and in the case of a satellite system this could be millions of homes. CA-system providers take measures to limit the usefulness of exploited session keys (for example, they are changed every few seconds), but if darknet latencies are low, or if encrypted broadcast data is cached, then the darknet could threaten CA-system revenues.

We observe that the exposure of the conditional access provider to losses due to piracy is proportional to the number of customers that share a session key. In this regard, cable-operators are in a safer position than satellite operators because a cable operator can narrowcast more cheaply.

3.2 DRM Systems

A classical-DRM system is one in which a client obtains content in protected (typically encrypted) form, with a license that specifies the uses to which the content may be put. Examples of licensing terms that are being explored by the industry are “play on these three hosts,” “play once,” “use computer program for one hour,” etc.

The license and the wrapped content are presented to the DRM system whose responsibility is to ensure that:

- a) The client cannot remove the encryption from the file and send it to a peer,
- b) The client cannot “clone” its DRM system to make it run on another host,
- c) The client obeys the rules set out in the DRM license, and,
- d) The client cannot separate the rules from the payload.

Advanced DRM systems may go further.

Some such technologies have been commercially very successful – the content scrambling system used in DVDs, and (broadly interpreted) the protection schemes used by conditional access system providers fall into this category, as do newer DRM systems that use the internet as a distribution channel and computers as rendering devices. These technologies are appealing because they promote the establishment of new businesses, and can reduce distribution costs. If costs and licensing terms are appealing to producers and consumers, then the vendor thrives. If the licensing terms are unappealing or inconvenient, the costs are too high, or competing systems exist, then the business will fail. The DivX “DVD” rental model failed on most or all of these metrics, but CSS-protected DVDs succeeded beyond the wildest expectations of the industry.

On personal computers, current DRM systems are software-only systems using a variety of tricks to make them hard to subvert. DRM enabled consumer electronics devices are also beginning to emerge.

In the absence of the darknet, the goal of such systems is to have comparable security to competing distribution systems – notably the CD and DVD – so that programmable computers can play an increasing role in home entertainment. We will speculate whether these strategies will be successful in the Sect. 5.

DRM systems strive to be BOBE (break-once, break everywhere)-resistant. That is, suppliers anticipate (and the assumptions of the darknet predict) that individual instances (clients) of all security-systems, whether based on hardware or software, will be subverted. If a client of a system is subverted, then all content protected by that DRM client can be unprotected. If the break can be applied to *any other* DRM client of that class so that all of those users can break their systems, then the DRM-scheme is BOBE-weak. If, on the other hand, knowledge gained breaking one client cannot be applied elsewhere, then the DRM system is BOBE-strong.

Most commercial DRM-systems have BOBE-exploits, and we note that the darknet applies to DRM-hacks as well. The CSS system is an exemplary BOBE-weak system. The knowledge and code that comprised the De-CSS exploit spread uncontrolled around the world on web-sites, newsgroups, and even T-shirts, in spite of the fact that, in principle, the Digital Millennium Copyright Act makes it a crime to develop these exploits.

A final characteristic of existing DRM-systems is *renewability*. Vendors recognize the possibility of exploits, and build systems that can be field-updated.

It is hard to quantify the effectiveness of DRM-systems for restricting the introduction of content into the darknet from experience with existing systems. Existing DRM-systems typically provide protection for months to years; however, the content available to such systems has to date been of minimal interest, and the content that *is* protected is also available in unprotected form. The

one system that was protecting valuable content (DVD video) was broken very soon after compression technology and increased storage capacities and bandwidth enabled the darknet to carry video content.

3.3 Software

The DRM-systems described above can be used to provide protection for software, in addition other objects (e.g. audio and video). Alternatively, copy protection systems for computer programs may embed the copy protection code in the software itself.

The most important copy-protection primitive for computer programs is for the software to be bound to a host in such a way that the program will not work on an unlicensed machine. Binding requires a machine ID: this can be a unique number on a machine (e.g. a network card MAC address), or can be provided by an external dongle.

For such schemes to be strong, two things must be true. First, the machine ID must not be “virtualizable.” For instance, if it is trivial to modify a NIC driver to return an invalid MAC address, then the software-host binding is easily broken. Second, the code that performs the binding checks must not be easy to patch. A variety of technologies that revolve around software tamper-resistance can help here [4].

We believe that binding software to a host is a more tractable problem than protecting passive content, as the former only requires tamper resistance, while the latter also requires the ability to hide and manage secrets. However, we observe that all software copy-protection systems deployed thus far have been broken. The definitions of BOBE-strong and BOBE-weak apply similarly to software. Furthermore, software is as much subject to the dynamics of the darknet as passive content.

4 Policing Hosts

If there are subverted hosts, then content will leak into the darknet. If the darknet is efficient, then content will be rapidly propagated to all interested peers. In the light of this, technologists are looking for alternative protection schemes. In this section we will evaluate watermarking and fingerprinting technologies.

4.1 Watermarking

Watermarking embeds an “indelible” invisible mark in content. A plethora of schemes exist for audio/video and still image content and computer programs.

There are a variety of schemes for exploiting watermarks for content-protection. Consider a rendering device that locates and interprets watermarks. If a watermark is found then special action is taken. Two common actions are:

- 1) *Restrict behavior:* For example, a bus-adaptor may refuse to pass content that has the “copy once” and “already copied once” bits set.
- 2) *Require a license to play:* For example, if a watermark is found indicating that content is rights-restricted then the renderer may demand a license indicating that the user is authorized to play the content.

Such systems were proposed for audio content – for example the secure digital music initiative (SDMI) [16], and are under consideration for video by the copy-protection technical working group (CPTWG) [12].

There are several reasons why it appears unlikely that such systems will ever become an effective anti-piracy technology. From a commercial point of view, building a watermark detector into a device renders it strictly less useful for consumers than a competing product that does not.

This argues that watermarking schemes are unlikely to be widely deployed, unless mandated by legislation. The recently proposed Hollings bill is a step along these lines [11].

We contrast watermark-based policing with classical DRM: If a general-purpose device is equipped with a classical DRM-system, it can play all content acquired from the darknet, *and* have access to new content acquired through the DRM-channel. This is in stark distinction to reduction of functionality inherent in watermark-based policing.

Even if watermarking systems were mandated, this approach is likely to fail due to a variety of technical inadequacies. The first inadequacy concerns the robustness of the embedding layer. We are not aware of systems for which simple data transformations cannot strip the mark or make it unreadable. Marks can be made more robust, but in order to recover marks after adversarial manipulation, the reader must typically search a large phase space, and this quickly becomes untenable. In spite of the proliferation of proposed watermarking schemes, it remains doubtful whether robust embedding layers for the relevant content types can be found.

A second inadequacy lies in unrealistic assumptions about key management. Most watermarking schemes require widely deployed cryptographic keys. Standard watermarking schemes are based on the normal cryptographic principles of a public algorithm and secret keys. Most schemes use a shared-key between marker and detector. In practice, this means that all detectors need a private key, and, typically, share a single private key. It would be naïve to assume that these keys will remain secret for long in an adversarial environment. Once the key or keys are compromised, the darknet will propagate them efficiently, and the scheme collapses. There have been proposals for public-key watermarking systems. However, so far, this work does not seem practical and the corresponding schemes do not even begin to approach the robustness of the cryptographic systems whose name they borrow.

A final consideration bears on the location of mandatory watermark detectors in client devices. On open computing devices (e.g. personal computers), these detectors could, in principle, be placed in software or in hardware. Placing detectors in software would be largely meaningless, as circumvention of the detector would be as simple as replacing it by a different piece of software. This includes detectors placed in the operating system, all of whose components can be easily replaced, modified and propagated over the darknet.

Alternatively, the detectors could be placed in hardware (e.g. audio and video cards). In the presence of the problems described this would lead to untenable renewability problems --- the hardware would be ineffective within days of deployment. Consumers, on the other hand, expect the hardware to remain in use for many years. Finally, consumers themselves are likely to rebel against “footing the bill” for these ineffective content protection systems. It is virtually certain, that the darknet would be filled with a continuous supply of watermark removal tools, based on compromised keys and weaknesses in the embedding layer. Attempts to force the public to “update” their hardware would not only be intrusive, but impractical.

In summary, attempts to mandate content protection systems based on watermark detection at the consumer’s machine suffer from commercial drawbacks and severe technical deficiencies. These schemes, which aim to provide content protection beyond DRM by attacking the darknet, are rendered entirely ineffective by the presence of even a moderately functional darknet.

4.2 Fingerprinting

Fingerprint schemes are based on similar technologies and concepts to watermarking schemes. However, whereas watermarking is designed to perform *a-priori* policing, fingerprinting is designed to provide *a-posteriori* forensics.

In the simplest case, fingerprinting is used for individual-sale content (as opposed to super-distribution or broadcast – although it can be applied there with some additional assumptions). When a client purchases an object, the supplier marks it with an individualized mark that identifies

the purchaser. The purchaser is free to use the content, but if it appears on a darknet, a policeman can identify the source of the content and the offender can be prosecuted.

Fingerprinting suffers from fewer technical problems than watermarking. The main advantage is that no widespread key-distribution is needed – a publisher can use whatever secret or proprietary fingerprinting technology they choose, and is entirely responsible for the management of their own keys.

Fingerprinting has one problem that is not found in watermarking. Since each fingerprinted copy of a piece of media is different, if a user can obtain several different copies, he can launch collusion attacks (e.g. averaging). In general, such attacks are very damaging to the fingerprint payload.

It remains to be seen whether fingerprinting will act as a deterrent to theft. There is currently no legal precedent for media fingerprints being evidence of crime, and this case will probably be hard to make – after all, detection is a statistical process with false positives, and plenty of opportunity for deniability. However, we anticipate that there will be uneasiness in sharing a piece of content that may contain a person's identity, and that ultimately leaves that person's control.

Note also that with widely distributed watermarking detectors, it is easy to see whether you have successfully removed a watermark. There is no such assurance for determining whether a fingerprint has been successfully removed from an object because users are not necessarily knowledgeable about the fingerprint scheme or schemes in use. However, if it turns out that the deterrence of fingerprinting is small (i.e. everyone shares their media regardless of the presence of marks), there is probably no reasonable legal response. Finally, distribution schemes in which objects must be individualized will be expensive.

5 Conclusions

There seem to be no technical impediments to darknet-based peer-to-peer file sharing technologies growing in convenience, aggregate bandwidth and efficiency. The legal future of darknet-technologies is less certain, but we believe that, at least for some classes of user, and possibly for the population at large, efficient darknets will exist. The rest of this section will analyze the implications of the darknet from the point of view of individual technologies and of commerce in digital goods.

5.1 Technological Implications

DRM systems are limited to protecting the content they contain. Beyond our first assumption about the darknet, the darknet is not impacted by DRM systems. In light of our first assumption about the darknet, DRM design details, such as properties of the tamper-resistant software may be strictly less relevant than the question whether the current darknet has a global database. In the presence of an infinitely efficient darknet – which allows instantaneous transmission of objects to all interested users – even sophisticated DRM systems are inherently ineffective. On the other hand, if the darknet is made up of isolated small worlds, even BOBE-weak DRM systems are highly effective. The interesting cases arise between these two extremes – in the presence of a darknet, which is connected, but in which factors, such as latency, limited bandwidth or the absence of a global database limit the speed with which objects propagate through the darknet. It appears that quantitative studies of the effective “diffusion constant” of different kinds of darknets would be highly useful in elucidating the dynamics of DRM and the darknet.

Proposals for systems involving mandatory watermark detection in rendering devices try to impact the effectiveness of the darknet directly by trying to detect and eliminate objects that originated in the darknet. In addition to severe commercial and social problems, these schemes suffer from several technical deficiencies, which, in the presence of an effective darknet, lead to their complete collapse. We conclude that such schemes are doomed to failure.

5.2 Business in the Face of the Darknet

There is evidence that the darknet will continue to exist and provide low cost, high-quality service to a large group of consumers. This means that in many markets, the darknet will be a competitor to legal commerce. From the point of view of economic theory, this has profound implications for business strategy: for example, increased security (e.g. stronger DRM systems) may act as a *disincentive* to legal commerce. Consider an MP3 file sold on a web site: this costs money, but the purchased object is as useful as a version acquired from the darknet. However, a securely DRM-wrapped song is strictly *less* attractive: although the industry is striving for flexible licensing rules, customers *will* be restricted in their actions if the system is to provide meaningful security. This means that a vendor will probably make more money by selling unprotected objects than protected objects. In short, if you are competing with the darknet, you must compete on the darknet's own terms: that is convenience and low cost rather than additional security.

Certain industries have faced this (to a greater or lesser extent) in the past. Dongle-protected computer programs lost sales to unprotected programs, or hacked versions of the program. Users have also refused to upgrade to newer software versions that are copy protected.

There are many factors that influence the threat of the darknet to an industry. We see the darknet having most direct bearing on mass-market consumer IP-goods. Goods sold to corporations are less threatened because corporations mostly try to stay legal, and will police their own intranets for illicit activities. Additionally, the cost-per-bit, and the total size of the objects have a huge bearing on the competitiveness of today's darknets compared with legal trade. For example, today's peer-to-peer technologies provide excellent service quality for audio files, but users must be very determined or price-sensitive to download movies from a darknet, when the legal competition is a rental for a few dollars.

References

- [1] E. Adar and B. A. Huberman, *Free Riding on Gnutella*, http://www.firstmonday.dk/issues/issue5_10/adar/index.html
- [2] W. Aiello, F. Chung and L. Lu, *Random evolution in massive graphs*, In Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science, pages 510—519, 2001.
- [3] R. Albert, H. Jeong and A.-L. Barabási, *Diameter of the world-wide web*, Nature 401, pages 130—131, 1999.
- [4] D. Aucsmith, *Tamper Resistant Software, An Implementation*, Information Hiding 1996, Proceedings: Springer 1998.
- [5] A.-L. Barabási, R. Albert, *Emergence of scaling in random networks*, Science 286, pages 509—512, 1999.
- [6] I. Clarke, O. Sandberg, B. Wiley and T. Hong, *Freenet: A distributed information storage and retrieval system*, International Workshop on Design Issues in Anonymity and Unobservability, 2000.
- [7] R. Clarke, *A defendant class action lawsuit* <http://www.kentlaw.edu/perritt/honorsscholars/clarke.html>
- [8] C. Cooper and A. Frieze, *A general model of web graphs*, Proceedings of ESA 2001, pages 500-511, 2001.
- [9] F. Dabek, E. Brunskill, M. F. Kaashoek, D. Karger, R. Morris, I. Stoica and H. Balakrishnan, *Building peer-to-peer systems with Chord, a distributed lookup service*, In Proceedings of the Eighth IEEE Workshop on Hot Topics in Operating Systems (HotOS-VIII), pages 81—86, 2001.
- [10] S. Hand and T. Roscoe, *Mnemosyne: peer-to-peer steganographic storage*, In Proceedings of the First International Workshop on Peer-to-Peer Systems, 2002.

- [11] Senator Fritz Hollings, *Consumer Broadband and Digital Television Promotion Act*.
- [12] <http://www.cptwg.org>
- [13] <http://www.gnutelladev.com/protocol/gnutella-protocol.html>
- [14] <http://www.napster.com>
- [15] <http://www.riaa.org>
- [16] <http://www.sdmi.org>
- [17] M. Javanovic, F. Annexstein and K. Berman, *Scalability Issues in Large Peer-to-Peer Networks - A Case Study of Gnutella*, ECECS Department, University of Cincinnati, Cincinnati, OH 45221
- [18] J. Kleinberg, Navigation in a small world, *Nature* 406, 2000.
- [19] J. Kleinberg, *Small-world phenomena and the dynamics of information*, *Advances in Neural Information Processing Systems (NIPS)* 14, 2001.
- [20] S. Milgram, The small world problem, *Psychology Today*, vol. 2, pages 60—67, 1967.
- [21] M. Newman, *Small worlds: the structure of social networks*, Santa Fe Institute, Technical Report 99-12-080, 1999.
- [22] M. Newman, D. Watts and S. Strogatz, *Random graph models of social networks*, *Proc. Natl. Acad. Sci. USA* 99, pages 2566—2572, 2002.
- [23] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, H. Balakrishnan, *CHORD: A scalable peer-to-peer lookup service for internet applications*, In *Proceedings of the ACM SIGCOMM 2001 Conference SIGCOMM-01*, pages 149—160, 2001.
- [24] D. J. Watts and S. H. Strogatz, *Collective dynamics of small-world networks*, *Nature*, 393:440-442, June 1998.

This is a copy of <http://www.dashes.com/ani/stuff/doctorow-drm-ms.html>, converted into pdf for inclusion in a course reader.

Microsoft Research DRM talk

by

Cory Doctorow (cory@eff.org)

June 17, 2004

This talk was originally given to Microsoft's Research Group and other interested parties from within the company at their Redmond offices on June 17, 2004. (See [public domain notice](#).)

EFF website: [EFF \(Donate\)](#)

Cory Doctorow's personal site: [Craphound](#)

Cory Doctorow's group weblog: [BoingBoing](#)

Introduction

Greetings fellow pirates! Arrrrr!

I'm here today to talk to you about copyright, technology and DRM, I work for the [Electronic Frontier Foundation](#) on copyright stuff (mostly), and I live in London. I'm not a lawyer—I'm a kind of mouthpiece/activist type, though occasionally they shave me and stuff me into my Bar Mitzvah suit and send me to a standards body or the UN to stir up trouble. I spend about three weeks a month on the road doing completely weird stuff like going to Microsoft to talk about DRM.

I lead a double life: I'm also a [science fiction writer](#). That means I've got a dog in this fight, because I've been dreaming of making my living from writing since I was 12 years old. Admittedly, my IP-based biz isn't as big as yours, but I guarantee you that it's every bit as important to me as yours is to you.

Here's what I'm here to convince you of:

1. That DRM systems [don't work](#)
2. That DRM systems are [bad for society](#)
3. That DRM systems are [bad for business](#)
4. That DRM systems are [bad for artists](#)
5. That DRM is a [bad business-move for MSFT](#)

It's a big brief, this talk. Microsoft has sunk a lot of capital into DRM systems, and spent a lot of time sending folks like Martha and Brian and Peter around to various smoke-filled rooms to make sure that Microsoft DRM finds a hospitable home in the future world. Companies like Microsoft steer like old Buicks, and this issue has a lot of forward momentum that will be hard to soak up without driving the engine block back into the driver's compartment. At best I think that Microsoft might convert some of that momentum on DRM into angular momentum, and in so doing, save all our asses.

Let's dive into it.

1. DRM systems don't work

This bit breaks down into two parts:

1. A quick refresher course in crypto theory

2. Applying that to DRM

Cryptography—secret writing—is the practice of keeping secrets. It involves three parties: a sender, a receiver and an attacker (actually, there can be more attackers, senders and recipients, but let's keep this simple). We usually call these people Alice, Bob and Carol.

Let's say we're in the days of the Caesar, the Gallic War. You need to send messages back and forth to your generals, and you'd prefer that the enemy doesn't get hold of them. You can rely on the idea that anyone who intercepts your message is probably illiterate, but that's a tough bet to stake your empire on. You can put your messages into the hands of reliable messengers who'll chew them up and swallow them if captured—but that doesn't help you if Brad Pitt and his men in skirts skewer him with an arrow before he knows what's hit him.

So you encipher your message with something like [ROT-13](#), where every character is rotated halfway through the alphabet. They used to do this with non-worksafe material on Usenet, back when anyone on Usenet cared about work-safe-ness—A would become N, B is O, C is P, and so forth. To decipher, you just add 13 more, so N goes to A, O to B yadda yadda.

Well, this is pretty lame: as soon as anyone figures out your algorithm, your secret is gonezored.

So if you're Caesar, you spend a lot of time worrying about keeping the existence of your messengers and their payloads secret. Get that? You're Augustus and you need to send a message to Brad without Caceous (a word I'm reliably informed means “cheese-like, or pertaining to cheese”) getting his hands on it. You give the message to Diatomaceous, the fleetest runner in the empire, and you encipher it with ROT-13 and send him out of the garrison in the pitchest hour of the night, making sure no one knows that you've sent it out. Caceous has spies everywhere, in the garrison and staked out on the road, and if one of them puts an arrow through Diatomaceous, they'll have their hands on the message, and then if they figure out the cipher, you're borked. So the existence of the message is a secret. The cipher is a secret. The ciphertext is a secret. That's a lot of secrets, and the more secrets you've got, the less secure you are, especially if any of those secrets are shared. Shared secrets aren't really all that secret any longer.

Time passes, stuff happens, and then Tesla invents the radio and Marconi takes credit for it. This is both good news and bad news for crypto: on the one hand, your messages can get to anywhere with a receiver and an antenna, which is great for the brave fifth columnists working behind the enemy lines. On the other hand, anyone with an antenna can listen in on the message, which means that it's no longer practical to keep the existence of the message a secret. Any time Adolf sends a message to Berlin, he can assume Churchill overhears it.

Which is OK, because now we have computers — big, bulky primitive mechanical computers, but computers still. Computers are machines for rearranging numbers, and so scientists on both sides engage in a fiendish competition to invent the most cleverest method they can for rearranging numerically represented text so that the other side can't unscramble it. The existence of the message isn't a secret anymore, but the cipher is.

But this is still too many secrets. If Bobby intercepts one of Adolf's Enigma machines, he can give Churchill all kinds of intelligence. I mean, this was good news for Churchill and us, but bad news for Adolf. And at the end of the day, it's bad news for anyone who wants to keep a secret.

Enter keys: a cipher that uses a key is still more secure. Even if the cipher is disclosed, even if the ciphertext is intercepted, without the key (or a break), the message is secret. Post-war, this is doubly important as we begin to realize what I think of as Schneier's Law: “any person can invent a security system so clever that she or he can't think of how to break it.” This means that the only experimental methodology for discovering if you've made mistakes in your cipher is to tell all the smart people you can about it and ask them to think of ways to break it. Without this critical step, you'll eventually end up living

in a fool's paradise, where your attacker has broken your cipher ages ago and is quietly decrypting all her intercepts of your messages, snickering at you.

Best of all, there's only one secret: the key. And with dual-key crypto it becomes a lot easier for Alice and Bob to keep their keys secret from Carol, even if they've never met. So long as Alice and Bob can keep their keys secret, they can assume that Carol won't gain access to their cleartext messages, even though she has access to the cipher and the ciphertext. Conveniently enough, the keys are the shortest and simplest of the secrets, too: hence even easier to keep away from Carol. Hooray for Bob and Alice.

Now, let's apply this to DRM.

In DRM, the attacker is *also the recipient*. It's not Alice and Bob and Carol, it's just Alice and Bob. Alice sells Bob a DVD. She sells Bob a DVD player. The DVD has a movie on it—say, *Pirates of the Caribbean*—and it's enciphered with an algorithm called CSS — Content Scrambling System. The DVD player has a CSS un-scrambler.

Now, let's take stock of what's a secret here: the cipher is well-known. The ciphertext is most assuredly in enemy hands, arrr. So what? As long as the key is secret from the attacker, we're golden.

But there's the rub. Alice wants Bob to buy *Pirates of the Caribbean* from her. Bob will only buy *Pirates of the Caribbean* if he can descramble the CSS-encrypted VOB—video object—on his DVD player. Otherwise, the disc is only useful to Bob as a drinks-coaster. So Alice has to provide Bob—the attacker—with the key, the cipher and the ciphertext.

Hilarity ensues.

DRM systems are broken in minutes, sometimes days. Rarely, months. It's not because the people who think them up are stupid. It's not because the people who break them are smart. It's not because there's a flaw in the algorithms. At the end of the day, all DRM systems share a common vulnerability: they provide their attackers with ciphertext, the cipher and the key. At this point, the secret isn't a secret anymore.

2. DRM systems are bad for society

Raise your hand if you're thinking something like, “But DRM doesn't have to be proof against smart attackers, only average individuals! It's like a speedbump!”

Put your hand down.

This is a fallacy for two reasons: one technical, and one social. They're both bad for society, though.

Here's the technical reason: I don't need to be a cracker to break your DRM. I only need to know how to search Google, or Kazaa, or any of the other general-purpose search tools for the cleartext that someone smarter than me has extracted.

Raise your hand if you're thinking something like, “But NGSCB can solve this problem: we'll lock the secrets up on the logic board and goop it all up with epoxy.”

Put your hand down.

Raise your hand if you're a co-author of the [Darknet paper](#).

Everyone in the first group, meet the co-authors of the Darknet paper. This is a paper that says, among other things, that DRM will fail for this very reason. Put your hands down, guys.

Here's the social reason that DRM fails: keeping an honest user honest is like keeping a tall user tall. DRM vendors tell us that their technology is meant to be proof against average users, not organized criminal gangs like the Ukrainian pirates who stamp out millions of high-quality counterfeits. It's not meant to be proof against sophisticated college kids. It's not meant to be proof against anyone who knows how to edit her registry, or hold down the shift key at the right moment, or use a search engine. At the end of the day, the user DRM is meant to defend against is the most unsophisticated and least capable among us.

Here's a true story about a user I know who was stopped by DRM. She's smart, college educated, and knows nothing about electronics. She has three kids. She has a DVD in the living room and an old VHS deck in the kids' playroom. One day, she brought home the *Toy Story* DVD for the kids. That's a substantial investment, and given the generally jam-smear character of everything the kids get their paws on, she decided to tape the DVD off to VHS and give that to the kids—that way she could make a fresh VHS copy when the first one went south. She cabled her DVD into her VHS and pressed play on the DVD and record on the VCR and waited.

Before I go farther, I want us all to stop a moment and marvel at this. Here is someone who is practically technophobic, but who was able to construct a mental model of sufficient accuracy that she figured out that she could connect her cables in the right order and dub her digital disc off to analog tape. I imagine that everyone in this room is the front-line tech support for someone in her or his family: would it be great if all our non-geek friends and relatives were this clever and imaginative?

I also want to point out that this is the proverbial honest user. She's not making a copy for the next door neighbors. She's not making a copy and selling it on a blanket on Canal Street. She's not ripping it to her hard-drive, DivX encoding it and putting it in her Kazaa sharepoint. She's doing something **honest** — moving it from one format to another. She's home taping.

Except she fails. There's a DRM system called Macrovision embedded – by law – in every DVD player and VHS that messes with the vertical blanking interval in the signal and causes any tape made in this fashion to fail. Macrovision can be defeated for about \$10 with a gadget readily available on eBay. But our infringer doesn't know that. She's "honest." Technically unsophisticated. Not stupid, mind you – just naive.

The Darknet paper addresses this possibility: it even predicts what this person will do in the long run: she'll find out about Kazaa and the next time she wants to get a movie for the kids, she'll download it from the net and burn it for them.

In order to delay that day for as long as possible, our lawmakers and big rights-holder interests have come up with a disastrous policy called anticircumvention.

Here's how anticircumvention works: if you put a lock – an access control – around a copyrighted work, it is illegal to break that lock. It's illegal to make a tool that breaks that lock. It's illegal to tell someone how to make that tool. It's illegal to tell someone where she can find out how to make that tool.

Remember [Schneier's Law](#)? Anyone can come up with a security system so clever that he can't see its flaws. The only way to find the flaws in security is to disclose the system's workings and invite public feedback. But now we live in a world where any cipher used to fence off a copyrighted work is off-limits to that kind of feedback. That's something that a Princeton engineering prof named Ed Felten discovered when he submitted a paper to an academic conference on the failings in the Secure Digital Music Initiative, a watermarking scheme proposed by the recording industry. The RIAA responded by threatening to sue his ass if he tried it. We fought them because Ed is the kind of client that impact litigators love: unimpeachable and clean-cut and the RIAA folded. Lucky Ed. Maybe the next guy isn't so lucky.

Matter of fact, the next guy wasn't. [Dmitry Sklyarov](#) is a Russian programmer who gave a talk at a hacker con in Vegas on the failings in Adobe's e-book locks. The FBI threw him in the slam for 30 days. He copped a plea, went home to Russia, and the Russian equivalent of the State Department issued a blanket warning to its researchers to stay away from American conferences, since we'd apparently turned into the kind of country where certain equations are illegal.

Anticircumvention is a powerful tool for people who want to exclude competitors. If you claim that your car engine firmware is a "copyrighted work," you can sue anyone who makes a tool for interfacing with it. That's not just bad news for mechanics—think of the hotrodders who want to chip their cars to tweak the performance settings. We have companies like Lexmark claiming that their printer cartridges contain copyrighted works—software that trips an "I am empty" flag when the toner runs out, and have sued a competitor who made a remanufactured cartridge that reset the flag. Even garage-door opener companies have gotten in on the act, claiming that their receivers' firmware are copyrighted works. Copyrighted cars, print carts and garage-door openers: what's next, copyrighted light-fixtures?

Even in the context of legitimate—excuse me, "traditional"—copyrighted works like movies on DVDs, anticircumvention is bad news. Copyright is a delicate balance. It gives creators and their assignees some rights, but it also reserves some rights to the public. For example, an author has no right to prohibit anyone from transcoding his books into assistive formats for the blind. More importantly, though, a creator has a very limited say over what you can do once you lawfully acquire her works. If I buy your book, your painting, or your DVD, it belongs to me. It's my property. Not my "intellectual property"—a whacky kind of pseudo-property that's swiss-cheesed with exceptions, easements and limitations—but real, no-fooling, actual tangible *property*—the kind of thing that courts have been managing through tort law for centuries.

But anticircumvention lets rightsholders invent new and exciting copyrights for themselves—to write private laws without accountability or deliberation—that expropriate your interest in your physical property to their favor. Region-coded DVDs are an example of this: there's no copyright here or in anywhere I know of that says that an author should be able to control where you enjoy her creative works, once you've paid for them. I can buy a book and throw it in my bag and take it anywhere from Toronto to Timbuktu, and read it wherever I am: I can even buy books in America and bring them to the UK, where the author may have an exclusive distribution deal with a local publisher who sells them for double the US shelf-price. When I'm done with it, I can sell it on or give it away in the UK. Copyright lawyers call this "First Sale," but it may be simpler to think of it as "Capitalism."

The keys to decrypt a DVD are controlled by an org called DVD-CCA, and they have a bunch of licensing requirements for anyone who gets a key from them. Among these is something called region-coding: if you buy a DVD in France, it'll have a flag set that says, "I am a French DVD." Bring that DVD to America and your DVD player will compare the flag to its list of permitted regions, and if they don't match, it will tell you that it's not allowed to play your disc.

Remember: there is no copyright that says that an author gets to do this. When we wrote the copyright statutes and granted authors the right to control display, performance, duplication, derivative works, and so forth, we didn't leave out "geography" by accident. That was on-purpose.

So when your French DVD won't play in America, that's not because it'd be illegal to do so: it's because the studios have invented a business-model and then invented a copyright law to prop it up. The DVD is your property and so is the DVD player, but if you break the region-coding on your disc, you're going to run afoul of anticircumvention.

That's what happened to [Jon Johansen](#), a Norwegian teenager who wanted to watch French DVDs on his Norwegian DVD player. He and some pals wrote some code to break the CSS so that he could do so. He's a wanted man here in America; in Norway the studios put the local fuzz up to bringing him up on charges of *unlawfully trespassing upon a computer system*. When his defense asked, "Which computer has Jon trespassed upon?" the answer was: "His own."

His no-fooling, real and physical property has been expropriated by the weird, notional, metaphorical intellectual property on his DVD: DRM only works if your record player becomes the property of whomever's records you're playing.

3. DRM systems are bad for biz

This is the worst of all the ideas embodied by DRM: that people who make record-players should be able to spec whose records you can listen to, and that people who make records should have a veto over the design of record-players.

We've never had this principle: in fact, we've always had just the reverse. Think about all the things that can be plugged into a parallel or serial interface, which were never envisioned by their inventors. Our strong economy and rapid innovation are byproducts of the ability of anyone to make anything that plugs into anything else: from the Flo-bee electric razor that snaps onto the end of your vacuum-hose to the octopus spilling out of your car's dashboard lighter socket, standard interfaces that anyone can build for are what makes billionaires out of nerds.

The courts affirm this again and again. It used to be illegal to plug anything that didn't come from AT&T into your phone-jack. They claimed that this was for the safety of the network, but really it was about propping up this little penny-ante racket that AT&T had in charging you a rental fee for your phone until you'd paid for it a thousand times over.

When that ban was struck down, it created the market for third-party phone equipment, from talking novelty phones to answering machines to cordless handsets to headsets — billions of dollars of economic activity that had been suppressed by the closed interface. Note that AT&T was one of the big beneficiaries of this: they *also* got into the business of making phone-kit.

DRM is the software equivalent of these closed hardware interfaces. [Robert Scoble](#) is a Softie who has an excellent blog, where he wrote an essay about the best way to protect your investment in the digital music you buy. Should you buy Apple iTunes music, or Microsoft DRM music? Scoble argued that Microsoft's music was a sounder investment, because Microsoft would have more downstream licensees for its proprietary format and therefore you'd have a richer ecosystem of devices to choose from when you were shopping for gizmos to play your virtual records on.

What a weird idea: that we should evaluate our record-purchases on the basis of which recording company will allow the greatest diversity of record-players to play its discs! That's like telling someone to buy the Betamax instead of the Edison Kinetoscope because Thomas Edison is a crank about licensing his patents; all the while ignoring the world's relentless march to the more open VHS format.

It's a bad business. DVD is a format where the guy who makes the records gets to design the record players. Ask yourself: how much innovation has there been over the past decade of DVD players? They've gotten cheaper and smaller, but where are the weird and amazing new markets for DVD that were opened up by the VCR? There's a company that's manufacturing the world's first HDD-based DVD jukebox, a thing that holds 30 movies, and they're charging \$30,000 for this thing. We're talking about a \$300 hard drive and a \$300 PC—all that other cost is the cost of anticompetition.

4. DRM systems are bad for artists

But what of the artist? The hardworking filmmaker, the ink-stained scribbler, the heroin-cured leathery rock-star? We poor slobs of the creative class are everyone's favorite poster-children here: the RIAA and MPAA hold us up and say, "Won't someone please think of the children?" File-sharers say, "Yeah, we're thinking about the artists, but the labels are The Man, who cares what happens to you?"

To understand what DRM does to artists, you need to understand how copyright and technology interact. Copyright is inherently technological, since the things it addresses—copying, transmitting, and so on—are inherently technological.

The piano roll was the first system for cheaply copying music. It was invented at a time when the dominant form of entertainment in America was getting a talented pianist to come into your living room and pound out some tunes while you sang along. The music industry consisted mostly of sheet-music publishers.

The player piano was a digital recording and playback system. Piano-roll companies bought sheet music and ripped the notes printed on it into 0s and 1s on a long roll of computer tape, which they sold by the thousands—the hundreds of thousands—the millions. They did this without a penny's compensation to the publishers. They were digital music pirates. Arrrrr!

Predictably, the composers and music publishers went nutso. Sousa showed up in Congress to say that:

These talking machines are going to ruin the artistic development of music in this country. When I was a boy ... in front of every house in the summer evenings, you would find young people together singing the songs of the day or old songs. Today you hear these infernal machines going night and day. We will not have a vocal chord left. The vocal chord will be eliminated by a process of evolution, as was the tail of man when he came from the ape.

The publishers asked Congress to ban the piano roll and to create a law that said that any new system for reproducing music should be subject to a veto from their industry association. Lucky for us, Congress realized what side of their bread had butter on it and decided not to criminalize the dominant form of entertainment in America.

But there was the problem of paying artists. The Constitution sets out the purpose of American copyright: to promote the useful arts and sciences. The composers had a credible story that they'd do less composing if they weren't paid for it, so Congress needed a fix. Here's what they came up with: anyone who paid a music publisher two cents would have the right to make one piano roll of any song that publisher published. The publisher couldn't say no, and no one had to hire a lawyer at \$200 an hour to argue about whether the payment should be two cents or a nickel.

This compulsory license is still in place today: when Joe Cocker sings "With a Little Help from My Friends," he pays a fixed fee to the Beatles' publisher and away he goes—even if Ringo hates the idea. If you ever wondered how Sid Vicious talked Anka into letting him get a crack at "My Way," well, now you know.

That compulsory license created a world where a thousand times more money was made by a thousand times more creators who made a thousand times more music that reached a thousand times more people.

This story repeats itself throughout the technological century, every ten or fifteen years. Radio was enabled by a voluntary blanket license—the music companies got together and asked for an antitrust exemption so that they could offer all their music for a flat fee. Cable TV took a compulsory: the only way cable operators could get their hands on broadcasts was to pirate them and shove them down the wire, and Congress saw fit to legalize this practice rather than screw around with their constituents' TVs.

Sometimes, the courts and Congress decided to simply take away a copyright — that’s what happened with the VCR. When Sony brought out the VCR in 1976, the studios had already decided what the experience of watching a movie in your living room would look like: they’d licensed out their programming for use on a machine called a Discovision, which played big LP-sized discs that disintegrated after a few plays. Proto-DRM.

The copyright scholars of the day didn’t give the VCR very good odds. Sony argued that their box allowed for a fair use, which is defined as a use that a court rules is a defense against infringement based on four factors: whether the use transforms the work into something new, like a collage; whether it uses all or some of the work; whether the work is artistic or mainly factual; and whether the use undercuts the creator’s business-model.

The Betamax failed on all four fronts: when you time-shifted or duplicated a Hollywood movie off the air, you made a non-transformative use of 100 percent of a creative work in a way that directly undercut the Discovision licensing stream.

Jack Valenti, the mouthpiece for the motion-picture industry, told Congress in 1982 that the VCR was to the American film industry “as the Boston Strangler is to a woman home alone.”

But the Supreme Court ruled against Hollywood in 1984, when it determined that any device capable of a substantial non-infringing use was legal. In other words, “We don’t buy this Boston Strangler business: if your business model can’t survive the emergence of this general-purpose tool, it’s time to get another business-model or go broke.”

Hollywood found another business model, as the broadcasters had, as the Vaudeville artists had, as the music publishers had, and they made more art that paid more artists and reached a wider audience.

There’s one thing that every new art business-model had in common: it embraced the medium it lived in.

This is the overweening characteristic of every single successful new medium: it is true to itself. The Luther Bible didn’t succeed on the axes that made a hand-copied monk Bible valuable: they were ugly, they weren’t in Church Latin, they weren’t read aloud by someone who could interpret it for his lay audience, they didn’t represent years of devoted-with-a-capital-D labor by someone who had given his life over to God. The thing that made the Luther Bible a success was its scalability: it was more popular because it was more proliferate: all success factors for a new medium pale beside its profligacy. The most successful organisms on earth are those that reproduce the most: bugs and bacteria, nematodes and virii. Reproduction is the best of all survival strategies.

Piano rolls didn’t sound as good as the music of a skilled pianist: but they *scaled better*. Radio lacked the social elements of live performance, but more people could build a crystal set and get it aimed correctly than could pack into even the largest Vaudeville house. MP3s don’t come with liner notes, they aren’t sold to you by a hipper-than-thou record store clerk who can help you make your choice, bad rips and truncated files abound: I once downloaded a twelve-second copy of “Hey Jude” from the original Napster. Yet MP3 is outcompeting the CD. I don’t know what to do with CDs anymore: I get them, and they’re like the especially nice garment bag they give you at the fancy suit shop: it’s nice and you feel like a goof for throwing it out, but Christ, how many of these things can you usefully own? I can put ten thousand songs on my laptop, but a comparable pile of discs, with liner notes and so forth—that’s a liability: it’s a piece of my monthly storage-locker costs.

Here are the two most important things to know about computers and the Internet:

1. A computer is a machine for rearranging bits
2. The Internet is a machine for moving bits from one place to another very cheaply and quickly

Any new medium that takes hold on the Internet and with computers will embrace these two facts, not regret them. A newspaper press is a machine for spitting out cheap and smeary newsprint at speed: if you try to make it output fine art lithos, you'll get junk. If you try to make it output newspapers, you'll get the basis for a free society.

And so it is with the Internet. At the heyday of Napster, record execs used to show up at conferences and tell everyone that Napster was doomed because no one wanted lossily compressed MP3s with no liner notes and truncated files and misspelled metadata.

Today we hear ebook publishers tell each other and anyone who'll listen that the barrier to ebooks is screen resolution. It's bollocks, and so is the whole sermonette about how nice a book looks on your bookcase and how nice it smells and how easy it is to slip into the tub. These are obvious and untrue things, like the idea that radio will catch on once they figure out how to sell you hotdogs during the intermission, or that movies will really hit their stride when we can figure out how to bring the actors out for an encore when the film's run out. Or that what the Protestant Reformation really needs is Luther Bibles with facsimile illumination in the margin and a rent-a-priest to read aloud from your personal Word of God.

New media don't succeed because they're like the old media, only better: they succeed because they're worse than the old media at the stuff the old media is good at, and better at the stuff the old media are bad at. Books are good at being paperwhite, high-resolution, low-infrastructure, cheap and disposable. Ebooks are good at being everywhere in the world at the same time for free in a form that is so malleable that you can just pastebomb it into your IM session or turn it into a page-a-day mailing list.

The only really successful epublishing—I mean, hundreds of thousands, millions of copies distributed and read—is the bookwarez scene, where scanned-and-OCR'd books are distributed on the darknet. The only legit publishers with any success at epublishing are the ones whose books cross the Internet without technological fetter: publishers like Baen Books and my own, Tor, who are making some or all of their catalogs available in ASCII and HTML and PDF.

The hardware-dependent ebooks, the DRM use-and-copy-restricted ebooks, they're cratering. Sales measured in the tens, sometimes the hundreds. Science fiction is a niche business, but when you're selling copies by the ten, that's not even a business, it's a hobby.

Every one of you has been riding a curve where you read more and more words off of more and more screens every day through most of your professional careers. It's zero-sum: you've also been reading fewer words off of fewer pages as time went by: the dinosauric executive who prints his email and dictates a reply to his secretary is info-roadkill.

Today, at this very second, people read words off of screens for every hour that they can find. Your kids stare at their Game Boys until their eyes fall out. Euroteens ring doorbells with their hypertrophied, SMS-twitching thumbs instead of their index fingers.

Paper books are the packaging that books come in. Cheap printer-binderies like the Internet Bookmobile that can produce a full bleed, four color, glossy cover, printed spine, perfect-bound book in ten minutes for a dollar are the future of paper books: when you need an instance of a paper book, you generate one, or part of one, and pitch it out when you're done. I landed at SEA-TAC on Monday and burned a couple CDs from my music collection to listen to in the rental car. When I drop the car off, I'll leave them behind. Who needs 'em?

Whenever a new technology has disrupted copyright, we've changed copyright. Copyright isn't an ethical proposition, it's a utilitarian one. There's nothing *moral* about paying a composer tuppence for the piano-roll rights, there's nothing *immoral* about not paying Hollywood for the right to videotape a movie off your TV. They're just the best way of balancing out so that people's physical property rights in their VCRs and

phonographs are respected and so that creators get enough of a dangling carrot to go on making shows and music and books and paintings.

Technology that disrupts copyright does so because it simplifies and cheapens creation, reproduction and distribution. The existing copyright businesses exploit inefficiencies in the old production, reproduction and distribution system, and they'll be weakened by the new technology. But new technology always gives us more art with a wider reach: that's what tech is *for*.

Tech gives us bigger pies that more artists can get a bite out of. That's been tacitly acknowledged at every stage of the copyfight since the piano roll. When copyright and technology collide, it's copyright that changes.

Which means that today's copyright—the thing that DRM nominally props up—didn't come down off the mountain on two stone tablets. It was created in living memory to accommodate the technical reality created by the inventors of the previous generation. To abandon invention now robs tomorrow's artists of the new businesses and new reach and new audiences that the Internet and the PC can give them.

5. DRM is a bad business-move for MSFT

When Sony brought out the VCR, it made a record player that could play Hollywood's records, even if Hollywood didn't like the idea. The industries that grew up on the back of the VCR—movie rentals, home taping, camcorders, even Bar Mitzvah videographers—made billions for Sony and its cohort.

That was good business—even if Sony lost the Betamax-VHS format wars, the money on the world-with-VCRs table was enough to make up for it.

But then Sony acquired a relatively tiny entertainment company and it started to massively screw up. When MP3 rolled around and Sony's walkman customers were clamoring for a solid-state MP3 player, Sony let its music business-unit run its show: instead of making a high-capacity MP3 walkman, Sony shipped its Music Clips, low-capacity devices that played brain-damaged DRM formats like Real and OpenAG. They spent good money engineering "features" into these devices that kept their customers from freely moving their music back and forth between their devices. Customers stayed away in droves.

Today, Sony is dead in the water when it comes to walkmen. The market leaders are poky Singaporean outfits like Creative Labs—the kind of company that Sony used to crush like a bug, back before it got borged by its entertainment unit—and PC companies like Apple.

That's because Sony shipped a product that there was no market demand for. No Sony customer woke up one morning and said, "Damn, I wish Sony would devote some expensive engineering effort in order that I may do less with my music." Presented with an alternative, Sony's customers enthusiastically jumped ship.

The same thing happened to a lot of people I know who used to rip their CDs to WMA. You guys sold them software that produced smaller, better-sounding rips than the MP3 rippers, but you also fixed it so that the songs you ripped were device-locked to their PCs. What that meant is that when they backed up their music to another hard-drive and reinstalled their OS (something that the spyware and malware wars has made more common than ever), they discovered that after they restored their music that they could no longer play it. The player saw the new OS as a different machine, and locked them out of their own music.

There is no market demand for this "feature." None of your customers want you to make expensive modifications to your products that make backing up and restoring even harder. And there is no moment when your customers will be less forgiving than the moment that they are recovering from catastrophic technology failures.

[I speak from experience](#). Because I buy a new Powerbook every ten months, and because I always order the new models the day they're announced, I get a lot of lemons from Apple. That means that I hit Apple's three-iTunes-authorized-computers limit pretty early on and found myself unable to play the hundreds of dollars' worth of iTunes songs I'd bought because one of my authorized machines was a lemon that Apple had broken up for parts, one was in the shop getting fixed by Apple, and one was my mom's computer, 3,000 miles away in Toronto.

If I had been a less good customer for Apple's hardware, I would have been fine. If I had been a less enthusiastic evangelist for Apple's products—if I hadn't shown my mom how iTunes Music Store worked—I would have been fine. If I hadn't bought so much iTunes music that burning it to CD and re-ripping it and re-keying all my metadata was too daunting a task to consider, I would have been fine.

As it was Apple rewarded my trust, evangelism and out-of-control spending by treating me like a crook and locking me out of my own music, at a time when my Powerbook was in the shop—i.e., at a time when I was hardly disposed to feel charitable to Apple.

I'm an edge case here, but I'm a *leading edge* case. If Apple succeeds in its business plans, it will only be a matter of time until even average customers have upgraded enough hardware and bought enough music to end up where I am.

You know what I would totally buy? A record player that let me play everybody's records. Right now, the closest I can come to that is an open source app called VLC, but it's clunky and buggy and it didn't come pre-installed on my computer.

Sony didn't make a Betamax that only played the movies that Hollywood was willing to permit—Hollywood asked them to do it, they proposed an early, analog broadcast flag that VCRs could hunt for and respond to by disabling recording. Sony ignored them and made the product they thought their customers wanted.

I'm a Microsoft customer. Like millions of other Microsoft customers, I want a player that plays anything I throw at it, and I think that you are just the company to give it to me.

Yes, this would violate copyright law as it stands, but Microsoft has been making tools of piracy that change copyright law for decades now. Outlook, Exchange and MSN are tools that abet widescale digital infringement.

More significantly, IIS and your caching proxies all make and serve copies of documents without their authors' consent, something that, if it is legal today, is only legal because companies like Microsoft went ahead and did it and dared lawmakers to prosecute.

Microsoft stood up for its customers and for progress, and won so decisively that most people never even realized that there was a fight.

Do it again! This is a company that looks the world's roughest, toughest anti-trust regulators in the eye and laughs. Compared to anti-trust people, copyright lawmakers are pantywaists. You can take them with your arm behind your back.

In Siva Vaidhyanathan's book *The Anarchist in the Library*, he talks about why the studios are so blind to their customers' desires. It's because people like you and me spent the 80s and the 90s telling them bad science fiction stories about impossible DRM technology that would let them charge a small sum of money every time someone looked at a movie — want to fast-forward? That feature costs another penny. Pausing is two cents an hour. The mute button will cost you a quarter.

When Mako Analysis issued their report last month advising phone companies to stop supporting Symbian phones, they were just writing the latest installment in this story. Mako says that phones like my P900, which can play MP3s as ringtones, are bad for the cellphone economy, because it'll put the extortionate ringtone sellers out of business. What Mako is saying is that just because you bought the CD doesn't mean that you should expect to have the ability to listen to it on your MP3 player, and just because it plays on your MP3 player is no reason to expect it to run as a ringtone. I wonder how they feel about alarm clocks that will play a CD to wake you up in the morning? Is that strangling the nascent "alarm tone" market?

The phone companies' customers want Symbian phones and for now, at least, the phone companies understand that if they don't sell them, someone else will.

The market opportunity for a truly capable devices is enormous. There's a company out there charging \$30,000 for a \$600 DVD jukebox — go and eat their lunch! Steve Jobs isn't going to do it: he's off at the D conference telling studio execs not to release hi-def movies until they're sure no one will make a hi-def DVD burner that works with a PC.

Maybe they won't buy into his BS, but they're also not much interested in what you have to sell. At the Broadcast Protection Discussion Group meetings where the Broadcast Flag was hammered out, the studios' position was, "We'll take anyone's DRM except Microsoft's and Philips'." When I met with UK broadcast wonks about the European version of the Broadcast Flag underway at the Digital Video Broadcasters' forum, they told me, "Well, it's different in Europe: mostly they're worried that some American company like Microsoft will get their claws into European television."

American film studios didn't want the Japanese electronics companies to get a piece of the movie pie, so they fought the VCR. Today, everyone who makes movies agrees that they don't want to let you guys get between them and their customers.

Sony didn't get permission. Neither should you. Go build the record player that can play everyone's records.

Because if you don't do it, someone else will.

This text is dedicated to the public domain, using a Creative Commons public domain dedication:

Copyright-Only Dedication (based on United States law).

The person or persons who have associated their work with this document (the "Dedicator") hereby dedicate the entire copyright in the work of authorship identified below (the "Work") to the public domain.

Dedicator makes this dedication for the benefit of the public at large and to the detriment of Dedicator's heirs and successors. Dedicator intends this dedication to be an overt act of relinquishment in perpetuity of all present and future rights under copyright law, whether vested or contingent, in the Work. Dedicator understands that such relinquishment of all rights includes the relinquishment of all rights to enforce (by lawsuit or otherwise) those copyrights in the Work.

Dedicator recognizes that, once placed in the public domain, the Work may be freely reproduced, distributed, transmitted, used, modified, built upon, or otherwise exploited by anyone for any purpose, commercial or non-commercial, and in any way, including by methods that have not yet been invented or conceived.

Bigger Monster, Weaker Chains:

The Growth of an American Surveillance Society



By Jay Stanley and Barry Steinhardt
January 2003



ACLU

AMERICAN CIVIL LIBERTIES UNION

Technology and
Liberty Program

Preface

There is no shortage of stories in the media today about the continuing assault on our privacy. But while the latest surveillance program or privacy-invading gadget always receives ample coverage, it is much rarer to find stories that connect the dots and describe the overall impact on privacy in the United States. And without that big picture, the importance of the individual pieces often gets lost.

This new report from the American Civil Liberties Union seeks to provide greater understanding of how our activities are increasingly being tracked and recorded, and how all that data could be drawn together from different sources to create a single high-resolution image of our private lives.

For decades, the notion of a “surveillance society,” where every facet of our private lives is monitored and recorded, has sounded abstract, paranoid or far-fetched to many people.

No more! The public’s recent introduction to the Pentagon’s “Total Information Awareness” project, which seeks to tie together every facet of our private lives in one big surveillance scheme, has provided a stunning lesson in the realities of the new world in which we live. The revelations about the Total Information Awareness program have given the public a sudden introduction to the concept of “data surveillance,” and an early glimmer of the technological potential for a surveillance society. It has also confirmed the national security and law enforcement establishments’ hunger for such surveillance.

Yet too many people still do not understand the danger, do not grasp just how radical an increase in surveillance by both the government and the private sector is becoming possible, or do not see that the danger stems not just from a single government program, but from a number of parallel developments in the worlds of technology, law, and politics. In this report, the ACLU seeks to flesh out these trends, and, by setting down various developments together in one place, to illuminate the overall danger and what can be done to eliminate it.

The surveillance monster is getting bigger and stronger by the day. But the American Civil Liberties Union believes that it is not too late to build a system of law that can chain it. It is not too late to take back our data.

Introduction

Privacy and liberty in the United States are at risk. A combination of lightning-fast technological innovation and the erosion of privacy protections threatens to transform Big Brother from an oft-cited but remote threat into a very real part of American life. We are at risk of turning into a Surveillance Society.

The explosion of computers, cameras, sensors, wireless communication, GPS, biometrics, and other technologies in just the last 10 years is feeding a surveillance monster that is growing silently in our midst. Scarcely a month goes by in which we don't read about some new high-tech way to invade people's privacy, from face recognition to implantable microchips, data-mining, DNA chips, and even "brain wave fingerprinting." The fact is, there are no longer any *technical* barriers to the Big Brother regime portrayed by George Orwell.

Even as this surveillance monster grows in power, we are weakening the legal chains that keep it from trampling our lives. We should be responding to intrusive new technologies by building stronger restraints to protect our privacy; instead, we are doing the opposite – loosening regulations on government surveillance, watching passively as private surveillance grows unchecked, and contemplating the introduction of tremendously powerful new surveillance infrastructures that will tie all this information together.

A gradual weakening of our privacy rights has been underway for decades, but many of the most startling developments have come in response to the terrorist attacks of September 11. But few of these hastily enacted measures are likely to increase our protection against terrorism. More often than not, September 11 has been used as a pretext to loosen constraints that law enforcement has been chafing under for years.

It doesn't require some apocalyptic vision of American democracy being replaced by dictatorship to worry about a surveillance society. There is a lot of room for the United States to become a meaner, less open and less just place without any radical change in government. All that's required is the continued construction of new surveillance technologies and the simultaneous erosion of privacy protections.

It's not hard to imagine how in the near future we might see scenarios like the following:

- An African-American man from the central city visits an affluent white suburb to attend a co-worker's barbeque. Later that night, a crime takes place elsewhere in the neighborhood. The police review surveillance camera images, use face recognition to identify the man, and pay him a visit at home the next day. His trip to the suburbs where he "didn't belong" has earned him an interrogation from suspicious police.
- A tourist walking through an unfamiliar city happens upon a sex shop. She stops to gaze at several curious items in the store's window before moving along. Unbeknownst to her, the store has set up the newly available "Customer Identification System," which detects a signal being emitted by a computer chip in her driver's license and records her identity and the date, time, and duration of her brief look inside the window. A week later, she gets a solicitation in the mail mentioning her "visit" and embarrassing her in front of her family.

Such possibilities are only the tip of the iceberg. The media faithfully reports the latest surveillance gadgets and the latest moves to soften the rules on government spying, but rarely provides the big picture. That is unfortunate, because each new threat to our privacy is much more significant as part of the overall trend than it seems when viewed in isolation. When these monitoring technologies and techniques are combined, they can create a surveillance network far more powerful than any single one would create on its own.

The good news is that these trends can be stopped. As the American people realize that each new development is part of this larger story, they will give more and more weight to protecting privacy, and support the measures we need to preserve our freedom.

The Growing Surveillance Monster

In the film *Minority Report*, which takes place in the United States in the year 2050, people called “Pre-cogs” can supposedly predict future crimes, and the nation has become a perfect surveillance society. The frightening thing is that except for the psychic Pre-cogs, the technologies of surveillance portrayed in the film already exist or are in the pipeline. Replace the Pre-cogs with “brain fingerprinting” – the supposed ability to ferret out dangerous tendencies by reading brain waves – and the film’s entire vision no longer lies far in the future. Other new privacy invasions are coming at us from all directions, from video and data surveillance to DNA scanning to new data-gathering gadgets.

Video Surveillance

Surveillance video cameras are rapidly spreading throughout the public arena. A survey of surveillance cameras in Manhattan, for example, found that it is impossible to walk around the city without being

Video surveillance may be on the verge of a revolutionary expansion in American Life.

recorded nearly every step of the way. And since September 11 the pace has quickened, with new cameras being placed not only in some of our most sacred public spaces, such as the National Mall in Washington and the Statue of Liberty in New York harbor, but on ordinary public streets all over America.

As common as video cameras have become, there are strong signs that, without public action, video surveillance may be on the verge of a revolutionary expansion in American life. There are three factors propelling this revolution:

1. **Improved technology.** Advances such as the digitization of video mean cheaper cameras, cheaper transmission of far-flung video feeds, and cheaper storage and retrieval of images.
2. **Centralized surveillance.** A new centralized surveillance center in Washington, DC is an early indicator of what technology may bring. It allows officers to view images from video cameras across the city – public buildings and streets, neighborhoods, Metro stations, and even schools. With the flip of a switch, officers can zoom in on people from cameras a half-mile away.¹

- 3. Unexamined assumptions that cameras provide security.** In the wake of the September 11 attacks, many embraced surveillance as the way to prevent future attacks and prevent crime. But it is far from clear how cameras will increase security. U.S. government experts on security technology, noting that “monitoring video screens is both boring and mesmerizing,” have found in experiments that after only 20 minutes of watching video monitors, “the attention of most individuals has degenerated to well below acceptable levels.”² In addition, studies of cameras’ effect on crime in Britain, where they have been extensively deployed, have found no conclusive evidence that they have reduced crime.³

These developments are creating powerful momentum toward pervasive video surveillance of our public spaces. If centralized video facilities are permitted in Washington and around the nation, it is inevitable that they will be expanded – not only in the number of cameras but also in their power and ability. It is easy to foresee inexpensive, one-dollar cameras being distributed throughout our cities and tied via wireless technology into a centralized police facility where the life of the city can be monitored. Those video signals could be stored indefinitely in digital form in giant but inexpensive databases, and called up with the click of a mouse at any time. With face recognition, the video records could even be indexed and searched based on who the systems identify – correctly, or all too often, incorrectly.

Several airports around the nation, a handful of cities, and even the National Park Service at the Statue of Liberty have installed face recognition. While not nearly reliable enough to be effective as a security application⁴, such a system could still violate the privacy of a significant percentage of the citizens who appeared before it (as well as the privacy of those who do not appear before it but are falsely identified as having done so). Unlike, say, an iris scan, face recognition doesn’t require the knowledge, consent, or participation of the subject; modern cameras can easily view faces from over 100 yards away.

Further possibilities for the expansion of video surveillance lie with unmanned aircraft, or drones, which have been used by the military and the CIA overseas for reconnaissance, surveillance, and targeting. Controlled from the ground, they can stay airborne for days at a time. Now there is talk of deploying them domestically. Senate Armed Services Committee Chairman John Warner (R, VA) said in December 2002 that he wants to explore their use in Homeland Security, and a number of domestic government agencies have expressed interest in deploying them. Drones are likely to be just one of many ways in which improving robotics technology will be applied to surveillance.⁵

The bottom line is that surveillance systems, once installed, rarely remain confined to their original purpose. Once the nation decides to go down the path of seeking security through video surveillance, the imperative to make it work will become overwhelming, and the monitoring of citizens in public places will quickly become pervasive.

Data Surveillance

An insidious new type of surveillance is becoming possible that is just as intrusive as video surveillance – what we might call

“data surveillance.” Data surveillance is *the collection of information about an identifiable individual, often from multiple sources, that can be assembled into a portrait of that person’s activities.*⁶ Most com-

It will soon be possible to recreate an individual’s activities with such detail that it becomes no different from being followed around with a video camera.

puters are programmed to automatically store and track usage data, and the spread of computer chips in our daily lives means that more and more of our activities leave behind “data trails.” It will soon be possible to combine information from different sources to recreate an individual’s activities with such detail that it becomes no different from being followed around all day by a detective with a video camera.

Some think comprehensive public tracking will make no difference, since life in public places is not “private” in the same way as life inside the home. This is wrong; such tracking would represent a radical change in American life. A woman who leaves her house, drives to a store, meets a friend for coffee, visits a museum, and then returns home may be in public all day, but her life is still private in that she is the only one who has an overall view of how she spent her day. In America, she does not expect that her activities are being watched or tracked in any systematic way – she expects to be left alone. But if current trends continue, it will be impossible to have any contact with the outside world that is not watched and recorded.

The Commodification of Information

A major factor driving the trend toward data surveillance forward is the commodification of personal information by corporations. As computer technology exploded in recent decades, making it much easier to collect information about what Americans buy and do, companies came to realize that such data is often very valuable. The expense of marketing efforts gives businesses a strong incentive to know as much about consumers as possible so they can focus on the most likely new customers. Surveys, sweepstakes questionnaires, loyalty programs and detailed product registration forms have proliferated in American life – all aimed at gathering information about consumers. Today, any consumer activity that is *not* being tracked and recorded is increasingly being viewed by businesses as money left on the table.

On the Internet, where every mouse click can be recorded, the tracking and profiling of consumers is even more prevalent. Web sites can not only track what consumers buy, but what they *look at* – and for how long, and in what order. With the end of the Dot Com era, personal information has become an even more precious source of hard cash for those Internet ventures that survive. And of course Americans use the Internet not just as a shopping mall, but to research topics of interest, debate political issues, seek support for personal problems, and many other purposes that can generate deeply private information about their thoughts, interests, lifestyles, habits, and activities.

Unlike other medical information, DNA is a unique combination: both difficult to keep confidential and extremely revealing about us.

Genetic Privacy

The relentless commercialization of information has also led to the breakdown of some longstanding traditions, such as doctor-patient confidentiality. Citizens share some of their most intimate and embarrassing secrets with their doctors on the old-fashioned assumption that their conversations are confidential. Yet those details are routinely shared with insurance companies, researchers, marketers, and employers. An insurance trade organization called the Medical Information Bureau even keeps a centralized medical database with records on millions of patients. Weak new medical privacy rules will do little to stop this behavior.

An even greater threat to medical privacy is looming: genetic information. The increase in DNA analysis for medical testing, research, and other purposes will accelerate sharply in coming years, and will increasingly be incorporated into routine health care.

Unlike other medical information, genetic data is a unique combination: both difficult to keep confidential and extremely revealing about us. DNA is very easy to acquire because we constantly slough off hair, saliva, skin cells and other samples of our DNA (household dust, for example, is made up primarily of dead human skin cells). That means that no matter how hard we strive to keep our genetic code private, we are always vulnerable to other parties' secretly testing samples of our DNA. The issue will be intensified by the development of cheap and efficient DNA chips capable of reading parts of our genetic sequences.

Gramm-Leach effectively gives financial institutions permission to sell their customers' financial data to anyone they choose.

Already, it is possible to send away a DNA sample for analysis. A testing company called Genelex reports that it has amassed 50,000 DNA samples, many gathered surreptitiously for paternity testing. "You'd be amazed," the company's CEO told U.S. News & World Report. "Siblings have sent in mom's discarded Kleenex and wax from her hearing aid to resolve the family rumors."

Not only is DNA easier to acquire than other medical information, revealing it can also have more profound consequences. Genetic markers are rapidly being identified for all sorts of genetic diseases, risk factors, and other characteristics. None of us knows what time bombs are lurking in our genomes.

The consequences of increased genetic transparency will likely include:

- **Discrimination by insurers.** Health and life insurance companies could collect DNA for use in deciding who to insure and what to charge them, with the result that a certain proportion of the population could become uninsurable. The insurance industry has already vigorously opposed efforts in Congress to pass meaningful genetic privacy and discrimination bills.
- **Employment discrimination.** Genetic workplace testing is already on the rise, and the courts have heard many cases. Employers desiring healthy, capable workers will always have an incentive to discriminate based on DNA – an incentive that will be even stronger as long as health insurance is provided through the workplace.
- **Genetic spying.** Cheap technology could allow everyone from schoolchildren to dating couples to nosy neighbors to routinely check out each other's genetic codes. A likely high-profile example: online posting of the genetic profiles of celebrities or politicians.

Financial privacy

Like doctor-patient confidentiality, the tradition of privacy and discretion by financial institutions has also collapsed; financial companies today routinely put the details of their customers' financial lives up for sale.

A big part of the problem is the Gramm-Leach-Bliley Act passed by Congress in 1999. Although Gramm-Leach is sometimes described as a “financial privacy law,” it created a very weak privacy standard – so weak, in fact, that far from protecting Americans’ financial privacy, the law has had the effect of ratifying the increasing abandonment of customer privacy by financial companies.

Gramm-Leach effectively gives financial institutions permission to sell their customers’ financial data to anyone they choose. That includes the date, amount, and recipient of credit card charges or checks a customer has written; account balances; and information about the flow of deposits and withdrawals through an account. Consumers provide a tremendous amount of information about themselves when they fill out applications to get a loan, buy insurance, or purchase securities, and companies can also share that information. In fact, the only information a financial company may NOT give out about you is your account number.

Under Gramm-Leach, you get no privacy unless you file complex paperwork, following a financial institution’s precise instructions before a deadline they set, and repeating the process for each and every financial service provider who may have data about you. And it is a process that many companies intentionally make difficult and cumbersome; few let consumers “opt out” of data sharing through a Web site or phone number, or even provide a self-addressed envelope.

Gramm-Leach is an excellent example of the ways that privacy protections are being weakened even as the potential for privacy invasion grows.

New Data-Gathering Technologies

The discovery by businesses of the monetary value of personal information and the vast new project of tracking the habits of consumers has been made possible by advances in computers, databases and the Internet. In the near future, other new technologies will continue to fill out the mosaic of information it is possible to collect on every individual. Examples include:

- **Cell phone location data.** The government has mandated that manufacturers make cell phones capable of automatically reporting their location when an owner dials 911. Of course, those phones are capable of tracking their location at other times as well. And in applying the rules that protect the privacy of telephone records to this location data, the government is weakening those rules in a way that allows phone companies to collect and share data about the location and movements of their customers.
- **Biometrics.** Technologies that identify us by unique bodily attributes such as our fingerprints, faces, iris patterns, or DNA are already being proposed for inclusion on national ID cards and to identify airline passengers. Face recognition is spreading. Fingerprint scanners have been introduced as security or payment mechanisms in office buildings, college campuses, grocery stores and even fast-food restaurants. And several companies are working on DNA chips that will be able to instantly identify individuals by the DNA we leave behind everywhere we go.
- **Black boxes.** All cars built today contain computers, and some of those computers are being programmed in ways that are not necessarily in the interest of owners. An increasing number of cars contain devices akin to the “black boxes” on aircraft that record details about a vehicle’s operation and movement. Those devices can “tattle” on car owners to the police or insurance

investigators. Already, one car rental agency tried to charge a customer for speeding after a GPS device in the car reported the transgression back to the company. And cars are just one example of how products and possessions can be programmed to spy and inform on their owners.

- **RFID chips.** RFID chips, which are already used in such applications as toll-booth speed passes, emit a short-range radio signal containing a unique code that identifies each chip. Once the cost of these chips falls to a few pennies each, plans are underway to affix them to products in stores, down to every can of soup and tube of toothpaste. They will allow everyday objects to “talk” to each other – or to anyone else who is listening. For example, they could let market researchers scan the contents of your purse or car from five feet away, or let police officers scan your identification when they pass you on the street.
- **Implantable GPS chips.** Computer chips that can record and broadcast their location have also been developed. In addition to practical uses such as building them into shipping containers, they can also serve as location “bugs” when, for example, hidden by a suspicious husband in a wife’s purse. And they can be implanted under the skin (as can RFID chips).

RFID chips will allow everyday objects to “talk” to each other – or anyone else who is listening.

If we do not act to reverse the current trend, data surveillance – like video surveillance – will allow corporations or the government to constantly monitor what individual Americans do every day. Data surveillance would cover *everyone*, with records of every transaction and activity squirreled away until they are sucked up by powerful search engines, whether as part of routine security checks, a general sweep for suspects in an unsolved crime, or a program of harassment against some future Martin Luther King.

Government Surveillance

Data surveillance is made possible by the growing ocean of privately collected personal data. But who would conduct that surveillance? There are certainly business incentives for doing so; companies called data aggregators (such as Acxiom and ChoicePoint) are in the business of compiling detailed databases on individuals and then selling that information to others. Although these companies are invisible to the average person, data aggregation is an enormous, multi-billion-dollar industry. Some databases are even “co-ops” where participants agree to contribute data about their customers in return for the ability to pull out cross-merchant profiles of customers’ activities.

The biggest threat to privacy, however, comes from the government. Many Americans are naturally concerned about corporate surveillance, but only the government has the power to take away liberty – as has been demonstrated starkly by the post-September 11 detention of suspects without trial as “enemy combatants.”

In addition, the government has unmatched power to centralize all the private sector data that is being generated. In fact, the distinction between government and private-sector privacy invasions is fading

quickly. The Justice Department, for example, reportedly has an \$8 million contract with data aggregator ChoicePoint that allows government agents to tap into the company's vast database of personal information on individuals.⁸ Although the Privacy Act of 1974 banned the government from maintaining information on citizens who are not the targets of investigations, the FBI can now evade that requirement by simply purchasing information that has been collected by the private sector. Other proposals – such as the Pentagon's "Total Information Awareness" project and airline passenger profiling programs – would institutionalize government access to consumer data in even more far-reaching ways (see below).

Government Databases

The government's access to personal information begins with the thousands of databases it maintains on the lives of Americans and others. For instance:

- The FBI maintains a giant database that contains millions of records covering everything from criminal records to stolen boats and databases with millions of computerized fingerprints and DNA records.
- The Treasury Department runs a database that collects financial information reported to the government by thousands of banks and other financial institutions.
- A "new hires" database maintained by the Department of Health and Human Services, which contains the name, address, social security number, and quarterly wages of every working person in the U.S.
- The federal Department of Education maintains an enormous information bank holding years worth of educational records on individuals stretching from their primary school years through higher education. After September 11, Congress gave the FBI permission to access the database without probable cause.
- State departments of motor vehicles of course possess millions of up-to-date files containing a variety of personal data, including photographs of most adults living in the United States.

The distinction between government and private-sector privacy invasions is fading quickly.

Communications Surveillance

The government also performs an increasing amount of eavesdropping on electronic communications. While technologies like telephone wiretapping have been around for decades, today's technologies cast a far broader net. The FBI's controversial "Carnivore" program, for example, is supposed to be used to tap into the e-mail traffic of a particular individual.

Unlike a telephone wiretap, however, it doesn't cover just one device but (because of how the Internet is built) filters through *all* the traffic on the Internet Service Provider to which it has been attached. The only thing keeping the government from trolling through all this traffic is software instructions that are written by the government itself. (Despite that clear conflict of interest, the FBI has refused to allow independent inspection and oversight of the device's operation.)

Another example is the international eavesdropping program codenamed Echelon. Operated by a part-

nership consisting of the United States, Britain, Canada, Australia, and New Zealand, Echelon reportedly grabs e-mail, phone calls, and other electronic communications from its far-flung listening posts across most of the earth. (U.S. eavesdroppers are not supposed to listen in on the conversations of Americans, but the question about Echelon has always been whether the intelligence agencies of participating nations can set up reciprocal, back-scratching arrangements to spy on each others' citizens.) Like Carnivore, Echelon may be used against particular targets, but to do so its operators must sort through massive amounts of information about potentially millions of people. That is worlds away from the popular conception of the old wiretap where an FBI agent listens to one line. Not only the volume of intercepts but the potential for abuse is now exponentially higher.

The "Patriot" Act

The potential for the abuse of surveillance powers has also risen sharply due to a dramatic post-9/11 erosion of legal protections against government surveillance of citizens. Just six weeks after the September 11 attacks, a panicked Congress passed the "USA PATRIOT Act," an overnight revision of the nation's surveillance laws that vastly expanded the government's authority to spy on its own citizens and reduced checks and balances on those powers, such as judicial oversight. The government never demonstrated that restraints on surveillance had contributed to the attack, and indeed much of the new legislation had nothing to do with fighting terrorism. Rather, the bill represented a successful use of the terrorist attacks by the FBI to roll back unwanted checks on its power. The most powerful provisions of the law allow for:

- **Easy access to records.** Under the PATRIOT Act, the FBI can force anyone to turn over records on their customers or clients, giving the government unchecked power to rifle through individuals' financial records, medical histories, Internet usage, travel patterns, or any other records. Some of the most invasive and disturbing uses permitted by the Act involve government access to citizens' reading habits from libraries and bookstores. The FBI does not have to show suspicion of a crime, can gag the recipient of a search order from disclosing the search to anyone, and is subject to no meaningful judicial oversight.
- **Expansion of the "pen register" exception in wiretap law.** The PATRIOT Act expands exceptions to the normal requirement for probable cause in wiretap law.⁹ As with its new power to search records, the FBI need not show probable cause or even reasonable suspicion of criminal activity, and judicial oversight is essentially nil.
- **Expansion of the intelligence exception in wiretap law.** The PATRIOT Act also loosens the evidence needed by the government to justify an intelligence wiretap or physical search. Previously the law allowed exceptions to the Fourth Amendment for these kinds of searches only if "the purpose" of the search was to gather foreign intelligence. But the Act changes "the purpose" to "a significant purpose," which lets the government circumvent the Constitution's probable cause requirement even when its main goal is ordinary law enforcement.¹⁰
- **More secret searches.** Except in rare cases, the law has always required that the subject of a search be notified that a search is taking place. Such notice is a crucial check on the government's power because it forces the authorities to operate in the open and allows the subject of searches to challenge their validity in court. But the PATRIOT Act allows the government to conduct searches without notifying the subjects until long after the search has been executed.

Under these changes and other authorities asserted by the Bush Administration, U.S. intelligence agents could conduct a secret search of an American citizen's home, use evidence found there to declare him an "enemy combatant," and imprison him without trial. The courts would have no chance to review these decisions – indeed, they might never even find out about them.¹¹

The "TIPS" Program

In the name of fighting terrorism, the Bush Administration has also proposed a program that would encourage citizens to spy on each other. The Administration initially planned to recruit people such as letter carriers and utility technicians, who, the White House said, are "well-positioned to recognize unusual events." In the face of fierce public criticism, the Administration scaled back the program, but continued

Attorney General John Ashcroft issued new guidelines that significantly increase the freedom of federal agents to conduct surveillance on Americans.

to enlist workers involved in certain key industries. In November 2002 Congress included a provision in the Homeland Security Act prohibiting the Bush Administration from moving forward with TIPS.

Although Congress killed TIPS, the fact that the Administration would pursue such a program reveals a disturbing disconnect with American values and a disturbing lack of awareness of the history of governmental abuses of power. Dividing citizen from citizen by encouraging mutual suspicion and reporting to the

government would dramatically increase the government's power by extending surveillance into every nook and cranny of American society. Such a strategy was central to the Soviet Union and other totalitarian regimes.

Loosened Domestic Spying Regulations

In May 2002, Attorney General John Ashcroft issued new guidelines on domestic spying that significantly increase the freedom of federal agents to conduct surveillance on American individuals and organizations. Under the new guidelines, FBI agents can infiltrate "any event that is open to the public," from public meetings and demonstrations to political conventions to church services to 12-step programs. This was the same basis upon which abuses were carried out by the FBI in the 1950s and 1960s, including surveillance of political groups that disagreed with the government, anonymous letters sent to the spouses of targets to try to ruin their marriages, and the infamous campaign against Martin Luther King, who was investigated and harassed for decades. The new guidelines are purely for spying on Americans; there is a separate set of Foreign Guidelines that cover investigations inside the U.S. of foreign powers and terrorist organizations such as Al Qaeda.

Like the TIPS program, Ashcroft's guidelines sow suspicion among citizens and extend the government's surveillance power into the capillaries of American life. It is not just the reality of government surveillance that chills free expression and the freedom that Americans enjoy. The same negative effects come when we are constantly forced to wonder whether we *might* be under observation – whether the person sitting next to us is secretly informing the government that we are "suspicious."

The Synergies of Surveillance

Multiple surveillance techniques added together are greater than the sum of their parts. One example is face recognition, which combines the power of computerized software analysis, cameras, and databases to seek matches between facial images. But the real synergies of surveillance come into play with data collection.

The growing piles of data being collected on Americans represent an enormous invasion of privacy, but our privacy has actually been protected by the fact that all this information still remains scattered across many different databases. As a result, there exists a pent-up capacity for surveillance in American life today – a capacity that will be fully realized if the government, landlords, employers, or other powerful forces gain the ability to *draw together* all this information. A particular piece of data about you – such as the fact that you entered your office at 10:29 AM on July 5, 2001 – is normally innocuous. But when enough pieces of that kind of data are assembled together, they add up to an extremely detailed and intrusive picture of an individual’s life and habits.

Data Profiling and “Total Information Awareness”

Just how real this scenario is has been demonstrated by another ominous surveillance plan to emerge from the effort against terrorism: the Pentagon’s “Total Information Awareness” program. The aim of this program is to give officials easy, unified access to every possible government and commercial database in the world.¹² According to program director John Poindexter, the program’s goal is to develop “ultra-large-scale” database technologies with the goal of “treating the world-wide, distributed, legacy databases as if they were one centralized database.” The program envisions a “full-coverage database containing all information relevant to identifying” potential terrorists and their supporters. As we have seen, the amount of available information is mushrooming by the day, and will soon be rich enough to reveal much of our lives.

Programs like TIA involve turning the defense capabilities of the United States inward and applying them to American people.

The TIA program, which is run by the Defense Advanced Research Projects Agency (DARPA), not only seeks to bring together the oceans of data that are already being collected on people, but would be designed to afford what DARPA calls “easy future scaling” to embrace new sources of data as they become available. It would also incorporate other work being done by the military, such as their “Human Identification at a Distance” program, which seeks to allow identification and tracking of people from a distance, and therefore without their permission or knowledge.¹³

Although it has not received nearly as much media attention, a close cousin of TIA is also being created in the context of airline security. This plan involves the creation of a system for conducting background checks on individuals who wish to fly and then separating out either those who appear to be the most trustworthy passengers (proposals known as “trusted traveler”) or flagging the least trustworthy (a proposal known as CAPS II, for Computer Assisted Passenger Screening) for special attention.

The *Washington Post* has reported that work is being done on CAPS II with the goal of creating a “vast air security screening system designed to instantly pull together every passenger’s travel history and liv-

ing arrangements, plus a wealth of other personal and demographic information” in the hopes that the authorities will be able to “profile passenger activity and intuit obscure clues about potential threats.” The government program would reportedly draw on enormous stores of personal information from data aggregators and other sources, including travel records, real estate histories, personal associations, credit card records, and telephone records. Plans call for using complex computer algorithms, including highly experimental technologies such as “neural networks,” to sort through the reams of new personal information and identify “suspicious” people.¹⁴

The dubious premise of programs like TIA and CAPS II – that “terrorist patterns” can be ferreted out from the enormous mass of American lives, many of which will inevitably be quirky, eccentric, or riddled with suspicious coincidences – probably dooms them to failure. But failure is not likely to lead these programs to be shut down – instead, the government will begin feeding its computers more and more personal information in a vain effort to make the concept work. We will then have the worst of both worlds: poor security and a super-charged surveillance tool that would destroy Americans’ privacy and threaten our freedom.

It is easy to imagine these systems being expanded in the future to share their risk assessments with other security systems. For example, CAPS could be linked to a photographic database and surveillance cameras equipped with face recognition software. Such a system might sound an alarm when a subject who has been designated as “suspicious” appears in public. The Suspicious Citizen could then be watched from a centralized video monitoring facility as he moves around the city.

In short, the government is working furiously to bring disparate sources of information about us together into one view, just as privacy advocates have been warning about for years. That would represent a radical branching off from the centuries-old Anglo-American tradition that the police conduct surveillance only where there is evidence of involvement in wrongdoing. It would seek to protect us by monitoring *everyone* for signs of wrongdoing – in short, by instituting a giant dragnet capable of sifting through the personal lives of Americans in search of “suspicious” patterns. The potential for abuse of such a system is staggering.

The massive defense research capabilities of the United States have always involved the search for ways of outwardly defending our nation. Programs like TIA¹⁵ involve turning those capabilities inward and applying them to the American people – something that should be done, if at all, only with extreme caution and plenty of public input, political debate, checks and balances, and Congressional oversight. So far, none of those things have been present with TIA or CAPS II.

National ID Cards

If Americans allow it, another convergence of surveillance technologies will probably center around a national ID card. A national ID would immediately combine new technologies such as biometrics and RFID chips along with an enormously powerful database (possibly distributed among the 50 states). Before long, it would become an overarching means of facilitating surveillance by allowing far-flung pools of information to be pulled together into a single, incredibly rich dossier or profile of our lives. Before long, office buildings, doctors’ offices, gas stations, highway tolls, subways and buses would incorporate the ID card into their security or payment systems for greater efficiency, and data that is currently scattered and disconnected will get organized around the ID and lead to the creation of what amounts to a national database of sensitive information about American citizens.

History has shown that databases created for one purpose are almost inevitably expanded to other uses; Social Security, which was prohibited by federal law from being used as an identifier when it was first created, is a prime example. Over time, a national ID database would inevitably contain a wider and wider range of information and become accessible to more and more people for more and more purposes that are further and further removed from its original justification.

The most likely route to a national ID is through our driver's licenses.

The most likely route to a national ID is through our driver's licenses. Since September 11, the American Association of Motor Vehicle Administrators has been forcefully lobbying Congress for funds to establish nationwide uniformity in the design and content of driver's licenses – and more importantly, for tightly interconnecting the databases that lie behind the physical licenses themselves.

An attempt to retrofit driver's licenses into national ID cards will launch a predictable series of events bringing us toward a surveillance society:

- Proponents will promise that the IDs will be implemented in limited ways that won't devastate privacy and other liberties.
- Once a limited version of the proposals is put in place, its limits as an anti-terrorism measure will quickly become apparent. Like a dam built halfway across a river, the IDs cannot possibly be effective unless their coverage is total.
- The scheme's ineffectiveness – starkly demonstrated, perhaps, by a new terrorist attack – will create an overwhelming imperative to “fix” and “complete” it, which will turn it into the totalitarian tool that proponents promised it would never become.

A perfect example of that dynamic is the requirement that travelers present driver's licenses when boarding airplanes, instituted after the explosion (now believed to have been mechanical in cause) that brought down TWA Flight 800 in 1996. On its own, the requirement was meaningless as a security measure, but after September 11 its existence quickly led to calls to begin tracking and identifying citizens on the theory that “we already have to show ID, we might as well make it mean something.”

Once in place, it is easy to imagine how national IDs could be combined with an RFID chip to allow for convenient, at-a-distance verification of ID. The IDs could then be tied to access control points around our public places, so that the unauthorized could be kept out of office buildings, apartments, public transit, and secure public buildings. Citizens with criminal records, poor CAPS ratings or low incomes could be barred from accessing airports, sports arenas, stores, or other facilities. Retailers might add RFID readers to find out exactly who is browsing their aisles, gawking at their window displays from the sidewalk or passing by without looking. A network of automated RFID listening posts on the sidewalks and roads could even reveal the location of all citizens at all times. Pocket ID readers could be used by FBI agents to sweep up the identities of everyone at a political meeting, protest march, or Islamic prayer service.

Conclusion

If we do not take steps to control and regulate surveillance to bring it into conformity with our values, we will find ourselves being tracked, analyzed, profiled, and flagged in our daily lives to a degree we can scarcely imagine today. We will be forced into an impossible struggle to conform to the letter of every rule, law, and guideline, lest we create ammunition for enemies in the government or elsewhere.

Our transgressions will become permanent Scarlet Letters visible to all and used by the powerful to increase their leverage over average people.

Our transgressions will become permanent Scarlet Letters that follow us throughout our lives, visible to all and used by the government, landlords, employers, insurance companies and other powerful parties to increase their leverage over average people.

Americans will not be able to engage in political protest or go about their daily lives without the constant awareness that we are –

or could be – under surveillance. We will be forced to constantly ask of even the smallest action taken in public, “Will this make me look suspicious? Will this hurt my chances for future employment? Will this reduce my ability to get insurance?” The exercise of free speech will be chilled as Americans become conscious that their every word may be reported to the government by FBI infiltrators, suspicious fellow citizens or an Internet Service Provider.

Many well-known commentators like Sun Microsystems CEO Scott McNealy have already pronounced privacy dead. The truth is that a surveillance society does loom over us, and privacy, while not yet dead, is on life support.

Heroic measures are required to save it.

Four main goals need to be attained to prevent this dark potential from being realized: a change in the terms of the debate, passage of comprehensive privacy laws, passage of new laws to regulate the powerful and invasive new technologies that have and will continue to appear, and a revival of the Fourth Amendment to the U.S. Constitution.

1. Changing the Terms of the Debate

In the public debates over every new surveillance technology, the forest too often gets lost for the trees, and we lose sight of the larger trend: the seemingly inexorable movement toward a surveillance society. It will always be important to understand and publicly debate every new technology and every new technique for spying on people. But unless each new development is also understood as just one piece of the larger surveillance mosaic that is rapidly being constructed around us, Americans are not likely to get excited about a given incremental loss of privacy like the tracking of cars through toll booths or the growing practice of tracking consumers’ supermarket purchases.

We are being confronted with fundamental choices about what sort of society we want to live in. But unless the terms of the debate are changed to focus on the forest instead of individual trees, too many

Americans will never even recognize the choice we face, and a decision against preserving privacy will be made by default.

2. Comprehensive Privacy Laws

Although broad-based protections against government surveillance, such as the wiretap laws, are being weakened, at least they exist. But surveillance is increasingly being carried out by the private sector – frequently at the behest of government – and the laws protecting Americans against non-governmental privacy invasions are pitifully weak.

In contrast to the rest of the developed world, the U.S. has no strong, comprehensive law protecting privacy – only a patchwork of largely inadequate protections. For example, as a result of many legislators’ discomfort over the disclosure of Judge Robert Bork’s video rental choices during his Supreme Court confirmation battle, video records are now protected by a strong privacy law. Medical records are governed by a separate, far weaker law that allows for widespread access to extremely personal information. Financial data is governed by yet another “privacy” law – Gramm-Leach – which as we have seen really amounts to a license to share financial information. Another law protects only the privacy of children under age 13 on the Internet. And layered on top of this sectoral approach to privacy by the federal government is a geographical patchwork of constitutional and statutory privacy protections in the states.

We are being confronted with fundamental choices about what sort of society we want to live in.

The patchwork approach to privacy is grossly inadequate. As invasive practices grow, Americans will face constant uncertainty about when and how these complex laws protect them, contributing to a pervasive sense of insecurity. With the glaring exception of the United States, every advanced industrialized nation in the world has enacted overarching privacy laws that protect citizens against private-sector abuses. When it comes to this fundamental human value, the U.S. is an outlaw nation. For example, the European Union bars companies from evading privacy rules by transferring personal information to other nations whose data-protection policies are “inadequate.” That is the kind of law that is usually applied to Third World countries, but the EU counts the United States in this category.

We need to develop a baseline of simple and clear privacy protections that crosses all sectors of our lives and give it the force of law. Only then can Americans act with a confident knowledge of when they can and cannot be monitored.

3. New Technologies and New Laws

The technologies of surveillance are developing at the speed of light, but the body of law that protects us is stuck back in the Stone Age. In the past, new technologies that threatened our privacy, such as telephone wiretapping, were assimilated over time into our society. The legal system had time to adapt and reinterpret existing laws, the political system had time to consider and enact new laws or regulations, and the culture had time to absorb the implications of the new technology for daily life. Today, however, change is happening so fast that none of this adaptation has time to take place – a problem

that is being intensified by the scramble to enact unexamined anti-terrorism measures. The result is a significant danger that surveillance practices will become entrenched in American life that would never be accepted if we had more time to digest them.

Since a comprehensive privacy law may never be passed in the U.S. – and certainly not in the near future – law and legal principles must be developed or adapted to rein in particular new technologies such as surveillance cameras, location-tracking devices, and biometrics. Surveillance cameras, for example, must be subject to force-of-law rules covering important details like when they will be used, how long images will be stored, and when and with whom they will be shared.

4. Reviving the Fourth Amendment

The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no warrants shall issue, but upon probable cause, supported by oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.

– Fourth Amendment to the U.S. Constitution

The Fourth Amendment, the primary Constitutional bulwark against Government invasion of our privacy, was a direct response to the British authorities' use of "general warrants" to conduct broad searches of the rebellious colonists.

Historically, the courts have been slow to adapt the Fourth Amendment to the realities of developing technologies. It took almost 40 years for the U.S. Supreme Court to recognize that the Constitution applies to the wiretapping of telephone conversations.¹⁶

In recent years – in no small part as the result of the failed "war on drugs" – Fourth Amendment principles have been steadily eroding. The circumstances under which police and other government officials may conduct warrantless searches has been rapidly expanding. The courts have allowed for increased surveillance and searches on the nation's highways and at our "borders" (the legal definition of which actually extends hundreds of miles inland from the actual border). And despite the Constitution's plain language covering "persons" and "effects," the courts have increasingly allowed for warrantless searches when we are outside of our homes and "in public." Here the courts have increasingly found we have no "reasonable expectation" of privacy and that therefore the Fourth Amendment does not apply.

But like other Constitutional provisions, the Fourth Amendment needs to be understood in contemporary terms. New technologies are endowing the government with the 21st Century equivalent of Superman's X-ray vision. Using everything from powerful video technologies that can literally see in the dark, to biometric identification techniques like face recognition, to "brain fingerprinting" that can purportedly read our thoughts, the government is now capable of conducting broad searches of our "persons and effects" while we are going about our daily lives – even while we are in "public."

The Fourth Amendment is in desperate need of a revival. The reasonable expectation of privacy cannot be defined by the power that technology affords the government to spy on us. Since that power is increasingly limitless, the “reasonable expectation” standard will leave our privacy dead indeed.

But all is not yet lost. There is some reason for hope. In an important pre-9/11 case, *Kyllo vs. U.S.*,¹⁷ the Supreme Court held that the reasonable expectation of privacy could not be determined by the power of new technologies. In a remarkable opinion written by conservative Justice Antonin Scalia, the Court held that without a warrant the police could not use a new thermal imaging device that searches for heat sources to conduct what was the functional equivalent of a warrantless search for marijuana cultivation in Danny Kyllo’s home.

The Court specifically declined to leave Kyllo “at the mercy of advancing technology.” While *Kyllo* involved a search of a home, it enunciates an important principle: the Fourth Amendment must adapt to new technologies. That principle can and should be expanded to general use. The Framers never expected the Constitution to be read exclusively in terms of the circumstances of 1791.

Notes

- 1 Jess Bravin, "Washington Police to Play 'I Spy' With Cameras, Raising Concerns," *Wall Street Journal*, Feb. 13, 2002.
- 2 See http://www.ncjrs.org/school/ch2a_5.html.
- 3 See <http://www.scotcrim.u-net.com/researchc2.htm>.
- 4 The success rate of face recognition technology has been dismal. The many independent findings to that effect include a trial conducted by the U.S. military in 2002, which found that with a reasonably low false-positive rate, the technology had less than a 20% chance of successfully identifying a person in its database who appeared before the camera. See http://www.aclu.org/issues/privacy/FINAL_1_Final_Steve_King.pdf, 17th slide.
- 5 Richard H.P. Sia, "Pilotless Aircraft Makers Seek Role For Domestic Uses," *CongressDaily*, December 17, 2002.
- 6 Data surveillance is often loosely referred to as "data mining." Strictly speaking, however, data mining refers to the search for hidden patterns in large, pre-existing collections of data (such as the finding that sales of both beer and diapers rise on Friday nights). Data mining need not involve personally identifiable information. Data surveillance, on the other hand, involves the collection of information about an identifiable individual. Note, however, that when data surveillance is carried out on a mass scale, a search for patterns in people's activities – data mining – can then be conducted as well. This is what appears to be contemplated in the Total Information Awareness and CAPS II programs (see below).
- 7 Dana Hawkins, "As DNA Banks Quietly Multiply, Who is Guarding the Safe?" *U.S. News & World Report*, Dec. 2, 2002.
- 8 Glenn R. Simpson, "Big Brother-in-Law: If the FBI Hopes to Get The Goods on You, It May Ask ChoicePoint" *Wall St. Journal*, April 13, 2001.
- 9 The expanded exception involves what are called "pen register/trap & trace" warrants that collect "addressing information" but not the content of a communication. Those searches are named after devices that were used on telephones to show a list of telephone numbers dialed and received (as opposed to tapping into actual conversations). The PATRIOT Act expands the pen register exception onto the Internet in ways that will probably be used by the government to collect the actual content of communications and that allow nonspecific "nationwide" warrants in violation of the Fourth Amendment's explicit requirement that warrants "must specify the place to be searched."
- 10 In August, the secret "FISA" court that oversees domestic intelligence spying released an opinion rejecting a Bush Administration attempt to allow criminal prosecutors to use intelligence warrants to evade the Fourth Amendment entirely. The court noted that agents applying for warrants had regularly filed false and misleading information. In November 2002, however, the FISA appeals court (three judges chosen by Supreme Court Chief Justice William Rehnquist), meeting for the first time ever, ruled in favor of the government.
- 11 See Charles Lane, "In Terror War, 2nd Track for Suspects," *Washington Post*, December 1, 2002. Online at <http://www.washingtonpost.com/wp-dyn/articles/A58308-2002Nov30.html>.
- 12 See "Pentagon Plans a Computer System That Would Peek at Personal Data of Americans," *New York Times*, Nov. 9, 2002; "US Hopes to Check Computers Globally," *Washington Post*, Nov. 12, 2002; "The Poindexter Plan," *National Journal*, Sept. 7, 2002.
- 13 Quotes are from the TIA homepage at <http://www.darpa.mil/iao/index.htm> and from public 8/2/02 remarks by Poindexter, online at <http://www.fas.org/irp/agency/dod/poindexter.html>.
- 14 Robert O'Harrow Jr., "Intricate Screening Of Fliers In Works," *Washington Post*, Feb. 1, 2002, p. A1.
- 15 The TIA is just one part of a larger post-9/11 expansion of federal research and development efforts. The budget for military R&D spending alone has been increased by 18% in the current fiscal year to a record \$58.8 billion. Bob Davis, "Massive Federal R&D Initiative To Fight Terror Is Under Way," *Wall Street Journal*, November 25, 2002.
- 16 In 1967 the Supreme Court finally recognized the right to privacy in telephone conversations in the case *Katz v. U.S.* (389 US 347), reversing the 1928 opinion *Olmstead v. U.S.* (277 US 438).
- 17 190 F.3d 1041, 2001.



Bhopal Lives

By Suketu Mehta

The Village Voice

December 3, 1996 (part 1) and December 10, 1996 (part 2),
the latter entitled “After Bhopal”

The 1984 Union Carbide Toxic-Gas Disaster Killed 10,000 People — and Has Changed Everything for Its Survivors

Next Tuesday, December 3, the International Medical Commission-Bhopal (IMCB) will release its final report on the current medical, social, and economic status of the Union Carbide disaster, a leak of toxic gas that claimed around 10,000 lives in Bhopal, India, 12 years ago.

The report, the culmination of a three-year study by a group of doctors affiliated with prestigious institutions in the U.S., Europe, and Asia, is the first comprehensive, peer-reviewed study of the chronic effects of the disaster that has been released publicly.

The commission found that up to 50,000 survivors are suffering from partial or total permanent disability as a consequence of the gas disaster. In addition the widely recognized lung and eye injuries, its report details medical conditions that have never been identified before, such as neurotoxicological effects (damage to the brain and central nervous system). They affect short-term memory, balance, and motor skills—they affect the survivors’ ability to hold jobs, and their children’s ability to read and write.

The study documents, for the first time, post-traumatic stress syndrome in the survivors. “People were buried alive,” says Dr. Rosalie Bertell, one of the commissioners. “Some of them actually were in a pile of bodies to be burned and came to—you can imagine the nightmares and panic attacks after that.”

According to earlier studies done by the Indian Council of Medical Research, descendants up to the third generation of survivors may sustain genetic damage leading to cancer and abnormalities in offspring. The new findings were not available to the Supreme Court of India when it imposed a settlement for damages in 1989, which the commission found to be “decidedly inadequate.” The report, therefore, should provide new grounds to reopen the case.


Bhopal has joined the roster of internationally recognized symbol-places—along with Hiroshima, Auschwitz, and Chernobyl—whose very names have become synonymous with the tragedies that have taken place within their precincts. Mention the word *Bhopal* to a person outside India, and they won’t think of a graceful city on the hills above two lakes with some of the most glorious Muslim architecture in India. They will think about what happened the night of December 2 and the early morning of December 3, 1984, when an accident at the chemical plant owned by Union Carbide of Danbury, Connecticut, led to history’s worst industrial disaster.

There is a pornography of images of disaster in the Third World—famine, floods, war, and earthquakes. Quick television interviews with the victims reinforce those images. And, as with all pornography, the net effect is this: the affected people lose their individuality, their humanity, and we, the viewers, who have no idea about their lives, begin to distance ourselves from them. As it is, they all look so foreign to us: all these brown or black people, poor things. A lot has been written about the bare facts of Bhopal disaster: how it might have happened, how many died, how many were injured. This article, the first of two parts, examines what has rarely been portrayed: the complexity of people’s individual responses to an enduring disaster.

The Night of the Gas

In May 1982, a Union Carbide inspection team from the Danbury headquarters visited the Bhopal plant and found 61 safety and maintenance problems, 30 of them major. A series of gas leaks had already resulted in the death of one factory worker and injuries to several others. Five months before the night of the accident, vital refrigeration and cooling systems had been shut down. Around the same time, the maintenance crew was reduced from six to two workers as part of a cost-cutting drive. Local lawyers and journalists had been warning Union Carbide for months that the plant could be dangerous to its neighbors. The company responded that such fears were “absolutely baseless.”

In the early morning hours of December 3, 1984, water entered under still disputed circumstance an underground storage tank containing 90,000 pounds of methyl isocyanate, a highly toxic chemical used to make pesticides. This set off the following reaction:



Forty-one tons of methyl isocyanate along with a stew of other highly toxic gases possibly including hydrogen cyanide boiled over and burst through the tank at a temperature of over 200 degrees Celsius and a rate of over 40,000 pound an hour. This was the birth of what scientists later name “Bhopal Toxic Gas.” The gas rose from the plant, then sedately, unhurriedly, floated out over the sleeping city.

Bhopalis have very personal relationships with “the gas.” Accounts of that night—again when in Bhopal someone says “that night,” they mean the night of December 2–3, 1984—describe how the gas was going toward Jahangirabad or Hamidia Road; how it hovered a few feet above the ground at some places or how it hugged the wet farm earth in others; how it killed buffalo and pigs but spared chickens and mosquitoes; how it made all the leaves of a peepul tree turn black and how it had a particular hunger for the tulsi plant; how it would travel down one side of a road but not the other, like rain falling a few feet from you while you’re standing in the sunshine. People know the gas like a member of their family—they know its smell, its color, its favorite foods, its predilections. One thing everybody remembers is the smell of chilies burning. Chilies are normally burned to ward off the evil eye, when, for example, a child is sick. People woke up and thought it must be a powerful evil eye that’s being driven away, the stink is so strong.

As people ran with their families, they saw their children falling beside them, and often had to choose which ones they would carry on their shoulders and save. This image comes up again and again in the dreams of the survivors: in the stampede, the sight of a hundred people walking over the body of their child.

Iftekhar Begum went out on the morning after the gas to help bury the Muslim dead. There were so many that she could not see the ground—she had to stand on the corpses to wash them. As she stood on the bodies, she noticed that many of the dead women had flowers in their hair. The gas had come on a Sunday, a night when people had dressed up to go out to a film or to someone’s house for dinner. The women had, as is common all over India, braided their hair with jasmine or mogra—small, fragrant flowers.

When Iftekhar Begum came back from the graveyard, all her fingertips were bleeding, she had sewn so man shrouds.

Arun’s Story

What would you do if you woke up one night when you were 13 years old and by the morning, seven of the 10 members of your immediate family were dead? How would your life change?

When I first meet the young man I will call Arun, to whom this happened, he is busy writing a wedding invitation card. Not his own. Not anybody’s, in fact; there will only be one copy of this invitation, and it will be show to the judge in the gas victims’ claims court. There is a Muslim woman with him. She has allotted 50,000 rupees (\$1249) in compensation for her injuries which the government has kept in a fixed-deposit bank account to prevent her from spending it all at once. To withdraw funds from her account, she has to demonstrate to the judge that she has some compelling need, like the wedding of a daughter. Arun is wise to the inscrutable ways of the authorities; for a consideration, he will help her get her money out. So he sits next to me making us this invitation to a wedding that will never be.

Arun's fee for writing up the affidavit and printing up one copy of the ceding card at a printing press (such costs him 100 rupees, or \$3) is 3000 rupees (\$86). This, he points out, is less than what a lawyer would charge, which is 10 percent, 5000 rupees (\$143). "The lawyers hate me," he crows.

The gas victim Arun loves his life. He wakes up at noon, massages himself with mustard oil, and spends the afternoon sitting on the newly constructed balcony of his house, chatting with friends. In the evenings, he drinks or goes to the Hotel International and asks to see the "special menu," which consists of several pages of pictures of the women they have for sale upstairs. On an occasional Sunday, he'll get partridges, which he kills with his own hands, cooks, and shares with his friends, who seem to be in awe of him. Three or four times a month he goes to the claims courts on behalf of someone, and that's enough money for him, mostly.

Arun first learned of the deaths of his parents and five siblings when he saw their photos stuck up on the wall by the side of the road. Till then people would tell him but he didn't believe them. Looking at the pictures the government had put up to alert survivors, Arun did not cry. Arun claims he has never once cried. "There were so many corpses. Who will you cry over? After a while, the heart becomes quiet."

On the night of the gas, Arun fell in love. As Arun and his family ran, as one by one his parents, brothers, sisters dropped to the ground or got separated from him, Arun felt someone holding his hand and leading him. On they ran, through the chaotic streets. That was the beginning of Arun's first love. The girl holding his hand lived in his neighborhood, and later on, she fed him and took care of him.

That girl was the first of his neighbors to adopt Arun and take care of him, but she was by no means the last. There were other families in the slum, his extended family in Lucknow, a rickshaw driver and his wife, and finally, the activist Satinath Sarangi, known with much love as "Sathyu" among the survivors. Arun moved into Sathyu's house and became a poster child of the activist movement; his story was widely used and he was recruited by all manner of groups, including the youth wing of the Communist Party of India, the state's major political parties, and almost all of the activist groups working on Bhopal. Arun became a kind of traveling victim, going on tours to talk about the tragedy that had devastated his family, not only all over India, but also, twice, to the United States. He was a natural. "At the age of 15 I learned to give such good answers that the journalists loved me," he recalls gleefully. On one of his trips to the U.S., Arun and a couple of the other survivors, while attempting to distribute literature in the Houston hotel where the annual meeting of Carbide's shareholders was being held, were arrested by the police and spent 10 hours in jail. Arun was impressed by the fact that the American jail was air-conditioned.

But gradually, Arun went from being a victim to something of a predator. Sundry scam inevitably pop up in any community where a large amount of money enters the scene all at once, and Arun has learned how to profit from them. So, for a commission, using an efficient system of bribes paid to every one from clerks to judges, Arun will extract the gas victims' compensation money from the clutches of the government. He is also a loan shark; he advances money at exorbitant rates of interest to illiterate migrants from the countryside, actively assists them in speeding it in the Bhopal bars, and beats them soundly if they cannot pay him. He has a gang, which will assault people's enemies for a price. He points to my knee—300 rupees (\$9) for breaking that—and then to my arm—460 rupees (\$10) for that.

Once, when Sathyu was remonstrating with Arun about his misdeeds, Arun responded, "Look at Warren Anderson [then Union Carbide's chairman]. He got away with killing so many people. If he can get away, so can I." Besides, Arun sometimes puts his potential for violence to good use. Though he is Hindu, he put his life on the line during the bloody Hindu-Muslim riots of 1992, when he stood guard outside Muslim homes with a sword.

Every year, on the anniversary of the gas leak, the chief minister holds a big commemorative public meeting and invites a number of victims. Arun will go this year and ask him for a favor—a coveted license to sell kerosene, which he'll divert to the black market. The chief minister, he tells me with a laugh, will never refuse such a famous orphan anything when there are so many journalists present.

Arun hates the term "gas victim." Once, in 1987, when he and other survivors were traveling to a demonstration, the train stopped at a station and the loudspeakers boomed out: "Now, all the gas victim children from Bhopal, go and play in the special waiting room." Arun sought out the government officer responsible for the announcement and swore: "Your mother's cunt."

"Is it stamped on my forehead, 'gas victim'?" he asks me. "Should I beg for pity, Hai Allah, help me, give me some food, I'm a gas victim?" Arun instructs his kid brother: "If a man thinks himself to be weak, he *will* be weak." Accordingly he insists the 12-year-old boy get up at six every morning to do calisthenics. There is a reason, Arun believes, that he himself has remained strong. "Gas? I shit gas out of my ass. You drink enough, you smoke enough, and there won't be any gas." To prove that he is stronger

than anybody, gas-affected or not, Arun steps in front of a passing minibus and looks at me. “Shall I beat up the driver?” he asks.

But Arun also tells me, matter-of-factly, that he’s been having *gabrahat*. This is a condition that is commonly reported by survivors, and there’s no exact English translation. All of a sudden, Arun’s heart will beat wildly, he’ll start sweating, and his mind will flood with anxiety. This lasts for about 10 minutes. Since most of the people affected by the gas lived in the poorer part of Bhopal, they were, by and large, not deemed worthy of psychiatric treatment or counseling. It’s certainly not anything the government will give Arun, or anyone, compensation for.

One night, three of us—Arun, his sidekick Ramdayal, and I—sit in the gas victims’ beer bar, a shed off the housing colony. Around us are gas victims, all of them men, drinking with the compensation money they should be spending to get treatment for their wives, education for their kids. As the evening progresses, Arun and Ramdayal are getting a lot more drunk than I am because they are drinking whiskey-and-beer cocktails. Presently, they get into theological argument: Was God present the night of the gas?

On the night of the gas, as his family was dying, as he was falling in love, Arun lost his faith in God. “Mother’s prick, six, seven people died—where the fuck was Ganesh? If I met him, I’d beat him with shoes and chase him off, mother’s prick, sister’s prick. The gas came, Ganesh fucked my mother, then ran away. If my mother were here I wouldn’t have a history.” I’ve never seen him so angry; he’s almost shouting, and finally he becomes completely incoherent and the gaps between the obscenities vanish and it’s all just obscenities: mother’s prick, sister’s prick. When he calms down, he says, “Only work is karma, work is the fruit.” Later I realize what he’s just said, in a single sentence: Krishna’s teaching to Arjuna in the Bhagavad Gita.

The Lifting of the Veils

In the years after the poison cloud came down from the factory, the veils covering the faces of the Muslim women of Bhopal started coming off.

The Bhopal Gas Peedit Mahila Udyog Sangathan (the Bhopal Gas-Affected Women Workers’ Organization), or GBPMUS, is the most remarkable and, after all these years, the most sustained movement to have sprung up in response to the disaster. The BGPMUS grew out of a group of sewing centers formed after the events to give poor women affected by the gas a means of livelihood. As they came together into the organization, the women participated in hundreds of demonstrations, hurried attorneys to fight the case against Carbide as well as the Indian government, and linked up with activist movements all over India and the world.

On any Saturday in Bhopal, you can go to the park opposite Lady Hospital and sit among an audience of several hundred women and watch all your stereotypes about traditional Indian women get shattered. I listened as a grandmother in her sixties got up and hurled abuse at the government with a vigor that Newt Gingrich would envy. She was followed by a woman in a plain sari who spoke for an hour about the role of multinationals in the third world, the wasteful expenditure of the government on sports stadiums, and the rampant corruption to be found everywhere in the country.

As the women of Bhopal got politicized after the gas, they became aware of other inequities in their lives too. Slowly, the Muslim women of the BGPMUS started coming out of the veil. They explained this to others and themselves by saying: look, we have to travel so much, give speeches, and this burkha, this long black curtain, is hot and makes our health worse.

But this was not a sudden process; great care was paid to social sensitivities. When Amida Bi wanted to give up her burkha, she asked her husband. “My husband took permission from his older brother and my parents.” Assent having been give all around, Amida Bi now goes all over the country without her veil, secure in the full support of her extended family.

Her daughters however, are another matter. Having been married out to other families, they still wear the burkha. But Amida Bi refuses to allow her own two daughters-in-law, over whom she has authority, to wear the veil at all. “I don’t think the burkha is bad,” she says. “But you can also do shameful things while wearing a burkha.”

Half of the Muslim women still attuning the rallies have folded up their burkhas forever.

Sajida Bano's Story

Sajida Bano never had to use a veil until her husband died. He was the first victim of the Carbide plant. In 1981, three years before the night of the gas, Ashraf was working in the factory when a valve malfunctioned and he was splashed with liquid phosgene. He was dead within 72 hours. After that, Sajida was forced to move with her two infant sons to a bad neighborhood, where if she went out without the burkha she was harassed. When she put it on, she felt shapeless, faceless, and anonymous: she could be anyone's mother, anyone's sister.

In 1984, Sajida took a trip to her mother's house in Kanpur, and happened to come back to Bhopal on the night of the gas. Her four-year-old son died in the waiting room of the train station, while his little brother held on to him. Sajida had passed out while looking for a taxi outside. The factory had killed the second of the three people Sajida loved most. She is left with her surviving son, now 14, who is sick in body and mind. For a long time, whenever he heard a train whistle, he would run outside, thinking that his brother was on that train.

Sajida Bano asked if I would carry a letter for her to "those Carbide people," whoever they are. She wrote it all in one night, without revision. She wants to eliminate distance, the food chain of activists, journalists, lawyers, and governments between her and the people in Danbury. Here, with her permission, are excerpts that I translated:

Sir,

Big people like you have snatched the peace and happiness of us poor people. You are living it up in big palaces and mansions. Moving around in cars. Have you ever thought that you have wiped away the marriage marks from our foreheads, emptied our laps of children, bathed us in poison, and we are sobbing, but death doesn't come. Like a living, walking corpse you have left us. At least tell us what our crime was, for which such a big punishment has been given. If with the strength of your money you had shot us all at once with bullets, then we wouldn't have to die such miserable sobbing deaths.

You put your hand on your heart and think, if you are a human being: if this happened to you, how would your wife and children feel? Only this one sentence must have caused you pain.

If this vampire Union Carbide factory would be quiet after eating my husband, if heartless people like you would have your eyes opened, then probably I would not have lost my child after the death of my husband. After my husband's death my son would have been my support. But before he would grow you uprooted him. I don't know myself why you have this enmity against me.

Why have you played with my life so much? What was I, a poor helpless woman, spoiling of yours that even after taking my husband you weren't content. You ate my child too. If you are a human being and have a human heart then tell me yourself what should be done with you people and with me. I am asking you only, tell me, what should I do?

Negative-Positive

The gas changed people's lives in ways big and small. Harishankar Magician used to be in the negative-positive business. It was a good business. He would sit on the pavement; hold up a small glass vial, and shout, "Negative to positive!" Then, hollering all the while, he would demonstrate. "It's very easy to put negative on paper. Take this chemical, take any negative, put it on any paper, rub it with this chemical, then put it in the sun for only 10 minutes. This is a process to make a positive from a negative." By this time a crowd would have gathered to watch the miraculous transformation of a plain film negative into an image on a postcard. In an hour and a half, Harishankar Magician could easily earn 50, 60 rupees (\$2) in this business. Then the gas came.

It killed his son and destroyed his lungs and his left leg. In the negative-positive business, he had to sit for hours. He couldn't do that now with his game leg, and he couldn't shout with his withered lungs. So Harishankar Magician looked for another business that didn't require standing and shouting. Now he wanders the city, pushing a bicycle that bears a box with a hand-painted sign: "ASTROLOGY BY ELECTRONICE MINI COMPUTER MACHIN."

Passersby, seeing the mysterious box, gather spontaneously to ask what it is. He invites them to put on the stethoscope, which is a pair of big padded headphones attached to the Machin. Then the front panel of the Machin comes alive with flashing Disco Lights, rows of red and yellow and green colored bulbs. The Machin, Harishankar Magician tells his customers, monitors their blood pressure, and then tells their fortune through the stethoscope. The fee is two rupees (six cents). Harishankar doesn't like this business; with this, unlike his previous trade, he thinks he is peddling a fraud. Besides, he can only do it for an hour and a half a day, and clears only about 15 rupees (43 cents).

Harishankar Magician is sad. He yearns for the negative-positive business. Once the activist Sathyu took a picture of Harishankar's son, who was born six days before the gas came. He died three years later. Harishankar and his wife have no photographs of their dead boy in their possession, and they ask Sathyu if he can find the negative of the photo he took. Then they will use the small vial of chemical to make a positive of their boy's negative, with only 10 minutes of sunlight.

The Plague of the Lawyers

Almost immediately after the disaster, the American lawyers started coming, by the dozens. Out they stepped from the plane, blinking and squinting in the strong Bhopal light, covering their noses with handkerchiefs as they stepped gingerly through the dung-strewn lanes of the slums, glad-handing the bereaved, pointing to their papers and telling their translators to tell the victims "MILLIONS of rupees, you understand? MILLIONS!" And so the people signed, putting their names down in Hindi, or just with their thumbprints.

In the Oriya slum, 11 years later, word spreads that a visitor from America has come, and a cluster of people come to meet me. A young man, Bhimraj, and his mother, Rukmini, approach me hesitantly, holding out a carefully preserved piece of paper. "The American government gave us this," he says. "Can you tell me what it says?"

I look at the document. It is a legal contract.

"Contract between law office of Pat Maloney, PC, of the city of San Antonio, Bexar County, Texas, and Suresh.

"Client agrees to pay attorney as attorney's fee for such representation one third (33%) of any gross recovery before action is filed, forty percent (40%) of any gross recover after action is filed but before the commencement of trial, and fifty percent (50%) of any gross recover after commencement of trial.

"This contract is performable in Bexar County, Texas."

On the night of the gas, Rukmini abandoned her three-year-old son, Raju, who was dead, and ran with her five-year-old daughter, Rajini, who died three days later. When the lawyers came, they got Rukmini's husband, Suresh, to put his name down in Hindi on this document. They took the family's pictures. "They didn't even send us a copy," says Rukmini. That was the last the family heard from the man they believed came on behalf of "the American government." So now they ask me, what should they do with this paper that they've been holding on to for 11 years?

"Tear it up and throw it away," I tell them. "It's junk." They look at me, their faces blank, not understanding.

(When I returned to America, I tried to contact attorney Pat Maloney. He did not return phone calls.)

Responding to such abuses, the Indian parliament passed a law declaring itself the sole legal representative of all the Bhopal gas victims. It sued Carbide in federal court in New York. The court held that the proper venue for the case should be in India; spectators were treated to the uniquely edifying spectacle of hearing the Indian government's lawyers argue the inadequacy of its own legal system, countering Carbide's lavish testaments to the excellence of the very same system. The reason was simple: everybody knew that any potential damage award given out by an Indian court would be considerably smaller than one awarded by a U.S. court. Had the victims succeeded in suing the company in its home country and winning, they would probably have bankrupted the giant corporation, much as the asbestos liability cases bankrupted the Manville Corporation and breast-implant litigation bankrupted Dow Corning.

As it transpired, after prolonged legal wrangling, the Indian Supreme Court unilaterally, without giving the victims a chance to make their case, imposed a settlement to the amount of \$470 million, with the government to make up any shortfall. The government had asked for \$3 billion from Carbide. Carbide

executives were delighted; they speedily transferred money to the government. That was in 1989. The first victim did not see the first rupee of Carbide's money until Christmas of 1992, eight years after the night of the gas. A total of 597,000 claims for compensation have been filed. As of May 1996, the government has passed rulings on only about half of them – 302,422 – and awarded compensation for injuries to 288,000 Bhopalis. Out of the total settlement amount of \$470 million plus interest since 1989, the government had, by May of 1996, only disbursed some \$241 million.

The Quantification of Loss

A government psychiatrist who has done a close study of the minds of the gas victims has come to this conclusion: they don't want to work. "You can't get domestic help in Bhopal nowadays," the doctor complained to me. "If a family has five affected people who get 200 rupees (\$6) each [in interim relief], that's a thousand rupees a month, so they don't want to work."

There is a widespread belief that the people destroyed by the gas—who tended to come from the poorer sections of Bhopal—aren't receiving deserved compensation for grievous injuries that they are legally and morally entitled to, but some sort of unearned windfall that's made them indolent. This belief is prevalent among the rich in Bhopal, government officials, and Carbide executives.

J.L. Ajmani is the secretary of the gas relief department of Madhya Pradesh state, and he won't give mean interview. Armani is a man of the 21st century. In his luxurious office, he has a computer, a bank of three phones, a sofa, a huge desk, and an executive chair in which he reposes under a big picture of Mahatma Gandhi. While brushing me off, he keeps tapping into his digital diary. I ask him about allegations of corruption in his department. He laughs fearlessly. "It's been 11 years. Volumes have been written. You also write."

Although the government isn't releasing figures about the average amount of rewards, the welfare commissioner's office told me that the maximum compensation awarded for deaths is 150,000 rupees (\$4286), except in a small handful of cases. Mohammed Laique, a local lawyer who has been representing claimants from the beginning, gave me the standard rates of compensation. For most deaths, the amount awarded is 100,000 rupees (\$2857). For personal injury cases, 90 percent get 25,000 rupees, or \$714 (the award bestowed on most of the survivors I spoke to directly).

Of these amounts, says Laique, "claimants lose between 15 percent and 20 percent at the outset in bribes. To get money out early, you pay another 10 percent." Then there are sundry small bribes. Clerks in government offices demand anywhere from 100 to 2000 rupees (\$57) to move papers, depending on the size of the awards. The payments the government has been disbursing since 1990 for interim relief (200 rupees, or \$6 a month) are also deducted from the awards. This means that from an award of 25,000 rupees, the maimed survivor in September 1995 could expect to receive as little as 7600 rupees. Two hundred and seventeen dollars.

Union Carbide claims that the compensation is "more than generous by any Indian standard." Is it really? For comparison, Laique pulls out the schedule of standard compensation set by Indian Railways for railway accidents. The schedule is gruesomely specific:

- In case of death: 200,000 minimum (\$5714)
- For disability of 1 leg: 120,000 (\$3429)
- If one or two hands are cut off: 200,000
- If one or two legs are severed: 200,000
- Thumb cut off: 60,000 (\$1714)
- If four fingers cut off from one hand: 100,000 (\$2857)
- 3 fingers cut off: 60,000
- 2 or 1 fingers cut off: 40,000 (\$1143)
- Breast cut off: 180,000 (\$5143)
- For problem with 1 eye: 80,000 (\$2286)
- Hip joint fracture: 40,000
- Minimum for bodily injury: 40,000

"And the railways give very fast decisions, plus interest after three months," adds Laique. During the bloody communal rioting that followed the destruction of the Babri Masjid mosque in Ayodhya in 1992,

the government gave a minimum of 200,000 rupees (\$5714) to the families of each person killed; these were people of the same socioeconomic status as Carbide's victims. It's clear that, if a Bhopali had any choice in the instrument of his death, it would be financially much more advantageous to be killed or maimed in a train wreck or at the hands of a religious fanatic than through an American multinational's gas cloud.

[December 10 continuation]

On the day after a Union Carbide plant leaked a toxic gas that would kill 10,000 people in Bhopal, India, Warren Anderson, Carbide's chairman at the time, flew to Bhopal to see the situation for himself and offer aid. The chairman was propelled by a visceral, human impulse, and acted against the advice of his lawyers and public relations people; he was promptly arrested, detained for several hours, and put on a plane to New Delhi. He was granted bail and flew home a few days later.

When he returned to Connecticut, Anderson met his real enemies—reporters, lawyers, shareholders, and consultants, hounding him with questions, offering advice. He fled with his wife and his mother-in-law and holed up for a week in a Stamford hotel, having all their meals sent up, “a grown man hiding in a hotel room,” as he later put it. After the accident, he had trouble sleeping. And well he might. Anderson is now wanted on charges of culpable homicide in India, and is rumored to be living quietly in Vero Beach, Florida.

Anderson's 1984 Bhopal expedition marked the last time a senior Carbide executive from Danbury got his shoes soiled in the city.

In the years after the tragedy, Carbide has admitted “moral responsibility” for the disaster. The company proposed a variety of small projects to aid the victims, including setting up a vocational center and contributing \$2 million toward relief efforts. After the assets of its Indian subsidiary were seized by Indian courts, Carbide made a virtue out of necessity and, at the Supreme Court's direction, announced that it would use the frozen assets to set up a trust to build a new hospital for the survivors. The company refuses to use any of its unencumbered assets toward this laudable endeavor.

Throughout, it has stoutly maintained that the disaster was a result of deliberate sabotage. The Carbide hypothesis goes like this: a disgruntled employee, upset about being demoted, deliberately introduced water into the methyl isocyanate tank, setting off the deadly chemical reaction. Subsequently, all the employees and supervisors on duty at the plant at that time decided, for reasons best known to them, to engage in a massive cover up of the real causes of the accident, and have successfully maintained their conspiracy through the 11 ½ years since.

“Much of the world's safety engineering community doubts the veracity of Carbide's sabotage evidence,” writes Wil Lepkowski, the American reporter who has most closely followed Bhopal, in *Chemical and Engineering News*. That evidence, Lepkowski points out, has never been subjected to scientific peer review or presented in court. Carbide will not name the saboteur, even though it promised to do so in court “at the appropriate time.” That was in 1986; a decade later, an appropriate time has still not been found.

At the moment, there is no Carbide employee in Bhopal. There is no executive, no secretary, no engineer personally supervising the setting up of their hospital; nobody walking through the slums to make sure that the people they visited their holocaust upon are being adequately taken care of.

Carbide is doing nothing to monitor the settlement amounts, to ensure that the victims' financial needs are being taken care of; its labs are doing no research, nor is the company funding any, on the long term effects of methyl isocyanate; and there is no monument in Danbury or at any other company site to the gas victims of Bhopal. As Carbide's chief of public relations Bob Berzok put it to me when refusing my request to talk to anyone but himself at the company, anyone at all from the president down to a cafeteria worker, “This does go back 10 years and I'm not interested in disrupting the business going on here. I inquired of several people and the feeling in general for those who were here 10 years ago was that there really was no interest in discussing their personal feelings [about Bhopal].”

Berzok himself has been to India some 15 times in connection with the Bhopal disaster, not to help the victims but to help the Indian subsidiary better manage its public relations. Staying at the posh guest house that Carbide used to own in Shamla Hills, Berzok has never once visited the slum colonies where the victims live and die; and he doesn't recall a single name or a single distinguishing feature of any of the victims. He saw some of them in the medical stations set up in the old city. “There were some people that were having difficulties breathing,” is what he remembers.

For Union Carbide, Bhopal was a hit and run accident.

“I wanted him to apologize,” says Syed Mohammed Irfan about Warren Anderson. Irfan lost his sister and his health because of the Carbide factory. Since the accident, his wife is terrified of living in Bhopal and has left him to live elsewhere in the state. “I wanted him to apologize, be humble. Say we made a mistake; get treatment, we’ll pay for it. We wouldn’t have hung him. This didn’t happen.” Carbide may have accepted “moral responsibility” for the disaster, but has never apologized to the people of Bhopal.

So Irfan’s views have changed. “Now if I meet Anderson in the street I’ll kill him.”

I have also met people who don’t think Carbide is to blame. A high school teacher who lost her niece, and has seen her own health suffer, told me, “I feel no anger toward Carbide. It’s the fault of the technology.” All of Bhopal is not a vengeful mob thirsting for revenge. Berzok emphasizes that whenever he was in Bhopal, traveling openly as a Carbide employee from the U.S., “I was treated very graciously, very hospitably, and that was true of all my visits over the years.”

Maybe if the victims saw their enemy in person, could put a human face on him, witnessed his genuine anguish and his tears, there could be some hope of forgiveness, or even of reconciliation. But as it is, the dehumanized structure of the multinational corporation works both ways; it makes it easier for individual officers of the corporation to avoid personal liability, and it makes it easier for outsiders to hate an abstract entity, a faceless monolith. Images of Anderson are drawn all over walls in Bhopal; they depict a stick figure with a top hat below the slogans “Hang Anderson” or “Killer Carbide.”

An activist, Satinath Sarangi, once gave the children of the survivors in the slum where he lives pens and paper, and asked them to draw pictures of Anderson. I saw the children’s drawings; most of them are depictions of the devil. But many of the horned figures are smiling and almost endearing, as if the young artists have not quite grasped the nature of evil.

Brian Mooney’s Story

In December 1984, Brian Mooney, one of six children of a Hackensack, New Jersey, shoe salesman, was working at the plush Park Avenue offices of Kelley Drye & Warren, “with people who belonged to country clubs and played squash.” Kelley Drye, one of the oldest and most prestigious law firms in New York, was also Union Carbide’s outside counsel. Mooney at the time was a few months out of law school, so when the Bhopal case broke, he was not one of the senior attorneys there. But the entire firm went into frenzied activity, with people working around the clock on the case. Mooney was put to work on legal research, principally insurance-coverage issues. Every morning that December he would open *The New York Times* and read gruesome accounts of the dead and dying and then take the subway to Park Avenue to put in a full day’s work preparing the defense of the corporation that had done this to them.

Mooney had to rationalize to himself the reasons why he was working for Carbide’s law firm. It was, he says, “a naive belief that people, especially people with suits on, are not capable of malice and wrongdoing, especially on such a large scale.” Also, in this case, the opposing side, in the courts at least, was the Indian government, “not a pristine entity either.”

But gradually Bhopal, and other cases he was working on that were even more untenable personally, dominated his thoughts. Mooney, who is gay and a former Catholic, used to celebrate mass on Saturday evenings at a Greenwich Village church with a gay Catholic group. New York’s Cardinal O’Connor forbade use of the church for the services. So some of the spurned worshipers started going to St. Patrick’s Cathedral on Sunday mornings, where, during the Cardinal’s sermons, they would stand up en masse and silently turn their backs on him.

The archdiocese of New York, through its legal counsel Kelley Drye & Warren, sued the protesters and obtained an injunction against them. A woman at the firm asked Mooney to serve the summons; he still isn’t sure if she knew that he was gay, but he laughed and said absolutely not. She never spoke to him again.

Mooney slowly realized that he had no remaining faith in the legal system, that it had an inefficiency woven into its warp and woof. Mooney quit Kelley Drye in 1988. “At the time, I didn’t have any idea of what I was going to do; I was very good at not thinking about myself because of my being gay.” He was 28, and started to ask himself questions that a 14-year-old would ask, about the purpose, meaning, and direction of his life. He knew that he wanted to help people, but he didn’t believe law was the way to do that.

After a few years of drifting, Mooney applied to graduate school at the University of Michigan at Ann Arbor for the doctoral program in anthropology. He was accepted and given a teaching assignment with a salary equivalent to one-tenth what he used to make at Kelley Drye. In the summer of 1995, Mooney decided to go to Bhopal and study the effects of the legal system on the very people his former employer had commanded him to battle against.

One day that summer, Mooney found himself in the park of the gas-affected women, at one of their Saturday rallies. He was there in his role of anthropologist/observer, ready to note down what the women were doing, why they were here, the structures that they lived within. Suddenly he felt a man tugging at his arm and heard an announcement to the effect that an American visitor would now be making a speech. Mooney was caught off guard, and extremely uncomfortable. "I shouldn't be speaking to them," he thought. "They should be speaking to me." But he found himself, willy-nilly, thrust onto the speaking platform, with a mike in front of him.

Mooney began in his halting Hindi, then switched to English. He told the women that he was studying to be a teacher, and the students that he was teaching at the moment didn't know anything about the rest of the world and they didn't know anything about corporate ethics. These were students who would later go to work for companies like Union Carbide. He was here to gather their stories, he told the women, so he could relate them to his students, so that maybe those students, uniquely powerful because American, would think twice about how the decisions they might make as corporate executives would affect the lives of people half way around the world. That's why he was here, in Bhopal: to gather their stories.

Mooney stopped and looked at the crowd. They applauded politely, smiled, but they didn't really understand. His translator's English was inadequate, and Mooney was left feeling extremely awkward. But, he realized at the same time, he had done something very important for himself: he had just defined his mission, the precise way in which he could help other people. He had been forced to think, and had found an answer to the most universal and least asked of questions: What am I doing here?

A Charge Against Earnings

If there's a happy ending to this story, it's for the Carbide executives and shareholders. Bhopal made the company prey for a takeover attempt a year after the disaster, which forced Carbide to divest itself of its consumer operations and concentrate on its highly profitable core chemical business. In the financial maneuverings that took place during the takeover battle, Carbide gave its shareholders a \$33 bonus dividend plus \$30 a share from the sale of its battery business, and gave its top executives a total of \$28 million in "golden parachutes" to foil future takeover attempts.

Of the \$470 million settlement, \$220 million came out of Carbide's insurance. After news of the settlement, Carbide's stock actually increased \$2 a share. If a person owning a single share of Carbide stock worth \$35 in December 1984 had reinvested all dividends and distribution rights, that share would have been worth more than \$700 a decade later. "Clearly, by any objective measure," says Arthur Sharplin, a management professor who studied these dealings, "Union Carbide Corporation and its managers benefited from the Bhopal incident. It is ironic that a disaster such as Bhopal [would] leave its victims devastated and other corporate stakeholders better off."

Before Bhopal, the worst industrial accident in world history, Union Carbide was involved in the worst industrial tragedy in American history, the death in the 1930s of up to 2000 of its workers due to silicosis during the building of the Hawks Nest Tunnel in West Virginia. Carbide makes no mention of that episode in its corporate histories.

When I went up to the Carbide headquarters in Danbury, Berzok proudly handed me an effusive Paine Webber report on the company, dated September 1995. It says, "We reinstate Carbide as our number one major chemical stock idea." Not once does the name Bhopal come up in the report. ♦

How To Help

The Bhopal Medical Appeal has set up an independent, community based clinic in Bhopal to provide day to day care, drugs, counseling, and physiotherapy. It also monitors the long term effects of the gas on survivors. Contributions or offers of volunteer medical services can be sent to the BMA, c/o Pesticide Action Network/ Bhopal, 116 New Montgomery Street, #810, San Francisco, CA 94105. -S.M.

healthy tissue. Relatively shallow tissue is treated with the accelerated electrons; to reach deeper tissue, the electron beam is converted into X-ray photons.

In the early 1970s, Atomic Energy of Canada Limited (AECL)¹ and a French company called CGR went into business together building linear accelerators. The products of this cooperation were (1) the Therac-6, a 6 million electron volt (MeV) accelerator capable of producing X-rays only and later (2) the Therac-20, a 20 MeV, dual-mode (X-rays or electrons) accelerator. Both were versions of older CGR machines, the Neptune and Sagittaire, respectively, which were augmented with computer control using a DEC PDP-11 minicomputer. We know that some of the old Therac-6 software routines were reused in the Therac-20 and that CGR developed the initial software.

Software functionality was limited in both machines: The computer merely added convenience to the existing hardware, which was capable of standing alone. Industry-standard hardware safety features and interlocks in the underlying machines were retained.

The business relationship between AECL and CGR faltered after the Therac-20 effort. Citing competitive pressures, the two companies did not renew their cooperative agreement when scheduled in 1981.

In the mid-1970s, AECL had developed a radical new “double pass” concept for electron acceleration. A double-pass accelerator needs much less space to develop comparable energy levels because it folds the long physical mechanism required to accelerate the electrons, and it is more economical to produce. Using this double-pass concept, AECL designed the Therac-25, a dual-mode linear accelerator that can deliver either photons at 25 MeV or electrons at various energy levels.

Compared with the Therac-20, the Therac-25 is notably more compact, more versatile, and arguably easier to use. The higher energy takes advantage of the phenomenon of *depth dose*: As the energy increases, the depth in the body at which maximum dose build-up occurs also increases, sparing the tissue above the target area. Economic advantages also come into play for the customer, since only one machine is required for both treatment modalities

¹AECL was an arms-length entity, called a crown corporation, of the Canadian government. Since the time of the incidents related in this paper, AECL Medical, a division of AECL, was privatized and is now called Theratronics International, Ltd. Currently, the primary business of AECL is the design and installation of nuclear reactors.

Medical Devices: The Therac-25*

Nancy Leveson
University of Washington

1 Introduction

Between June 1985 and January 1987, a computer-controlled radiation therapy machine, called the Therac-25, massively overdosed six people. These accidents have been described as the worst in the 35-year history of medical accelerators [6].

A detailed accident investigation, drawn from publicly available documents, can be found in Leveson and Turner [4]. The following account is taken from this report and includes both the factors involved in the overdoses themselves and the attempts by the users, manufacturers, and governments to deal with them. Because this accident was never officially investigated, some information on the Therac-25 software development, management, and quality control procedures are not available. What is included below has been gleaned from law suits and depositions, government records, and copies of correspondence and other material obtained from the U.S. Food and Drug Administration (FDA), which regulates these devices.

2 Background

Medical linear accelerators (linacs) accelerate electrons to create high-energy beams that can destroy tumors with minimal impact on the surrounding

*This appendix is taken from Nancy Leveson, *Software: System Safety and Computers*, Addison-Wesley, 1995. Copyright 1995. All rights reserved.

(electrons and photons).


Several features of the Therac-25 are important in understanding the accidents. First, like the Therac-6 and the Therac-20, the Therac-25 is controlled by a PDP-11 computer. However, AECL designed the Therac-25 to take advantage of computer control from the outset; they did not build on a stand-alone machine. The Therac-6 and Therac-20 had been designed around machines that already had histories of clinical use without computer control.

In addition, the Therac-25 software has more responsibility for maintaining safety than the software in the previous machines. The Therac-20 has independent protective circuits for monitoring the electron-beam scanning plus mechanical interlocks for policing the machine and ensuring safe operation. The Therac-25 relies more on software for these functions. AECL took advantage of the computer's abilities to control and monitor the hardware and decided not to duplicate all the existing hardware safety mechanisms and interlocks.

Some software for the machines was interrelated or reused. In a letter to a Therac-25 user, the AECL quality assurance manager said, "The same Therac-6 package was used by the AECL software people when they started the Therac-25 software. The Therac-20 and Therac-25 software programs were done independently starting from a common base" [4]. The reuse of Therac-6 design features or modules may explain some of the problematic aspects of the Therac-25 software design. The quality assurance manager was apparently unaware that some Therac-20 routines were also used in the Therac-25; this was discovered after a bug related to one of the Therac-25 accidents was found in the Therac-20 software.

AECL produced the first hardwired prototype of the Therac-25 in 1976, and the completely computer-controlled commercial version was available in late 1982.

Figure 1: Upper turntable assembly.



Turntable Positioning. The Therac-25 turntable design plays an important role in the accidents. The upper turntable (see Figure 1) rotates accessory equipment into the beam path to produce two therapeutic modes: electron mode and photon mode. A third position (called the field light position) involves no beam at all, but rather is used to facilitate correct positioning of the patient. Because the accessories appropriate to each mode

PATIENT NAME : TEST	BEAM TYPE: X	ENERGY (MeV): 25
TREATMENT MODE : FX	ACTUAL	PRESCRIBED
	0	200
UNIT RATE/MINUTE	50	50
MONITOR UNITS	0.27	100
TIME (MIN)		
GANTRY ROTATION (DEG)	0.0	0
COLLIMATOR ROTATION (DEG)	359.2	359
COLLIMATOR X (CM)	14.2	14.3
COLLIMATOR Y (CM)	27.2	27.3
WEDGE NUMBER	1	1
ACCESSORY NUMBER	0	0
DATE : 84-OCT-26	SYSTEM : BEAM READY	OP. MODE : TREAT AUTO
TIME : 12:55: 8	TREAT : TREAT PAUSE	X-RAY 173777
OPR ID : T25V02-R03	REASON : OPERATOR	COMMAND:

Figure 2: Operator interface screen layout.

The Operator Interface. The description of the operator interface here applies to the version of the software used during the accidents. Changes made as a result of an FDA recall are described later.

The Therac-25 operator controls the machine through a DEC VT100 terminal. In the general case, the operator positions the patient on the treatment table, manually sets the treatment field sizes and gantry rotation, and attaches accessories to the machine. Leaving the treatment room, the operator returns to the console to enter the patient identification, treatment prescription (including mode or beam type, energy level, dose, dose rate, and time), field sizing, gantry rotation, and accessory data. The system then compares the manually set values with those entered at the console. If they match, a *verified* message is displayed and treatment is permitted. If they do not match, treatment is not allowed to proceed until the mismatch is corrected. Figure 2 shows the screen layout.

When the system was first built, operators complained that it took too

are physically attached to the turntable, proper operation of the Therac-25 is heavily dependent on the turntable position, which is monitored by three microswitches.

The raw, highly concentrated accelerator beam is dangerous to living tissue. In electron therapy, the computer controls the beam energy (from 5 to 25 MeV) and current, while scanning magnets are used to spread the beam to a safe, therapeutic concentration. These scanning magnets are mounted on the turntable and moved into proper position by the computer. Similarly, an ion chamber to measure electrons is mounted on the turntable and also moved into position by the computer. In addition, operator-mounted electron trimmers can be used to shape the beam if necessary.

For X-ray (or photon) therapy, only one energy level is available: 25 MeV. Much greater electron-beam current is required for X-ray mode (some 100 times greater than that for electron therapy) [6] to produce comparable output. Such a high dose-rate capability is required because a “beam flattener” is used to produce a uniform treatment field. This flattener, which resembles an inverted ice cream cone, is a very efficient attenuator; thus, to get a reasonable treatment dose rate out of the flattener, a very high input dose rate is required. If the machine should produce a photon beam with the beam flattener not in position, a high output dose to the patient results. This is the basic hazard of dual-mode machines: If the turntable is in the wrong position, the beam flattener will not be in place.

In the Therac-25, the computer is responsible for positioning the turntable (and for checking the turntable position) so that a target, flattening filter, and X-ray ion chamber are directly in the beam path. With the target in place, electron bombardment produces X-rays. The X-ray beam is shaped by the flattening filter and measured by the X-ray ion chamber.

No accelerator beam is expected in the third or field light turntable position. A stainless steel mirror is placed in the beam path and a light simulates the beam. This lets the operator see precisely where the beam will strike the patient and make necessary adjustments before treatment starts. There is no ion chamber in place at this turntable position, since no beam is expected.

Traditionally, electromechanical interlocks have been used on these types of equipment to ensure safety—in this case, to ensure that the turntable and attached equipment are in the correct position when treatment is started. In the Therac-25, software checks were substituted for many of the traditional hardware interlocks.

“It was not out of the ordinary for something to stop the machine.... It would often give a low dose rate in which you would turn the machine back on. ... They would give messages of low dose rate, V-tilt, H-tilt, and other things; I can't remember all the reasons it would stop, but there was a lot of them.”

A radiation therapist at another clinic reported that an average of 40 dose-rate malfunctions, attributed to underdoses, occurred on some days. The operator further testified that during instruction she had been taught that there were “so many safety mechanisms” that she understood it was virtually impossible to overdose a patient.

Hazard Analysis. In March 1983, AECL performed a safety analysis on the Therac-25. This analysis was in the form of a fault tree and apparently excluded the software. According to the final report, the analysis made several assumptions about the computer and its software:

1. Programming errors have been reduced by extensive testing on a hardware simulator and under field conditions on teletherapy units. Any residual software errors are not included in the analysis.
2. Program software does not degrade due to wear, fatigue, or reproduction process.
3. Computer execution errors are caused by faulty hardware components and by “soft” (random) errors induced by alpha particles and electromagnetic noise.

The fault tree resulting from this analysis does appear to include computer failure, although apparently, judging from the basic assumptions above, it considers hardware failures only. For example, in one OR gate leading to the event of getting the wrong energy, a box contains “Computer selects wrong energy,” and a probability of 10^{-11} is assigned to this event. For “Computer selects wrong mode,” a probability of 4×10^{-9} is given. The report provides no justification of either number.

long to enter the treatment plan. In response, AECL modified the software before the first unit was installed: Instead of reentering the data at the keyboard, operators could simply use a carriage return to copy the treatment site data [5]. A quick series of carriage returns would thus complete the data entry. This modification was to figure in several of the accidents.

The Therac-25 could shut down in two ways after it detected an error condition. One was a *treatment suspend*, which required a complete machine reset to restart. The other, not so serious, was a *treatment pause*, which only required a single key command to restart the machine. If a *treatment pause* occurred, the operator could press the $\text{\textcircled{P}}$ key to “proceed” and resume treatment quickly and conveniently. The previous treatment parameters remained in effect, and no reset was required. This feature could be invoked a maximum of five times before the machine automatically suspended treatment and required the operator to perform a system reset.

Error messages provided to the operator were cryptic, and some merely consisted of the word MALFUNCTION followed by a number from 1 to 64 denoting an analog/digital channel number. According to an FDA memorandum written after one accident:

The operator's manual supplied with the machine does not explain nor even address the malfunction codes. The Maintenance [sic] Manual lists the various malfunction numbers but gives no explanation. The materials provided give no indication that these malfunctions could place a patient at risk.

The program does not advise the operator if a situation exists wherein the ion chambers used to monitor the patient are saturated, thus are beyond the measurement limits of the instrument. This software package does not appear to contain a safety system to prevent parameters being entered and intermixed that would result in excessive radiation being delivered to the patient under treatment.

An operator involved in one of the accidents testified that she had become insensitive to machine malfunctions. Malfunction messages were commonplace and most did not involve patient safety. Service technicians would fix the problems or the hospital physicist would realign the machine and make it operable again. She said,

date, although some rather odd questions had been posed by Tim Still, the hospital physicist.” The physicist at a hospital in Tyler, Texas, where a later accident occurred, reported, “According to Tim Still, the patient filed suit in October 1985 listing the hospital, manufacturer and service organization responsible for the machine. AECL was notified informally about the suit by the hospital, and AECL received official notification of a law suit in November 1985.”

Because of the lawsuit (filed November 13, 1985), some AECL administrators must have known about the Marietta accident—although no investigation occurred at this time. FDA memos point to the lack of a mechanism in AECL to follow up reports of suspected accidents [4].

The patient went home, but shortly afterward she developed a reddening and swelling in the center of the treatment area. Her pain had increased to the point that her shoulder “froze,” and she experienced spasms. She was admitted to a hospital in Atlanta, but her oncologists continued to send her to Kennestone for Therac-25 treatments. Clinical explanation was sought for the reddening of the skin, which at first her oncologist attributed to her disease or to normal treatment reaction.

About two weeks later, the Kennestone physicist noticed that the patient had a matching reddening on her back as though a burn had gone right through her body, and the swollen area had begun to slough off layers of skin. Her shoulder was immobile, and she was apparently in great pain. It was now obvious that she had a radiation burn, but the hospital and her doctors could provide no satisfactory explanation.

The Kennestone physicist later estimated that the patient received one or two doses of radiation in the 15,000 to 20,000 rad (radiation absorbed dose) range. He did not believe her injury could have been caused by less than 8,000 rads. To understand the magnitude of this, consider that typical single therapeutic doses are in the 200 rad range. Doses of 1,000 rads can be fatal if delivered to the whole body; in fact, 500 rads is the accepted figure for whole-body radiation that will cause death in 50 percent of the cases. The consequences of an overdose to a smaller part of the body depend on the tissue’s radio-sensitivity. The director of radiation oncology at the Kennestone facility explained their confusion about the accident as due to the fact that they had never seen an overtreatment of that magnitude before [7].

Eventually, the patient’s breast had to be removed because of the radiation burns. Her shoulder and arm were paralyzed, and she was in constant

3 Events

Eleven Therac-25s were installed: five in the United States and six in Canada. Six accidents occurred between 1985 and 1987, when the machine was finally recalled to make extensive design changes. These changes include adding hardware safeguards against software errors.

Related problems were found in the Therac-20 software, but they were not recognized until after the Therac-25 accidents because the Therac-20 includes hardware safety interlocks. Thus, no injuries resulted.

3.1 Kennestone Regional Oncology Center, June 1985

Details of this accident in Marietta, Georgia, are sketchy because it was never investigated. There was no admission that the injury was caused by the Therac-25 until long after the occurrence, despite claims by the patient that she had been injured during treatment, the obvious and severe radiation burns the patient suffered, and the suspicions of the radiation physicist involved.

After undergoing a lumpectomy to remove a malignant breast tumor, a 61-year-old woman was receiving follow-up radiation treatment to nearby lymph nodes on a Therac-25 at the Kennestone facility in Marietta. The Therac-25 had been operating at Kennestone for about six months; other Therac-25s had been operating, apparently without incident, since 1983.

On June 3, 1985, the patient was set up for a 10 MeV electron treatment to the clavicle area. When the machine turned on, she felt a “tremendous force of heat...this red-hot sensation.” When the technician came in, she said, “You burned me.” The technician replied that that was impossible. Although there were no marks on the patient at the time, the treatment area felt “warm to the touch.”

It is unclear exactly when AECL learned about this incident. Tim Still, the Kennestone physicist, said that he contacted AECL to ask if the Therac-25 could operate in electron mode without scanning to spread the beam. Three days later the engineers at AECL called the physicist back to explain that improper scanning was not possible.

In an August 19, 1986 letter from AECL to the FDA, the AECL quality assurance manager said, “In March of 1986 AECL received a lawsuit from the patient involved... This incident was never reported to AECL prior to this

pecting the machine to deliver the proper dose this time. This was standard operating procedure, and Therac-25 operators had become accustomed to frequent malfunctions that had no untoward consequences for the patient. Again, the machine shut down in the same manner. The operator repeated this process four times after the original attempt—the display showing NO DOSE delivered to the patient each time. After the fifth pause, the machine went into treatment suspend, and a hospital service technician was called. The technician found nothing wrong with the machine. According to a Therac-25 operator, this scenario also was not unusual.

After the treatment, the patient complained of a burning sensation, described as an “electric tingling shock” to the treatment area in her hip. Six other patients were treated later that day without incident. She came back for further treatment on July 29 and complained of burning, hip pain, and excessive swelling in the region of treatment. The patient was hospitalized for the condition on July 30, and the machine was taken out of service.

AECL was informed of the apparent radiation injury and sent a service engineer to investigate. The U.S. FDA, the then Canadian Radiation Protection Bureau (RPB),² and users were informed that there was a problem, although the users claim that they were never informed that a patient injury had occurred. Users were told that they should visually confirm the proper turntable alignment until further notice (which occurred three months later).

The patient died on November 3, 1985, of an extremely virulent cancer. An autopsy revealed the cause of death as the cancer, but it was noted that had she not died, a total hip replacement would have been necessary as a result of the radiation overexposure. An AECL technician later estimated the patient had received between 13,000 and 17,000 rads.

3.2.1 Manufacturer’s Response

AECL could not reproduce the malfunction that had occurred, but suspected a transient failure in the microswitch used to determine the turntable position. During the investigation of the accident, AECL hardwired the error conditions they assumed were necessary for the malfunction and, as a result, found some turntable positioning design weaknesses and potential mechanical problems.

²On April 1, 1986, the Radiation Protection Bureau and the Bureau of Medical Devices were merged to form the Bureau of Radiation and Medical Devices (BRMD).

pain. She had suffered a serious radiation burn, but the manufacturer and operators of the machine refused to believe that it could have been caused by the Therac-25. The treatment prescription printout feature of the computer was disabled at the time of the accident, so there was no hardcopy of the treatment data. The lawsuit was eventually settled out of court.

From what we can determine, the accident was not reported to the FDA until *after* further accidents in 1986. The reporting requirements for medical device incidents at that time applied only to equipment manufacturers and importers, not users. The regulations required that manufacturers and importers report deaths, serious injuries, or malfunctions that could result in those consequences, but health-care professionals and institutions were not required to report incidents to manufacturers. The comptroller general of the U.S. Government Accounting Office (GAO), in testimony before Congress on November 6, 1989, expressed great concern about the viability of the incident-reporting regulations in preventing or spotting medical device problems. According to a 1990 GAO study, the FDA knew of less than 1 percent of deaths, serious injuries, or equipment malfunctions that occurred in hospitals [2]. The law was amended in 1990 to require health-care facilities to report incidents to the manufacturer and to the FDA.

At this point, the other Therac-25 users were also unaware that anything untoward had occurred and did not learn about any problems with the machine until after subsequent accidents. Even then, most of their information came through personal communication among themselves.

3.2 Ontario Cancer Foundation, July 1985

The second in this series of accidents occurred about seven weeks after the Kenestone patient was overdosed. At that time, the Therac-25 at the Ontario Cancer Foundation in Hamilton, Ontario (Canada), had been in use for more than six months. On July 26, 1985, a forty-year-old patient came to the clinic for her twenty-fourth Therac-25 treatment for carcinoma of the cervix. The operator activated the machine, but the Therac shut down after five seconds with an **FAULT** error message. The Therac-25’s console display read **NO DOSE** and indicated a **TREATMENT PAUSE**.

Since the machine did not suspend and the control display indicated no dose was delivered to the patient, the operator went ahead with a second attempt at treatment by pressing the **P** key (the *proceed* command), ex-

The computer senses and controls turntable position by reading a 3-bit signal about the status of three microswitches in the turntable switch assembly. Essentially, AECL determined that a 1-bit error in the microswitch codes (which could be caused by a single open-circuit fault on the switch lines) could produce an ambiguous position message to the computer. The problem was exacerbated by the design of the mechanism that extends a plunger to lock the turntable when it is in one of the three cardinal positions: The plunger could be extended when the turntable was way out of position, thus giving a second false position indication. AECL devised a method to indicate turntable position that tolerated a 1-bit error so that the code would still unambiguously reveal correct position with any one microswitch failure. In addition, AECL altered the software so that the computer checked for “in transit” status of the switches to keep further track of the switch operation and turntable position and to give additional assurance that the switches were working and the turntable was moving.

As a result of these improvements, AECL claimed in its report and correspondence with hospitals that “analysis of the hazard rate of the new solution indicates an improvement over the old system by at least 5 orders of magnitude [emphasis added].” However, in its final incident report to the FDA, AECL concluded that they “cannot be firm on the exact cause of the accident but can only suspect . . . ,” which underscored their inability to determine the cause of the accident with any certainty. The AECL quality assurance manager testified that they could not reproduce the switch malfunction and that testing of the microswitch was “inconclusive.” The similarity of the errant behavior and the patient injuries in this accident and a later one in Yakima, Washington, provide good reason to believe that the Hamilton overdose was probably related to software error rather than to a microswitch failure.

3.2.2 Government and User Response

The Hamilton accident resulted in a voluntary recall by AECL, and the FDA termed it a Class II recall. Class II means “a situation in which the use of, or exposure to, a violative product may cause temporary or medically reversible adverse health consequences or where the probability of serious adverse health consequences is remote.” The FDA audited AECL’s subsequent modifications, and after the modifications were made, the users were told they could return to normal operating procedures.

As a result of the Hamilton accident, the head of advanced X-ray systems in the Canadian RPB, Gordon Symonds, wrote a report that analyzed the design and performance characteristics of the Therac-25 with respect to radiation safety. Besides citing the flawed microswitch, the report faulted both hardware and software components of the Therac’s design. It concluded with a list of four modifications to the Therac-25 necessary for compliance with Canada’s Radiation Emitting Devices (RED) Act. The RED law, enacted in 1971, gives government officials power to ensure the safety of radiation-emitting devices.

The modifications specified in the Symonds report included redesigning the microswitch and changing the way the computer handled malfunction conditions. In particular, treatment was to be terminated in the event of a dose-rate malfunction, giving a treatment “suspend.” This change would have removed the option to proceed simply by pressing the ⊕ key. The report also made recommendations regarding collimator test procedures and message and command formats. A November 8, 1985 letter, signed by the director of the Canadian RPB, asked that AECL make changes to the Therac-25 based on the Symonds report “to be in compliance with the RED act.”

Although, as noted above, AECL did make the microswitch changes, they did not comply with the directive to change the malfunction pause behavior into treatment suspends, instead reducing the maximum number of retries from five to three. According to Symonds, the deficiencies outlined in the RPB letter of November 8 were still pending when the next accident happened five months later.

Immediately after the Hamilton accident, the Ontario Cancer Foundation hired an independent consultant to investigate. He concluded in a September 1985 report that an independent system (beside the computer) was needed to verify the turntable position and suggested the use of a potentiometer. The RPB wrote a letter to AECL in November 1985 requesting that AECL install such an independent interlock on the Therac-25. Also, in January 1986, AECL received a letter from the attorney representing the Hamilton clinic. The letter said that there had been continuing problems with the turntable, including four incidents at Hamilton, and requested the installation of an independent system (potentiometer) to verify the turntable position. AECL did not comply: No independent interlock was installed by AECL on the Therac-25s at this time. The Hamilton Clinic, however, decided to install one themselves on their machine.

3.3 Yakima Valley Memorial Hospital, December 1985

In this accident, as in the Kennestone overdose, machine malfunction was not acknowledged until after later accidents were understood.

The Therac-25 at Yakima, Washington, had been modified by AECL in September 1985 in response to the overdose at Hamilton. During December 1985, a woman treated with the Therac-25 developed erythema (excessive reddening of the skin) in a parallel striped pattern on her right hip. Despite this, she continued to be treated by the Therac-25, as the cause of her reaction was not determined to be abnormal until January 1986. On January 6, her treatments were completed.

The staff monitored the skin reaction closely and attempted to find possible causes. The open slots in the blocking trays in the Therac-25 could have produced such a striped pattern, but by the time the skin reaction was determined to be abnormal, the blocking trays had been discarded, so the blocking arrangement and tray striping orientation could not be reproduced. A reaction to chemotherapy was ruled out because that should have produced reactions at the other treatment sites and would not have produced stripes. When the doctors discovered that the woman slept with a heating pad, they thought maybe the burn pattern had been caused by the parallel wires that deliver the heat in such pads. The staff X-rayed the heating pad but discovered that the wire pattern did not correspond to the erythema pattern on the patient's hip.

The hospital staff sent a letter to AECL on January 31, and they also spoke on the phone with the AECL technical support supervisor. On February 24, the AECL technical support supervisor sent a written response to the director of radiation therapy at Yakima saying, "After careful consideration we are of the opinion that this damage could not have been produced by any malfunction of the Therac-25 or by any operator error." The letter goes on to support this opinion by listing two pages of technical reasons why an overdose by the Therac-25 was impossible, along with the additional argument that there have "apparently been no other instances of similar damage to this or other patients." The letter ends, "In closing, I wish to advise that is normal practice."

The hospital staff eventually ascribed the patient's skin reaction to "cause unknown." In a report written on this first Yakima incident after another

Yakima overdose a year later, the medical physicist involved wrote:

At that time, we did not believe that [the patient] was overdosed because the manufacturer had installed additional hardware and software safety devices to the accelerator.

In a letter from the manufacturer dated 16-Sep-85, it is stated that "Analysis of the hazard rate resulting from these modifications indicates an improvement of at least five orders of magnitude"! With such an improvement in safety (10,000,000%) we did not believe that there could have been any accelerator malfunction. These modifications to the accelerator were completed on 5.6-Sep-85.

Even with fairly sophisticated physics support, the hospital staff, as users, did not have the ability to investigate the possibility of machine malfunction further. They were not aware of any other incidents and, in fact, were told that there had been none, so there was no reason for them to pursue the matter. No further investigation of this incident was done by the manufacturer or by any government agencies (who did not know about it).

About a year later (February 1987), after the second Yakima overdose led the hospital staff to suspect that this first injury had been due to a Therac-25 fault, the staff investigated and found that the first overdose victim had a chronic skin ulcer, tissue necrosis (death) under the skin, and was in continual pain. The damage was surgically repaired, skin grafts were made, and the symptoms relieved. The patient is alive today with minor disability and some scarring related to the overdose. The hospital staff concluded that the dose accidentally delivered in the first accident must have been much lower than in the second, as the reaction was significantly less intense and necrosis did not develop until six or eight months after exposure. Some other factors related to the place on the body where the overdose occurred also kept her from having more significant problems.

3.4 East Texas Cancer Center, March 1986

More is known about the Tyler, Texas, accidents than the others because of the diligence of the Tyler hospital physicist, Fritz Hager, without whose efforts the understanding of the software problems may have been delayed even further.

The Therac-25 had been at the East Texas Cancer Center (ETCC) for two years before the first serious accident, and more than 500 patients had been treated. On March 21, 1986, a male patient came into ETCC for his ninth treatment on the Therac-25, one of a series prescribed as followup to the removal of a tumor from his back.

This treatment was to be a 22 MeV electron beam treatment of 180 rads on the upper back and a little to the left of his spine, for a total of 6,000 rads over six and a half weeks. He was taken into the treatment room and placed face down on the treatment table. The operator then left the treatment room, closed the door, and sat at the control terminal.

The operator had held this job for some time, and her typing efficiency had increased with experience. She could quickly enter prescription data and change it conveniently with the Therac’s editing features. She entered the patient’s prescription data quickly, then noticed that she had typed “x” (for X-ray) when she had intended “e” (for electron) mode. This was a common mistake as most of the treatments involved X-rays, and she had gotten used to typing this. The mistake was easy to fix; she merely used the $\text{\textcircled{1}}$ key to edit the mode entry.

Because the other parameters she had entered were correct, she hit the return key several times and left their values unchanged. She reached the bottom of the screen, where it was indicated that the parameters had been VERIFIED and the terminal displayed BEAM READY, as expected. She hit the one-key command, $\text{\textcircled{B}}$ for *beam on*, to begin the treatment. After a moment, the machine shut down and the console displayed the message MALFUNCTION 54. The machine also displayed a TREATMENT PAUSE, indicating a problem of low priority. The sheet on the side of the machine explained that this malfunction was a “dose input 2” error. The ETCC did not have any other information available in its instruction manual or other Therac-25 documentation to explain the meaning of MALFUNCTION 54. An AECL technician later testified that “dose input 2” meant that a dose had been delivered that was either too high or too low. The messages had been expected to be used only during internal company development.

The machine showed a substantial underdose on its dose monitor display—6 monitor units delivered whereas the operator had requested 202 monitor units. She was accustomed to the quirks of the machine, which would frequently stop or delay treatment; in the past, the only consequences had been inconvenience. She immediately took the normal action when the machine

merely paused, which was to hit the $\text{\textcircled{P}}$ key to proceed with the treatment. The machine promptly shut down with the same MALFUNCTION 54 error and the same underdose shown by the dosimetry.

The operator was isolated from the patient, since the machine apparatus was inside a shielded room of its own. The only way that the operator could be alerted to patient difficulty was through audio and video monitors. On this day, the video display was unplugged and the audio monitor was broken.

After the first attempt to treat him, the patient said that he felt, as if he had received an electric shock or that someone had poured hot coffee on his back. He felt a thump and heat and heard a buzzing sound from the equipment. Since this was his ninth treatment, he knew that this was not normal. He began to get up from the treatment table to go for help. It was at this moment that the operator hit the $\text{\textcircled{P}}$ key to proceed with the treatment. The patient said that he felt like his arm was being shocked by electricity and that his hand was leaving his body. He went to the treatment room door and pounded on it. The operator was shocked and immediately opened the door for him. He appeared visibly shaken and upset.

The patient was immediately examined by a physician, who observed intense reddening of the treatment area, but suspected nothing more serious than electric shock. The patient was discharged and sent home with instructions to return if he suffered any further reactions. The hospital physicist was called in, and he found the machine calibration within specifications. The meaning of the malfunction message was not understood. The machine was then used to treat patients for the rest of the day.

In actuality, but unknown to anyone at that time, the patient had received a massive overdose, concentrated in the center of the treatment location. After-the-fact simulations of the accident revealed possible doses of 16,500 to 25,000 rads in less than 1 second over an area of about 1 cm.

Over the weeks following the accident, the patient continued to have pain in his neck and shoulder. He lost the function of his left arm and had periodic bouts of nausea and vomiting. He was eventually hospitalized for radiation-induced myelitis of the cervical cord causing paralysis of his left arm and both legs, left vocal cord paralysis (which left him unable to speak), neurogenic bowel and bladder, and paralysis of the left diaphragm. He also had a lesion on his left lung and recurrent herpes simplex skin infections. He died from complications of the overdose five months after the accident.

again. The operator rushed into the treatment room, hearing her patient moaning for help. He began to remove the tape that had held his head in position and said something was wrong. She asked him what he felt, and he replied, “fire” on the side of his face. She immediately went to the hospital physicist and told him that another patient appeared to have been burned. Asked by the physicist to describe what had happened, the patient explained that something had hit him on the side of the face, he saw a flash of light, and he heard a sizzling sound reminiscent of frying eggs. He was very agitated and asked, “What happened to me, what happened to me?”

This patient died from the overdose on May 1, 1986, three weeks after the accident. He had disorientation, which progressed to coma, fever to 104°F, and neurological damage. An autopsy showed an acute high-dose radiation injury to the right temporal lobe of the brain and the brain stem.

3.5.1 User and Manufacturer Response

After this second Tyler accident, the ETCC physicist immediately took the machine out of service and called AECL to alert them to this second apparent overexposure. The physicist then began a careful investigation of his own. He worked with the operator, who remembered exactly what she had done on this occasion. After a great deal of effort, they were eventually able to elicit the MALFUNCTION 54 message. They determined that data entry speed during editing was the key factor in producing the error condition: If the prescription data was edited at a fast pace (as is natural for someone who has repeated the procedure a large number of times), the overdose occurred. It took some practice before the physicist could repeat the procedure rapidly enough to elicit the MALFUNCTION 54 message at will.

The next day, an engineer from AECL called and said that he could not reproduce the error. After the ETCC physicist explained that the procedure had to be performed quite rapidly, AECL could finally produce a similar malfunction on its own machine. Two days after the accident, AECL said it had measured the dosage (at the center of the field) to be 25,000 rads. An AECL engineer explained that the frying sound heard by the patients was the ion chambers being saturated.

In one law suit that resulted from the Tyler accidents, the AECL quality control manager testified that a “cursor up” problem had been found in the service (maintenance) mode at other clinics in February or March of 1985 and

3.4.1 User and Manufacturer Response

The Therac-25 was shut down for testing the day after this accident. One local AECL engineer and one from the home office in Canada came to ETCC to investigate. They spent a day running the machine through tests, but could not reproduce a Malfunction 54. The AECL engineer from the home office reportedly explained that it was not possible for the Therac-25 to overdose a patient. The ETCC physicist claims that he asked AECL at this time if there were any other reports of radiation overexposure and that AECL personnel (including the quality assurance manager) told him that AECL knew of no accidents involving radiation overexposure by the Therac-25. This seems odd since AECL was surely at least aware of the Hamilton accident that had occurred seven months before and the Yakima accident, and, even by their account, learned of the Georgia law suit around this time (which had been filed four months earlier). The AECL engineers then suggested that an electrical problem might have caused the problem.

The electric shock theory was checked out thoroughly by an independent engineering firm. The final report indicated that there was no electrical grounding problem in the machine, and it did not appear capable of giving a patient an electrical shock. The ETCC physicist checked the calibration of the Therac-25 and found it to be satisfactory. He put the machine back into service on April 7, 1986, convinced that it was performing properly.

3.5 East Texas Cancer Center, April 1986

Three weeks later, on April 11, 1986, another male patient was scheduled to receive an electron treatment at ETCC for a skin cancer on the side of his face. The prescription was for 10 MeV. The same technician who had treated the first Tyler accident victim prepared this patient for treatment. Much of what follows is from the operator’s deposition.

As with her former patient, she entered the prescription data and then noticed an error in the mode. Again she used the edit \ominus key to change the mode from X-ray to electron. After she finished editing, she pressed the RETURN key several times to place the cursor on the bottom of the screen. She saw the BEAM READY message displayed and turned the beam on.

Within a few seconds the machine shut down, making a loud noise audible via the (now working) intercom. The display showed MALFUNCTION 54

also in the summer of 1985. Both times, AECL thought that the software problems had been fixed. There is no way to determine whether there is any relationship between these problems and the Tyler accidents.

3.5.2 Related Therac-20 Problems

The software for both the Therac-25 and Therac-20 “evolved” from the Therac-6 software. Additional functions had to be added because the Therac-20 (and Therac-25) operate in both X-ray and electron mode, while the Therac-6 has only X-ray mode. CGR modified the software for the Therac-20 to handle the dual modes. When the Therac-25 development began, AECL engineers adapted the software from the Therac-6, but they also borrowed software routines from the Therac-20 to handle electron mode, which was allowed under their cooperative agreements.

After the second Tyler, Texas, accident, a physicist at the University of Chicago Joint Center for Radiation Therapy heard about the Therac-25 software problem and decided to find out whether the same thing could happen with the Therac-20. At first, the physicist was unable to reproduce the error on his machine, but two months later he found the link.

The Therac-20 at the University of Chicago is used to teach students in a radiation therapy school conducted by the center. The center’s physicist, Frank Borger, noticed that whenever a new class of students started using the Therac-20, fuses and breakers on the machine tripped, shutting down the unit. These failures, which had been occurring ever since the school had acquired the machine, might happen three times a week while new students operated the machine and then disappear for months. Borger determined that new students make many different types of mistakes and use “creative methods” of editing parameters on the console. Through experimentation, he found that certain editing sequences correlated with blown fuses and determined that the same computer bug (as in the Therac-25 software) was responsible. The physicist notified the FDA, which notified Therac-20 users [3].

The software error is just a nuisance on the Therac-20 because this machine has independent hardware protective circuits for monitoring the electron beam scanning. The protective circuits do not allow the beam to turn on, so there is no danger of radiation exposure to a patient. While the Therac-20 relies on mechanical interlocks for monitoring the machine, the

Therac-25 relies largely on software.

3.5.3 The Software “Bug”

A lesson to be learned from the Therac-25 story is that focusing on particular software “bugs” is not the way to make a safe system. Virtually all complex software can be made to behave in an unexpected fashion under some conditions. The basic mistakes here involved poor software engineering practices and building a machine that relies on the software for safe operation. Furthermore, the particular coding error is not as important as the general unsafe design of the software overall. Examining the part of the code blamed for the Tyler accidents is instructive, however, in demonstrating the overall software design flaws. First the software design is described and then the errors believed to be involved in the Tyler accidents and perhaps others.

Therac-25 Software Development and Design. AECL claims proprietary rights to its software design. However, from voluminous documentation regarding the accidents, the repairs, and the eventual design changes, we can build a rough picture of it.

The software is responsible for monitoring the machine status, accepting input about the treatment desired, and setting the machine up for this treatment. It turns the beam on in response to an operator command (assuming that certain operational checks on the status of the physical machine are satisfied) and also turns the beam off when treatment is completed, when an operator commands it, or when a malfunction is detected. The operator can print out hardcopy versions of the CRT display or machine setup parameters.

The treatment unit has an interlock system designed to remove power to the unit when there is a hardware malfunction. The computer monitors this interlock system and provides diagnostic messages. Depending on the fault, the computer either prevents a treatment from being started or, if the treatment is in progress, creates a pause or a suspension of the treatment.

There are two basic operational modes: treatment mode and service mode. Treatment mode controls the normal treatment process. In service mode, the unit can be operated with some of the operational and treatment interlocks bypassed, and additional operational commands and characteristics may be selected. Service mode is entered only through the use of a password at the service keyboard.

- The servo task controls gun emission, dose rate (pulse repetition frequency), symmetry (beam steering), and machine motions. The servo task also sets up the machine parameters and monitors the beam-tilt error and the flatness-error interlocks.
- The housekeeper task takes care of system status interlocks and limit checks and displays appropriate messages on the CRT display. It decodes some information and checks the setup verification.

Noncritical tasks include

- Checksum processor (scheduled to run periodically)
- Treatment console keyboard processor (scheduled to run only if it is called by other tasks or by keyboard interrupts). This task acts as the communication interface between the other software and the operator.
- Treatment console screen processor (run periodically). This task lays out appropriate record formats for either CRT displays or hard copies.
- Service keyboard processor (run on demand). This task arbitrates non-treatment-related communication between the therapy system and the operator.
- Snapshot (run periodically by the scheduler). Snapshot captures pre-selected parameter values and is called by the treatment task at the end of a treatment.
- Hand control processor (run periodically).
- Calibration processor. This task is responsible for a package of tasks that let the operator examine and change system setup parameters and interlock limits.

It is clear from the AECL documentation on the modifications that the software allows concurrent access to shared memory, that there is no real synchronization aside from data that are stored in shared variables, and that the “test” and “set” for such variables are not indivisible operations. Race conditions resulting from this implementation of multitasking played an important part in the accidents.

Specific Design Errors. The following explanation of the specific software problems found at this time is taken from the description AECL provided to the FDA, but clarified somewhat. The description leaves some unanswered questions, but it is the best that can be done with the information available.

The manufacturer describes the Therac-25 software as having a stand-alone, real-time treatment operating system. The system does not use a standard operating system or executive. Rather, the real-time executive was written especially for the Therac-25 and runs on a 32K PDP-11/23. Cycles are allocated to the critical and noncritical tasks using a preemptive scheduler.

The software, written in PDP-11 assembly language, has four major components: stored data, a scheduler, a set of critical and noncritical tasks, and interrupt services. The stored data includes calibration parameters for the accelerator setup as well as patient-treatment data. The interrupt routines include

- A clock interrupt service routine
- A scanning interrupt service routine
- Traps (for software overflow and computer hardware generated interrupts)
- Power up (initiated at power up to initialize the system and pass control to the scheduler)
- Treatment console screen interrupt handler
- Treatment console keyboard interrupt handler
- Service printer interrupt handler
- Service keyboard interrupt handler

The scheduler controls the sequencing of all noninterrupt events and coordinates all concurrent processes. Tasks are initiated every 0.1 second, with the critical tasks executed first and the noncritical tasks executed in any remaining cycle time. Critical tasks include the following:

- The treatment monitor (Treat) directs and monitors patient setup and treatment via eight operating phases. These are called as subroutines, depending on the value of the Tphase control variable. Following the execution of a particular subroutine, Treat reschedules itself. Treat interacts with the keyboard processing task, which handles operator console communication. The prescription data is cross-checked and verified by other tasks (such as keyboard processor or parameter setup sensor) that inform the treatment task of the verification status via shared variables.

mainline. Treat will then reschedule itself, essentially rescheduling the Datent subroutine.

The command line at the lower right-hand corner of the screen (see Figure 2) is the cursor's normal position when the operator has completed all the necessary changes to the prescription. Prescription editing is signified by moving the cursor off the command line. As the program was originally designed, the Data Entry Complete variable by itself is not sufficient because it does not ensure that the cursor is located on the command line; under the right circumstances, the data entry phase can be exited before all edit changes are made on the screen.

The keyboard handler parses the mode and energy level specified by the operator and places an encoded result in another shared variable, the 2-byte Mode/Energy Offset variable (MEOS). The low-order byte of this variable is used by another task (Hand) to set the collimator/turntable to the proper position for the selected mode and energy. The high-order byte of the MEOS variable is used by Datent to set several operating parameters.

Initially, the data-entry process forces the operator to enter the mode and energy except when the photon mode is selected, in which case the energy defaults to 25 MeV. The operator can later edit the mode and energy separately. If the keyboard handler sets the Data Entry Complete flag before the operator changes the data in MEOS, Datent will not detect the changes because it has already exited and will not be reentered again. The upper collimator (turntable), on the other hand, is set to the position dictated by the low-order byte of MEOS by another concurrently running task (Hand) and can therefore be inconsistent with the parameters set in accordance with the information in the high-order byte. The software appears to contain no checks to detect such an incompatibility.

The first thing Datent does when it is entered is to check whether the keyboard handler has set the mode and energy in MEOS. If so, it uses the high-order byte to index into a table of preset operating parameters and places them in the digital-to-analog output table. The contents of this output table are transferred to the digital-to-analog converter during the next clock cycle. Once the parameters are all set, Datent calls the subroutine Magnet, which sets the bending magnets. The following shows a simplified pseudocode description of relevant parts of the software:

Datent:

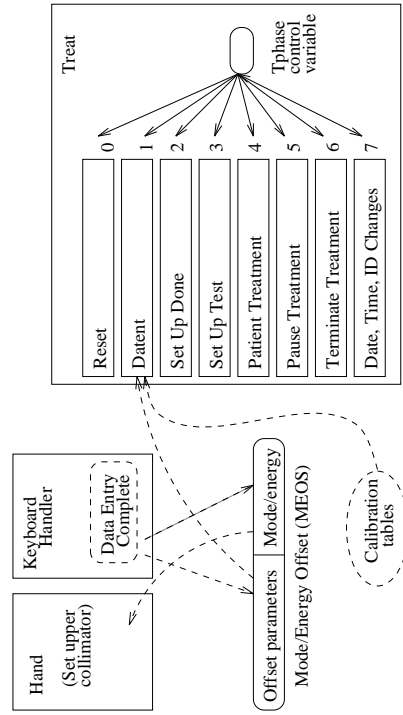


Figure 3: Tasks and subroutines in the code blamed for the Tyler accidents.

The treatment monitor task (Treat) controls the various phases of treatment by executing its eight subroutines. The treatment phase indicator variable (Tphase) is used to determine which subroutine should be executed (Figure 3). Following the execution of a particular subroutine, Treat reschedules itself.

One of Treat's subroutines, called Datent (data entry), communicates with the keyboard handler task (a task that runs concurrently with Treat) via a shared variable (Data Entry Complete flag) to determine whether the prescription data has been entered. The keyboard handler recognizes the completion of data entry and changes the Data Entry Complete variable to denote this. Once this variable is set, the Datent subroutine detects the variable's change in status and changes the value of Tphase from 1 (Datent) to 3 (Set Up Test). In this case, the Datent subroutine exits back to the Treat subroutine, which will reschedule itself and begin execution of the Set Up Test subroutine. If the Data Entry Complete variable has not been set, Datent leaves the value of Tphase unchanged and exits back to Treat's

```

if mode/energy specified then
  begin
    calculate table index
  repeat
    fetch parameter
    output parameter
    point to next parameter
  until all parameters set
  call Magnet
  if mode/energy changed then return
  end
if data entry is complete then set Tphase to 3
if data entry is not complete then
  if reset command entered then set Tphase to 0
return

```

Magnet:

```

Set bending magnet flag
repeat
  Set next magnet
  call Ptime
  if mode/energy has changed, then exit
until all magnets are set
return

```

Ptime:

```

repeat
  if bending magnet flag is set then
    if editing taking place then
      if mode/energy has changed then exit
    hysteresis delay has expired
  Clear bending magnet flag
return

```

Setting the bending magnets takes about eight seconds. Magnet calls a subroutine called Ptime to introduce a time delay. Since several magnets need

to be set, Ptime is entered and exited several times. A flag to indicate that the bending magnets are being set is initialized upon entry to the Magnet subroutine and cleared at the end of Ptime. Furthermore, Ptime checks a shared variable, set by the keyboard handler, that indicates the presence of any editing requests. If there are edits, then Ptime clears the bending magnet variable and exits to Magnet, which then exits to Datent. But the edit change variable is checked by Ptime only if the bending magnet flag is set. Because Ptime clears it during its first execution, any edits performed during each succeeding pass through Ptime will not be recognized. Thus, an edit change of the mode or energy, although reflected on the operator's screen and the mode/energy offset variable, will not be sensed by Datent so it can index the appropriate calibration tables for the machine parameters.

Recall that the Tyler error occurred when the operator made an entry indicating the mode and energy, went to the command line, then moved the cursor up to change the mode or energy and returned to the command line all within eight seconds. Because the magnet setting takes about eight seconds and Magnet does not recognize edits after the first execution of Ptime, the editing had been completed by the return to Datent, which never detected that it had occurred. Part of the problem was fixed after the accident by clearing the bending magnet variable at the end of Magnet (after *all* the magnets have been set) instead of at the end of Ptime.

But this is not the only problem. Upon exit from the Magnet subroutine, the data entry subroutine (Datent) checks the Data Entry Complete variable. If it indicates that data entry is complete, Datent sets Tphase to 3 and Datent is not entered again. If it is not set, Datent leaves Tphase unchanged, which means it will eventually be rescheduled. But the Data Entry Complete variable only indicates that the cursor has been down to the command line, not that it is still there. A potential race condition is set up. To fix this, AECL introduced another shared variable controlled by the keyboard handler task that indicates the cursor is not positioned on the command line. If this variable is set, then prescription entry is still in progress and the value of Tphase is left unchanged.

3.5.4 The Government and User Response

The FDA does not approve each new medical device on the market. All medical devices go through a classification process that determines the level

To avoid accidental use of this key, the key cap must be removed and the switch contacts fixed in the open position with electrical tape or other insulating material. For assistance with the latter you should contact your local AECL service representative.

Disabling this key means that if any prescription data entered is incorrect then a 'R' reset command must be used and the whole prescription reentered.

For those users of the Multipoint option it also means that editing of dose rate, dose and time will not be possible between ports.

On May 2, 1986, the FDA declared the Therac defective, demanded a CAP, and required renotification of all the Therac customers. In the letter from the FDA to AECL, the Director of Compliance, Center for Devices and Radiological Health, wrote:

We have reviewed [AECL's] April 15 letter to purchasers and have concluded that it does not satisfy the requirements for notification to purchasers of a defect in an electronic product. Specifically, it does not describe the defect nor the hazards associated with it. The letter does not provide any reason for disabling the cursor key and the tone is not commensurate with the urgency for doing so. In fact, the letter implies the inconvenience to operators outweighs the need to disable the key. We request that you immediately renotify purchasers.

AECL promptly made a new notice to users and also requested an extension to produce a CAP. The FDA granted this request.

About this time, the Therac-25 users created a user's group and held their first meeting at the annual conference of the American Association of Physicists in Medicine. At the meeting, users discussed the Tyler accident and heard an AECL representative present the company's plans for responding to it. AECL promised to send a letter to all users detailing the CAP.

Several users described additional hardware safety features that they had added to their own machines to provide additional protection. An interlock (that checked gun current values), which the Vancouver clinic had previously added to their Therac-25, was labeled as redundant by AECL; the users

of FDA approval necessary. Medical accelerators follow a procedure called pre-market notification before commercial distribution. In this process, the firm must establish that the product is substantially equivalent in safety and effectiveness to a product already on the market. If that cannot be done to the FDA's satisfaction, a pre-market approval is required. For the Therac-25, the FDA required only a pre-market notification. After the Therac-25 accidents, new procedures for approval of software-controlled devices were adopted.

The agency is basically reactive to problems and requires manufacturers to report serious ones. Once a problem is identified in a radiation-emitting product, the FDA is responsible for approving the corrective action plan (CAP).

The first reports of the Tyler incidents came to the FDA from the State of Texas Health Department, and this triggered FDA action. The FDA investigation was well under way when AECL produced a medical device report to discuss the details of the radiation overexposures at Tyler. The FDA declared the Therac-25 defective under the Radiation Control for Health and Safety Act and ordered the firm to notify all purchasers, investigate the problem, determine a solution, and submit a corrective action plan for FDA approval.

The final CAP consisted of more than twenty changes to the system hardware and software, plus modifications to the system documentation and manuals. Some of these changes were unrelated to the specific accidents, but were improvements to the general safety of the machine. The full CAP implementation, including an extensive safety analysis, was not complete until more than two years after the Tyler accidents.

AECL made their accident report to the FDA on April 15, 1986. On that same date, AECL sent out a letter to each Therac user recommending a temporary "fix" to the machine that would allow continued clinical use. The letter (shown in its complete form) stated:

SUBJECT: CHANGE IN OPERATING PROCEDURES FOR THE THERAC 25 LINEAR ACCELERATOR

Effective immediately, and until further notice, the key used for moving the cursor back through the prescription sequence (i.e., cursor 'UP' inscribed with an upward pointing arrow) must not be used for editing or any other purpose.

disagreed. There were further discussions of poor design and other problems that caused a 10- to 30-percent underdosing in both modes. The meeting notes said

There was a general complaint by all users present about the lack of information propagation. The users were not happy about receiving incomplete information. The AECL representative countered by stating that AECL does not wish to spread rumors and that AECL has no policy to 'keep things quiet'. The consensus among the users was that an improvement was necessary.

After the first user's group meeting, there were two user's group newsletters. The first, dated fall 1986, contained letters from Tim Still, the Kennebec stone physicist, who complained about what he considered to be eight major problems he had experienced with the Therac-25. These problems included poor screen-refresh subroutines that leave trash and erroneous information on the operator console and some tape-loading problems upon startup that he discovered involved the use of "phantom tables" to trigger the interlock system in the event of a load failure instead of using a checksum. He asked the question, "Is programming safety relying too much on the software interlock routines?" The second user's group newsletter, in December 1986, further discussed the implications of the phantom table problem.

AECL produced its first CAP on June 13, 1986. The FDA asked for changes and additional information about the software, including a software test plan. AECL responded on September 26 with several documents describing the software and its modifications but no test plan. They explained how the Therac-25 software evolved from the Therac-6 software and stated that "no single test plan and report exists for the software since both hardware and software were tested and exercised separately and together over many years." AECL concluded that the current CAP improved "machine safety by many orders of magnitude and virtually eliminates the possibility of lethal doses as delivered in the Tyler incident."

An FDA internal memo dated October 20 commented on these AECL submissions, raising several concerns:

Unfortunately, the AECL response also seems to point out an apparent lack of documentation on software specifications and a software test plan.

...concerns include the question of previous knowledge of problems by AECL, the apparent paucity of software quality assurance at the manufacturing facility, and possible warnings and information dissemination to others of the generic type problems.

...As mentioned in my first review, there is some confusion on whether the manufacturer should have been aware of the software problems prior to the ARO's [Accidental Radiation Overdoses] in Texas. AECL had received official notification of a law suit in November 1985 from a patient claiming accidental over-exposure from a Therac-25 in Marietta, Georgia.... If knowledge of these software deficiencies were known beforehand, what would be the FDA's posture in this case?

...The materials submitted by the manufacturer have not been in sufficient detail and clarity to ensure an adequate software quality assurance program currently exists. For example, a response has not been provided with respect to the software part of the CAP to the CDRH's [FDA Center for Devices and Radiological Health] request for documentation on the revised requirements and specifications for the new software. In addition, an analysis has not been provided, as requested, on the interaction with other portions of the software to demonstrate the corrected software does not adversely affect other software functions.

The July 23 letter from the CDRH requested a documented test plan including several specific pieces of information identified in the letter. This request has been ignored up to this point by the manufacturer. Considering the ramifications of the current software problem, changes in software QA attitudes are needed at AECL.

AECL also planned to retain the malfunction codes, but the FDA required better warnings for the operators. Furthermore, AECL had not planned on any quality assurance testing to ensure exact copying of software, but the FDA insisted on it. The FDA further requested assurances that rigorous testing would become a standard part of AECL's software modification procedures.

We also expressed our concern that you did not intend to perform the protocol to future modifications to software. We believe that

the rigorous testing must be performed each time a modification is made in order to ensure the modification does not adversely affect the safety of the system.

AECL was also asked to draw up an installation test plan to ensure that both hardware and software changes perform as designed when installed.

AECL submitted CAP Revision 2 and supporting documentation on December 22, 1986. They changed the CAP to have dose malfunctions suspend treatment and included a plan for meaningful error messages and highlighted dose error messages. They also expanded their diagrams of software modifications and expanded their test plan to cover hardware and software.

3.6 Yakima Valley Memorial Hospital, January 1987

On Saturday, January 17, 1987, the second patient of the day was to be treated for a carcinoma. This patient was to receive two film verification exposures of 4 and 3 rads plus a 79-rad photon treatment (for a total exposure of 86 rads.)

Film was placed under the patient and 4 rads were administered. After the machine paused to open the collimator jaws further, the second exposure of 3 rads was administered. The machine paused again.

The operator entered the treatment room to remove the film and verify the patient's precise position. He used the hand control in the treatment room to rotate the turntable to the field light position, which allowed him to check the alignment of the machine with respect to the patient's body in order to verify proper beam position. He then either pressed the *set* button on the hand control or left the room and typed a set command at the console to return the turntable to the proper position for treatment; there is some confusion as to exactly what transpired. When he left the room, he forgot to remove the film from underneath the patient. The console displayed "beam ready," and the operator hit the **B** key to turn the beam on.

The beam came on, but the console displayed no dose or dose rate. After five or six seconds, the unit shut down with a pause and displayed a message. The message "may have disappeared quickly"; the operator was unclear on this point. However, since the machine merely paused, he was able to push the **C** key to proceed with treatment.

The machine paused again, this time displaying FLATNESS on the reason

line. The operator heard the patient say something over the intercom, but could not understand him. He went into the room to speak with the patient, who reported "feeling a burning sensation" in the chest. The console displayed only the total dose of the two film exposures (7 rads) and nothing more.

Later in the day, the patient developed a skin burn over the entire treatment area. Four days later, the redness developed a striped pattern matching the slots in the blocking tray. The striped pattern was similar to the burn a year earlier at this same hospital, which had first been ascribed to a heating pad and later officially labeled by the hospital as "cause unknown."

AECL began an investigation, and users were told to confirm the turntable position visually before turning on the beam. All tests run by the AECL engineers indicated that the machine was working perfectly. From the information that had been gathered to that point, it was suspected that the electron beam had come on when the turntable was in the field light position. But the investigators could not reproduce the fault condition.

On the following Thursday, AECL sent in an engineer from Ottawa to investigate. The hospital physicist had, in the meantime, run some tests himself. He placed a film in the Therac's beam and then ran two exposures of X-ray parameters with the turntable in field light position. The film appeared to match the film that was left (by mistake) under the patient during the accident.

After a week of checking the hardware, AECL determined that the "incorrect machine operation was probably not caused by hardware alone." After checking the software, AECL engineers discovered a flaw (described below) that could explain the erroneous behavior. The coding problems explaining this accident are completely different from those associated with the Tyler accidents.

Preliminary dose measurements by AECL indicated that the dose delivered under these conditions—that is, when the turntable is in the field light position—is on the order of 4,000 to 5,000 rads. After two attempts, the patient could have received 8,000 to 10,000 instead of the 86 rads prescribed. AECL again called users on January 26 (nine days after the accident) and gave them detailed instructions on how to avoid this problem. In an FDA internal report on the accident, the AECL quality assurance manager investigating the problem is quoted as saying that the software and hardware changes to be retrofitted following the Tyler accident nine months earlier

(but which had not yet been installed) would have prevented the Yakima accident.

The patient died in April from complications related to the overdose. He had a terminal form of cancer, but a lawsuit was initiated by his survivors alleging that he died sooner than he would have and endured unnecessary pain and suffering due to the radiation overdose. The suit, like all the others, was settled out of court.

3.6.1 The Yakima Software “Bug”

The software problem for the second Yakima accident is fairly well-established and different from that implicated in the Tyler accidents. There is no way to determine what particular software design errors were related to the Kennebstone, Hamilton, and first Yakima accidents. Given the unsafe programming practices exhibited in the code, unknown race conditions or errors could have been responsible for them. There is speculation, however, that the Hamilton accident was the same as this second Yakima overdose. In a report of a conference call on January 26, 1987, between the AECL quality assurance manager and Ed Miller of the FDA discussing the Yakima accident, Miller

notes

This situation probably occurred in the Hamilton, Ontario accident a couple of years ago. It was not discovered at that time and the cause was attributed to intermittent interlock failure. The subsequent recall of the multiple microswitch logic network did not really solve the problem.

The second Yakima accident was again attributed to a type of race condition in the software — this one allowed the device to be activated in an error setting (a “failure” of a software interlock). The Tyler accidents were related to problems in the data-entry routines that allowed the code to proceed to Set Up Test before the full prescription had been entered and acted upon. The Yakima accident involved problems encountered later in the logic after the treatment monitor Treat reaches Set Up Test.

The Therac-25’s field light feature allows very precise positioning of the patient for treatment. The operator can control the machine right at the treatment site using a small hand control that offers certain limited functions for patient setup, including setting gantry, collimator, and table motions.




Figure 4: The Yakima software flaw.

Normally, the operator enters all the prescription data at the console (outside the treatment room) before the final setup of all machine parameters is completed in the treatment room. This gives rise to an UNVERIFIED condition at the console. The operator then completes patient setup in the treatment room, and all relevant parameters now VERIFY. The console displays a message to PRESS SET BUTTON while the turntable is in the field light position. The operator now presses the set button on the hand control or types “set” at the console. That should set the collimator to the proper position for treatment.

In the software, after the prescription is entered and verified by the Data-ent routine, the control variable **Tphase** is changed so that the **Set Up Test** routine is entered (Figure 4). Every pass through the **Set Up Test** rou-

time increments the upper collimator position check, a shared variable called Class3. If Class3 is nonzero, there is an inconsistency and treatment should not proceed. A zero value for Class3 indicates that the relevant parameters are consistent with treatment, and the software does not inhibit the beam.

After setting the Class3 variable, Set Up Test next checks for any malfunctions in the system by checking another shared variable (set by a routine that actually handles the interlock checking) called F\$mal to see if it has a nonzero value. A nonzero value in F\$mal indicates that the machine is not ready for treatment, and the Set Up Test subroutine is rescheduled. When F\$mal is zero (indicating that everything is ready for treatment), the Set Up Test subroutine sets the Tphase variable equal to 2, which results in next scheduling the Set Up Done subroutine and the treatment is allowed to continue.

The actual interlock checking is performed by a concurrent Housekeeper task (Hkper). The upper collimator position check is performed by a subroutine of Hkper called Lmtchk (analog-to-digital limit checking). Lmtchk first checks the Class3 variable. If Class3 contains a non-zero value, Lmtchk calls the Check Collimator (Chkcol) subroutine. If Class3 contains zero, Chkcol is bypassed and the upper collimator position check is not performed. The Chkcol subroutine sets or resets bit 9 of the F\$mal shared variable, depending on the position of the upper collimator—which in turn is checked by the Set Up Test subroutine of Treat to decide whether to reschedule itself or to proceed to Set Up Done.

During machine setup, Set Up Test will be executed several hundred times because it reschedules itself waiting for other events to occur. In the code, the Class3 variable is incremented by one in each pass through Set Up Test. Since the Class3 variable is one byte, it can only contain a maximum value of 255 decimal. Thus, on every 256th pass through the Set Up Test code, the variable will overflow and have a zero value. That means that on every 256th pass through Set Up Test, the upper collimator will not be checked and an upper collimator fault will not be detected.

The overexposure occurred when the operator hit the “set” button at the precise moment that Class3 rolled over to zero. Thus, Chkcol was not executed and F\$mal was not set to indicate that the upper collimator was still in the field-light position. The software turned on the full 25 MeV without the target in place and without scanning. A highly concentrated electron beam resulted, which was scattered and deflected by the stainless

steel mirror that was in the path.

The technical “fix” implemented for this particular software flaw is described by AECL as simple: the program is changed so that the Class3 variable is set to some fixed nonzero value each time through Set Up Test instead of being incremented.

3.6.2 Manufacturer, Government, and User Response

On February 3, 1987, after interaction with the FDA and others, including the user’s group, AECL announced to its customers

1. A new software release to correct both the Tyler and Yakima software problems
2. A hardware single-pulse shutdown circuit
3. A turntable potentiometer to independently monitor turntable position
4. A hardware turntable interlock circuit

The second item, a hardware single-pulse shutdown circuit, essentially acts as a hardware interlock to prevent overdosing by detecting an unsafe level of radiation and halting beam output after one pulse of high energy and current. This interlock effectively provides an independent way to protect against a wide range of potential hardware failures and software errors. The third item, a turntable potentiometer, was the safety device recommended by several groups after the Hamilton accident.

After the second Yakima accident, the FDA became concerned that the use of the Therac-25 during the CAP process, even with AECL’s interim operating instructions, involved too much risk to patients. The FDA concluded that the accidents demonstrated that the software alone could not be relied upon to assure safe operation of the machine. In a February 18, 1987, internal FDA memorandum, the Director of the Division of Radiological Products wrote:

It is impossible for CDRH to find all potential failure modes and conditions of the software. AECL has indicated the “simple software fix” will correct the turntable position problem displayed at Yakima. We have not yet had the opportunity to evaluate that modification. Even if it does, based upon past history, I am not

convinced that there are not other software glitches that could result in serious injury.

... We are in the position of saying that the proposed CAP can reasonably be expected to correct the deficiencies for which they were developed (Tyler). We cannot say that we are reasonable [sic] confident about the safety of the entire system to prevent or minimize exposure from other fault conditions.

On February 6, 1987, Ed Miller of the FDA called Pavel Dvorak of Canada's Health and Welfare to advise him that the FDA would recommend that all Therac-25s be shutdown until permanent modifications could be made. According to Miller's notes on the phone call, Dvorak agreed and indicated that Health and Welfare would coordinate their actions with the FDA.

AECL responded on April 13 with an update on the Therac CAP status and a schedule of the nine action items pressed by the users at a user's group meeting in March. This unique and highly productive meeting provided an unusual opportunity to involve the users in the CAP evaluation process. It brought together all concerned parties in one place and at one time so that a course of action could be decided upon and approved as quickly as possible. The attendees included representatives from

- The manufacturer (AECL)
- All users, including their technical and legal staffs
- The FDA and the Canadian Bureau of Radiation and Medical Devices
- the Canadian Atomic Energy Control Board
- the Province of Ontario
- the Radiation Regulations Committee of the Canadian Association of Physicists

According to Gordon Symonds, from the Canadian BRMD, this meeting was very important to the resolution of the problems, since the regulators, users, and manufacturer arrived at a consensus in one day.

At this second user's meeting, the participants carefully reviewed all the six known major Therac-25 accidents to that date and discussed the elements of the CAP along with possible additional modifications. They came up with a prioritized list of modifications they wanted included in the CAP and

expressed concerns about the lack of independent evaluation of the software and the lack of a hardcopy audit trail to assist in diagnosing faults.

The AECL representative, who was the quality assurance manager, responded that tests had been done on the CAP changes, but that the tests were not documented and that independent evaluation of the software "might not be possible." He claimed that two outside experts had reviewed the software, but he could not provide their names. In response to user requests for a hard copy audit trail and access to source code, he explained that memory limitations would not permit including such options and that source code would not be made available to users.

On May 1, AECL issued CAP Revision 4 as a result of the FDA comments and the user's meeting input. The FDA response on May 26 approved the CAP subject to submission of the final test plan results and an independent safety analysis, distribution of the draft revised manual to customers, and completion of the CAP by June 30, 1987. The FDA concluded by rating this a Class I recall: a recall in which there is a reasonable probability that the use of, or exposure to, a violative product will cause serious adverse health consequences or death [1].

AECL sent more supporting documentation to the FDA on June 5, 1987, including the CAP test plan, a draft operator's manual, and the draft of the new safety analysis. This time the analysis included the software in the fault trees but used a "generic failure rate" of 10^{-4} for software events. This number was justified as being based on the historical performance of the Therac-25 software. The final report on the safety analysis states that many of the fault trees had a computer malfunction as a causative event, and the outcome for quantification was therefore dependent on the failure rate chosen for the software. Assuming that all software errors are equally likely seems rather strange.

A close inspection of the code was also conducted during this safety analysis to "obtain more information on which to base decisions." An outside consultant performed the inspection, which included a detailed examination of the implementation of each function, a search for coding errors, and a qualitative assessment of the software's reliability. No information is provided in the final safety report about whether any particular methodology or tools were used in the software inspection or whether someone just read the code looking for errors.

AECL planned a fifth revision of the CAP to include the testing and final safety analysis results. Referring to the test plan at this, the final stage of the CAP process, an FDA reviewer said,

Amazingly, the test data presented to show that the software changes to handle the edit problems in the Therac-25 are appropriate prove the exact opposite result. A review of the data table in the test results indicates that the final beam type and energy (edit change) has no effect on the initial beam type and energy. I can only assume that either the fix is not right or the data was entered incorrectly. The manufacturer should be admonished for this error. Where is the QC [Quality Control] review for the test program? AECL must: (1) clarify this situation, (2) change the test protocol to prevent this type of error from occurring, and (3) set up appropriate QC control on data review.

A further FDA memo indicated:

[The AECL quality assurance manager] could not give an explanation and will check into the circumstances. He subsequently called back and verified that the technician completed the form incorrectly. Correct operation was witnessed by himself and others. They will repeat and send us the correct data sheet.

At the American Association of Physicists in Medicine meeting in July 1987, a third user's meeting was held. The AECL representative described the status of the latest CAP and explained that the FDA had given verbal approval and that he expected full implementation by the end of August 1987. He went on to review and comment on the prioritized concerns of the last meeting. Three of the user-requested hardware changes had been included in the CAP. Changes to tape load error messages and checksums on the load data would wait until after the CAP was done. Software documentation was described as a lower priority task that needed definition and would not be available to the FDA in any form for over a year.

On July 6, 1987, AECL sent a letter to all users to update them on the FDA's verbal approval of the CAP and to delineate how AECL would proceed. Finally, on July 21, 1987, AECL issued the final and fifth CAP revision. The major features of the final CAP are these:

- All interruptions related to the dosimetry system will go to a treatment suspend, not a treatment pause. Operators will not be allowed to restart the machine without reentering all parameters.
- A software single-pulse shutdown will be added.
- An independent hardware single-pulse shutdown will be added.
- Monitoring logic for turntable position will be improved to ensure that the turntable is in one of the three legal positions.
- A potentiometer will be added to the turntable. The output is used to monitor exact turntable location and provide a visible position signal to the operator.
- Interlocking with the 270-degree bending magnet will be added to ensure that the target and beam flattener are in position if the X-ray mode is selected.
- Beam-on will be prevented if the turntable is in the field light or any intermediate position.
- Cryptic malfunction messages will be replaced with meaningful messages and highlighted dose-rate messages.
- Editing keys will be limited to *cursor up*, *backspace*, and *return*. All other keys will be inoperative.
- A motion-enable footswitch (a type of deadman switch) will be added. The operator will be required to hold this switch closed during movement of certain parts of the machine to prevent unwanted motions when the operator is not in control.
- Twenty three other changes will be made to the software to improve its operation and reliability, including disabling of unused keys, changing the operation of the *set* and *reset* commands, preventing copying of the control program on site, changing the way various detected hardware faults are handled, eliminating errors in the software that were detected during the review process, adding several additional software interlocks, disallowing changes in the service mode while a treatment is in progress, and adding meaningful error messages.
- The known software problems associated with the Tyler and Yakima accidents will be fixed.
- The manuals will be fixed to reflect the changes.

Figure 5 shows a typical Therac-25 installation after the CAP changes were made.




Figure 5: A typical Therac-25 facility after the final CAP.

Ed Miller, the director of the Division of Standards Enforcement, Center for Devices and Radiological Health at the FDA, wrote in 1987:

FDA has performed extensive review of the Therac-25 software and hardware safety systems. We cannot say with absolute certainty that all software problems that might result in improper dose have been found and eliminated. However, we are confident that the hardware and software safety features recently added will prevent future catastrophic consequences of failure.

No Therac-25 accidents have been reported since the final corrective action plan was implemented.

4 Causal Factors

Many lessons can be learned from this series of accidents. A few are considered here.

Overconfidence in Software. A common mistake in engineering, in this case and in many others, is to put too much confidence in software. There seems to be a feeling among nonsoftware professionals that software will not or cannot fail, which leads to complacency and overreliance on computer functions.

A related tendency among engineers is to ignore software. The first safety analysis on the Therac-25 did not include software—although nearly full responsibility for safety rested on it. When problems started occurring, it was assumed that hardware had caused them, and the investigation looked only at the hardware.

Confusing Reliability with Safety. This software was highly reliable. It worked tens of thousands of times before overdosing anyone, and occurrences of erroneous behavior were few and far between. AEC/L assumed that their software was safe because it was reliable, and this led to complacency.

Lack of Defensive Design. The software did not contain self-checks or other error-detection and error-handling features that would have detected

the inconsistencies and coding errors. Audit trails were limited because of a lack of memory. However, today larger memories are available and audit trails and other design techniques must be given high priority in making tradeoff decisions.

Patient reactions were the only real indications of the seriousness of the problems with the Therac-25; there were no independent checks that the machine and its software were operating correctly. Such verification cannot be assigned to operators without providing them with some means of detecting errors: The Therac-25 software “lied” to the operators, and the machine itself was not capable of detecting that a massive overdose had occurred. The ion chambers on the Therac-25 could not handle the high density of ionization from the unscanned electron beam at high beam current; they thus became saturated and gave an indication of a low dosage. Engineers need to design for the worst case.

Failure to Eliminate Root Causes. One of the lessons to be learned from the Therac-25 experiences is that focusing on particular software design errors is not the way to make a system safe. Virtually all complex software can be made to behave in an unexpected fashion under some conditions: There will always be another software bug. Just as engineers would not rely on a design with a hardware single point of failure that could lead to catastrophe, they should not do so if that single point of failure is software.

The Therac-20 contained the same software error implicated in the Tyler deaths, but this machine included hardware interlocks that mitigated the consequences of the error. Protection against software errors can and should be built into both the system and the software itself. We cannot eliminate all software errors, but we can often protect against their worst effects, and we can recognize their likelihood in our decision making.

One of the serious mistakes that led to the multiple Therac-25 accidents was the tendency to believe that the cause of an accident had been determined (e.g., a microswitch failure in the case of Hamilton) without adequate evidence to come to this conclusion and without looking at all possible contributing factors. Without a thorough investigation, it is not possible to determine whether a sensor provided the wrong information, the software provided an incorrect command, or the actuator had a transient failure and did the wrong thing on its own. In the case of the Hamilton accident, a

transient microswitch failure was assumed to be the cause even though the engineers were unable to reproduce the failure or to find anything wrong with the microswitch.

In general, it is a mistake to patch just one causal factor (such as the software) and assume that future accidents will be eliminated. Accidents are unlikely to occur in exactly the same way again. If we patch only the symptoms and ignore the deeper underlying causes, or if we fix only the specific cause of one accident, we are unlikely to have much effect on future accidents. The series of accidents involving the Therac-25 is a good example of exactly this problem: Fixing each individual software flaw as it was found did not solve the safety problems of the device.

Complacency. Often it takes an accident to alert people to the dangers involved in technology. A medical physicist wrote about the Therac-25 accidents:

In the past decade or two, the medical accelerator “industry” has become perhaps a little complacent about safety. We have assumed that the manufacturers have all kinds of safety design experience since they’ve been in the business a long time. We know that there are many safety codes, guides, and regulations to guide them and we have been reassured by the hitherto excellent record of these machines. Except for a few incidents in the 1960’s (e.g., at Hammersmith, Hamburg) the use of medical accelerators has been remarkably free of serious radiation accidents until now. Perhaps, though we have been spoiled by this success [6].

This problem seems to be common in all fields.

Unrealistic Risk Assessments. The first hazard analyses initially ignored software, and then they treated it superficially by assuming that all software errors were equally likely. The probabilistic risk assessments generated undue confidence in the machine and in the results of the risk assessment themselves. When the first Yakima accident was reported to AEC/L, the company did not investigate. Their evidence for their belief that the radiation burn could not have been caused by their machine included a probabilistic risk assessment showing that safety had increased by five orders of magnitude as a result of the microswitch fix.

The belief that safety had been increased by such a large amount seems hard to justify. Perhaps it was based on the probability of failure of the microswitch (typically 10⁻⁵) AND-ed with the other interlocks. The problem with all such analyses is that they typically make many independence assumptions and exclude aspects of the problem—in this case, software—that are difficult to quantify but which may have a larger impact on safety than the quantifiable factors that are included.

Inadequate Investigation or Followup on Accident Reports. Every company building safety-critical systems should have audit trails and incident analysis procedures that are applied whenever any hint of a problem is found that might lead to an accident. The first phone call by Tim Sill should have led to an extensive investigation of the events at Kennestone. Certainly, learning about the first lawsuit should have triggered an immediate response.

Inadequate Software Engineering Practices. Some basic software engineering principles that apparently were violated in the case of the Therac-25 include the following:

- Software specifications and documentation should not be an afterthought.
- Rigorous software quality assurance practices and standards should be established.
- Designs should be kept simple and dangerous coding practices avoided.
- Ways to detect errors and get information about them, such as software audit trails, should be designed into the software from the beginning.
- The software should be subjected to extensive testing and formal analysis at the module and software level; system testing alone is not adequate. Regression testing should be performed on all software changes.
- Computer displays and the presentation of information to the operators, such as error messages, along with user manuals and other documentation need to be carefully designed.

The manufacturer said that the hardware and software were “tested and exercised separately or together over many years.” In his deposition for one of the lawsuits, the quality assurance manager explained that testing was done in two parts. A “small amount” of software testing was done on

a simulator, but most of the testing was done as a system. It appears that unit and software testing was minimal, with most of the effort directed at the integrated system test. At a Therac-25 user’s meeting, the same man stated that the Therac-25 software was tested for 2,700 hours. Under questioning by the users, he clarified this as meaning “2700 hours of use.” The FDA difficulty in getting an adequate test plan out of the company and the lack of regression testing are evidence that testing was not done well.

The design is unnecessarily complex for such critical software. It is untestable in the sense that the design ensured that the known errors (there may very well be more that have just not been found) would most likely not have been found using standard testing and verification techniques. This does not mean that software testing is not important, only that software must be designed to be testable and that simple designs may prevent errors in the first place.

Software Reuse. Important lessons about software reuse can be found in these accidents. A naive assumption is often made that reusing software or using commercial off-the-shelf software will increase safety because the software will have been exercised extensively. Reusing software modules does not guarantee safety in the new system to which they are transferred and sometimes leads to awkward and dangerous designs. Safety is a quality of the system in which the software is used; it is not a quality of the software itself. Rewriting the entire software in order to get a clean and simple design may be safer in many cases.

Safe versus Friendly User Interfaces. Making the machine as easy as possible to use may conflict with safety goals. Certainly, the user interface design left much to be desired, but eliminating multiple data entry and assuming that operators would check the values carefully before pressing the return key was unrealistic.

User and Government Oversight and Standards. Once the FDA got involved in the Therac-25, their response was impressive, especially considering how little experience they had with similar problems in computer-controlled medical devices. Since the Therac-25 events, the FDA has moved to improve the reporting system and to augment their procedures and guide-

App. A – Medical Devices: The Therac-25

49

lines to include software. The input and pressure from the user group was also important in getting the machine fixed and provides an important lesson to users in other industries.

References

- [1] C.A. Bowsheer. Medical device recalls: Examination of selected cases. Technical Report GAO Report GAO/PEMD-90-6, U.S. Government Accounting Organization, October 1990.
- [2] C.A. Bowsheer. Medical devices: The public health at risk. Technical Report GAO Report GAO/T-PEMD-90-2, U.S. Government Accounting Organization, 1990.
- [3] M. Kivel, editor. *Radiological Health Bulletin*, volume XX:8. Center for Devices and Radiological Health, Food and Drug Administration, Rockville, Maryland, December 1986.
- [4] Nancy G. Leveson and Clark S. Turner. An investigation of the Therac-25 accidents. *IEEE Computer*, 26(7):18–41, July 1993.
- [5] Ed Miller. The Therac-25 experience. In *Conference of State Radiation Control Program Directors*, 1987.
- [6] J.A. Rawlinson. Report on the Therac-25. In *OCTRF/OCI Physicists Meeting*, Kingston, Ontario, May 1987.
- [7] R. Saltos. Man killed by accident with medical radiation. *Boston Globe*, June 20 1986.

CITY PERILS
THE FIFTY-NINE-STORY CRISIS

THE NEW YORKER, MAY 29, 1995, pp 45-53

What's an engineer's worst nightmare? To realize that the supports he designed for a skyscraper like Citicorp Center are flawed—and hurricane season is approaching.

by **JOE MORGENSTERN**

On a warm June day in 1978, William J. LeMessurier, one of the nation's leading structural engineers, received a phone call at his headquarters, in Cambridge, Massachusetts, from an engineering student in New Jersey. The young man, whose name has been lost in the swirl of subsequent events, said that his professor had assigned him to write a paper on the Citicorp tower, the slash-topped silver skyscraper that had become, on its completion in Manhattan the year before, the seventh-tallest building in the world.


LeMessurier found the subject hard to resist, even though the call caught him in the middle of a meeting. As a structural consultant to the architect Hugh Stubbins, Jr., he had designed the twenty-five-thousand-ton steel skeleton beneath the tower's sleek aluminum skin. And, in a field where architects usually get all the credit, the engineer, then fifty-two, had won his own share of praise for the tower's technical elegance and singular grace; indeed, earlier that year he had been elected to the National Academy of Engineering, the highest honor his profession bestows. Excusing himself from the meeting, LeMessurier asked his caller how he could help.

The student wondered about the columns—there are four—that held the building up. According to his professor, LeMessurier had put them in the wrong place.

"I was very nice to this young man," LeMessurier recalls. "But I said, 'Listen, I want you to tell your teacher that he doesn't know what the hell he's talking about, because he doesn't know the problem that had to be solved.' I promised to call back after my meeting and explain the whole thing."

The problem had been posed by a church. When planning for Citicorp Center began, in the early nineteen-seventies, the site of choice was on the east side of Lexington Avenue between Fifty-third and Fifty-fourth Streets, directly across the street from Citicorp's headquarters. But the northwest corner of that block was occupied by St Peter's Church, a decaying Gothic structure built in 1905. Since St. Peter's owned the corner, and one of the world's biggest banking corporations wanted the whole block, the church was able to strike a deal that seemed heaven-sent: its old building would be demolished and a new one built as a free-standing part of Citicorp Center.

To clear space for the new church, Hugh Stubbins and Bill LeMessurier (he pronounces his name "LeMeasure") set their fifty-nine-story tower on four massive, nine-story-high stilts, and positioned them at the center of each side, rather than at each corner. This daring scheme allowed the designers to cantilever the building's corners



seventy-two feet out over the church, on the northwest, and over a plaza on the southwest. The columns also produced high visual drama: a nine-hundred-and-fourteen-foot monolith that seemed all but weightless as it hovered above the street.

When LeMessurier called the student back, he related this with the pride of a master builder and the elaborate patience of a pedagogue; he, too, taught a structural-engineering class, to architecture students at Harvard. Then he explained how the peculiar geometry of the building, far from constituting a mistake, put the columns in the strongest position to resist what sailors call quartering winds—those which come from a diagonal and, by flowing across two sides of a building at once, increase the forces on both. For further enlightenment on the matter, he referred the student to a technical article written by LeMessurier's partner in New York, an engineer named Stanley Goldstein. LeMessurier recalls, "I gave him a lot of information, and I said, 'Now you really have something on your professor, because you can explain all of this to him yourself.'"

Later that day, LeMessurier decided that the information would interest his own students; like sailors, designers of tall buildings must know the wind and respect its power. And the columns were only part of the tower's defense against swaying in severe winds. A classroom lecture would also look at the tower's unusual system of wind braces, which LeMessurier had first sketched out, in a burst of almost ecstatic invention, on a napkin in a Greek restaurant in Cambridge: forty-eight braces, in six tiers of eight, arrayed like giant chevrons behind the building's curtain of aluminum and glass. ("I'm very vain," LeMessurier says. "I would have liked my stuff to be expressed on the outside of the building, but Stubbins wouldn't have it. In the end, I told myself I didn't give a damn—the structure was there, it'd be seen by God.")

LeMessurier had long since established the strength of those braces in perpendicular winds—the only calculation required by New York City's building code. Now, in the spirit of intellectual play, he wanted to see if they were just as strong in winds hitting from forty-five degrees. His new calculations surprised him. In four of the eight chevrons in each tier, a quartering wind increased the strain by forty per cent. Under normal circumstances, the wind braces would have absorbed the extra load without so much as a tremor. But the circumstances were not normal. A few weeks before, during a meeting in his office, LeMessurier had learned of a crucial change in the way the braces were joined.

The meeting had been called, during the month of May, to review plans for two new skyscrapers in Pittsburgh. Those towers, too, were designed by Hugh Stubbins with LeMessurier as structural consultant, and the plans called for wind braces similar to those used in Citicorp Center, with the same specifications for welded joints. This was top of the-line engineering; two structural members joined by a skilled welder become as strong as one. But welded joints, which are labor-intensive and therefore expensive, can be needlessly strong; in most cases, bolted joints are more practical and equally safe. That was the position taken at the May meeting by a man from U.S. Steel, a potential bidder on the contract to erect the Pittsburgh towers. If welded joints were a condition, the project might be too expensive and his firm might not want to take it on.

To reassure him, LeMessurier put in a call to his office in New York. "I spoke to Stanley Goldstein and said, 'Tell me about your success with those welded joints in Citicorp.' And Stanley said, 'Oh, didn't you know? They were changed—they were never welded at all, because Bethlehem Steel came to us and said they didn't think we needed to do it.'" Bethlehem, which built the Citicorp tower, had made the same objection—welds were stronger than necessary, bolts were the right way to do the job. On August 1, 1974, LeMessurier's New York office—actually a venture in conjunction with an old-line Manhattan firm called the Office of James Ruderman—had accepted Bethlehem's proposal.

This news gave LeMessurier no cause for concern in the days immediately following the meeting. The choice of bolted joints was technically sound and professionally correct. Even the failure of his associates to flag him on the design change was justifiable; had every decision on the site in Manhattan waited for approval from Cambridge, the building would never have been finished. Most important, modern skyscrapers are so strong that catastrophic collapse is not considered a realistic prospect; when engineers seek to limit a building's sway, they do so for the tenants' comfort.

Yet now, a month after the May meeting, the substitution of bolted joints raised a troubling question. If the bracing system was unusually sensitive to quartering winds, as LeMessurier had just discovered, so were the joints that held it together. The question was whether the Manhattan team had considered such winds when it designed the bolts. “I didn’t go into a panic over it,” LeMessurier says. “But I was haunted by a hunch that it was something I’d better look into,”

On July 24th, he flew to New York, where his hunch was soon confirmed: his people had taken only perpendicular winds into account. And he discovered another “subtle conceptual error,” as he calls it now—one that threatened to make the situation much worse.

To understand why, one must look at the interplay of opposing forces in a windblown building. The wind causes tension in the structural members—that is, it tries to blow the building down. At the same time, some of that tension, measured in thousands, or even millions, of pounds, is offset by the force of gravity, which, by pressing the members together, tends to hold the building in place. The joints must be strong enough to resist the differential between these forces—the amount of wind tension minus the amount of compression.

Within this seemingly simple computation, however, lurks a powerful multiplier. At any given level of the building, the compression figure remains constant; the wind may blow harder, but the structure doesn’t get any heavier. Thus, immense leverage can result from higher wind forces. In the Citicorp tower, the forty-percent increase in tension produced by a quartering wind became a hundred-and-sixty-per-cent increase on the building’s bolts.

Precisely because of that leverage, a margin of safety is built into the standard formulas for calculating how strong a joint must be; these formulas are contained in an American Institute of Steel Construction specification that deals with joints in structural columns. What LeMessurier found in New York, however, was that the people on his team had disregarded the standard. They had chosen to define the diagonal wind braces not as columns but as trusses, which are exempt from the safety factor. As a result, the bolts holding the joints together were perilously few. “By then,” LeMessurier says, “I was getting pretty shaky.”

He later detailed these mistakes in a thirty-page document called “Project SERENE”; the acronym, both rueful and apt, stands for “Special Engineering Review of Events Nobody Envisioned.” What emerges from this document, which has been confidential until now, and from interviews with LeMessurier and other principals in the events, is not malfeasance, or even negligence, but a series of miscalculations that flowed from a specific mindset. In the case of the Citicorp tower, the first event that nobody envisioned had taken place when LeMessurier sketched, on a restaurant napkin, a bracing system with an inherent sensitivity to quartering winds. None of his associates identified this as a problem, let alone understood that they were compounding it with their fuzzy semantics. In the stiff, angular language of “Project SERENE,” “consideration of wind from non-perpendicular directions on ordinary rectangular buildings is generally not discussed in the literature or in the classroom.”

LeMessurier tried to take comfort from another element of Citicorp’s advanced design: the building’s tuned mass damper. This machine, built at his behest and perched where the bells would have been if the Citicorp tower had been a cathedral, was essentially a four-hundred-and-ten-ton block of concrete, attached to huge springs and floating on a film of oil. When the building swayed, the block’s inertia worked to damp the movement and calm tenants’ queasy stomachs. Reducing sway was of special importance, because the Citicorp tower was an unusually lightweight building; the twenty-five thousand tons of steel in its skeleton contrasted with the Empire State Building’s sixty-thousand-ton superstructure. Yet the damper, the first of its kind in a large building was never meant to be a safety device. At best, the machine might reduce the danger, not dispel it.

Before making a final judgment on how dangerous the bolted joints were, LeMessurier turned to a Canadian engineer named Alan Davenport, the director of the Boundary Layer Wind Tunnel Laboratory, at the University of Western Ontario, and a world authority on the behavior of buildings in

high winds. During the Citicorp tower's design, Davenport had run extensive tests on scale models of the structure. Now LeMessurier asked him and his deputy to retrieve the relevant files and magnetic tapes. "If we were going to think about such things as the possibility of failure," LeMessurier says—the word "failure" being a euphemism for the Citicorp tower's falling down—"we would think about it in terms of the best knowledge that the state of the art can produce, which is what these guys could provide for me."

On July 26th, he flew to London, Ontario, and met with Davenport. Presenting his new calculations, LeMessurier asked the Canadians to evaluate them in the light of the original data. "And you have to tell me the truth," he added. "Don't go easy if it doesn't come out the right way." It didn't, and they didn't. The tale told by the wind-tunnel experts was more alarming than LeMessurier had expected. His assumption of a forty-per-cent increase in stress from diagonal winds was theoretically correct, but it could go higher in the real world, when storms lashed at the building and set it vibrating like a tuning fork. "Oh, my God," he thought, "now we've got that on top of an error from the bolts being under-designed." Refining their data further, the Canadians teased out wind-tunnel forces for each structural member in the building, with and without the tuned mass damper in operation; it remained for LeMessurier to interpret the numbers' meaning.

First, he went to Cambridge, where he talked to a trusted associate, and then he called his wife at their summer house in Maine. "Dorothy knew what I was up to," he says. "I told her, 'I think we've got a problem here, and I'm going to sit down and try to think about it.'" On July 28th, he drove to the northern shore of Sebago Lake, took an outboard motorboat a quarter of a mile across the water to his house on a twelve-acre private island, and worked through the wind-tunnel numbers, joint by joint and floor by floor.

The weakest joint, he discovered, was at the building's thirtieth floor; if that one gave way, catastrophic failure of the whole structure would follow. Next, he took New York City weather records provided by Alan Davenport and calculated the probability of a storm severe enough to tear that joint apart. His figures told him that such an event had a statistical probability of occurring as often as once every sixteen years—what meteorologists call a sixteen-year storm.

"That was very low, awesomely low," LeMessurier said, his voice hushed as if the horror of discovery were still fresh. "To put it another way, there was one chance in sixteen in any year, including that one." When the steadying influence of the tuned mass damper was factored in, the probability dwindled to one in fifty-five—a fifty-five-year storm. But the machine required electric current, which might fail as soon as a major storm hit.

As an experienced engineer, LeMessurier liked to think he could solve most structural problems, and the Citicorp tower was no exception. The bolted joints were readily accessible, thanks to Hugh Stubbins' insistence on putting the chevrons inside the building's skin rather than displaying them outside. With money and materials, the joints could be reinforced by welding heavy steel plates over them, like giant Band-Aids. But time was short; this was the end of July, and the height of the hurricane season was approaching. To avert disaster, LeMessurier would have to blow the whistle quickly on himself. That meant facing the pain of possible protracted litigation, probable bankruptcy, and professional disgrace. It also meant shock and dismay for Citicorp's officers and shareholders when they learned that the bank's proud new corporate symbol, built at a cost of a hundred and seventy-five million dollars, was threatened with collapse.

On the island, LeMessurier considered his options. Silence was one of them; only Davenport knew the full implications of what he had found, and he would not disclose them on his own. Suicide was another, if LeMessurier drove along the Maine Turnpike at a hundred miles an hour and steered into a bridge abutment, that would be that. But keeping silent required betting other people's lives against the odds, while suicide struck him as a coward's way out and—although he was passionate about nineteenth-century classical music—unconvincingly melodramatic. What seized him an instant later was entirely convincing, because it was so unexpected almost giddy sense of power. "I had information that nobody else in the world had," LeMessurier recalls. "I had power in my hands to effect extraordinary events that only I could initiate. I mean, sixteen years to failure—that was very simple, very clear-cut. I almost said, 'thank you, dear Lord, for making this problem so sharply defined that there's no choice to make.'"

At his office in Cambridge on the morning of Monday, July 31st, LeMessurier tried to reach Hugh Stubbins whose firm was upstairs in the same building, but Stubbins was in California and unavailable by phone. Then he called Stubbins' lawyer, Carl Sapers, and outlined the emergency over lunch. Sapers advised him against telling Citicorp until he had consulted with his own company's liability insurers, the Northbrook Insurance Company, in Northbrook, Illinois. When LeMessurier called Northbrook, which represented the Office of James Ruderman as well, someone there referred him to the company's attorneys in New York and warned him not to discuss the matter with anyone else.

At 9 A.M. on Tuesday, in New York, LeMessurier faced a battery of lawyers who, he says, "wanted to meet me to find out if I was nutty." Being lawyers, not engineers, they were hard put to reconcile his dispassionate tone with the apocalyptic thrust of his prophecy. They also bridled at his carefully qualified answers to seemingly simple questions. When they asked how big a storm it would take to blow the building down, LeMessurier confined himself to statistical probabilities—a storm that might occur once in sixteen years.

When they pressed him for specific wind velocities—would the wind have to be at eighty miles per hour, or ninety, or ninety-five?—he insisted that such figures were not significant in themselves, since every structure was uniquely sensitive to certain winds; an eighty-five-mile-per-hour wind that blew for sixteen minutes from the northwest might pose less of a threat to a particular building than an eighty-mile-per-hour wind that blew for fourteen minutes from the southwest.

But the lawyers certainly understood that they had a crisis on their hands, so they sent for an expert adviser they trusted: Leslie Robertson, an engineer who had been a structural consultant for the World Trade Center. "I got a phone call out of the blue from some lawyer summoning me to a meeting," Robertson says. "What's it about?" "You'll find out when you get there." "Sorry, I have other things to do—I don't attend meetings on that basis." A few minutes later, I got another call, from another lawyer, who said there'd been a problem with Citicorp Center. I went to the meeting that morning, and I didn't know anybody there but Bill. He stood up and explained what he perceived were the difficulties with the building, and everyone, of course, was very concerned. Then they turned to me and said, 'Well?' I said, 'Look, if this is in fact the case, you have a very serious problem.' "

The two structural engineers were peers, but not friends. LeMessurier was a visionary with a fondness for heroic designs, though he was also an energetic manager. Robertson was a stickler for technical detail, a man fascinated by how things fit together. LeMessurier, older by two years, was voluble and intense, with a courtly rhetorical style. Robertson was tall, trim, brisk, and edgily funny, but made no effort to hide his impatience with things that didn't interest him.

In addition to his engineering expertise, Robertson brought to the table a background in disaster management. He had worked with such groups as the National Science Foundation and the National Research Council on teams that studied the aftermaths of earthquakes, hurricanes, and floods. (In 1993, he worked with the F.B.I. on the World Trade Center bombing.) For the liability lawyers, this special perspective enhanced his stature as a consultant, but it unsettled LeMessurier from the start. As he remembers it, "Robertson predicted to everybody present that within hours of the time Citicorp heard about this the whole building would be evacuated. I almost fainted. I didn't want that to happen." (For his part, Robertson recalls making no such dire prediction.)

LeMessurier didn't think an evacuation would be necessary. He believed that the building was safe for occupancy in all but the most violent weather, thanks to the tuned mass damper, and he insisted that the damper's reliability in a storm could be assured by installing emergency generators. Robertson conceded the importance of keeping the damper running—it had performed flawlessly since it became operational earlier that year—but, because, in his view, its value as a safety device was unproved, he flatly refused to consider it as a mitigating factor. (In a conversation shortly after the World Trade Center bombing, Robertson noted dryly that the twin towers' emergency generators "lasted for fifteen minutes.")

One point on which everyone agreed was that LeMessurier, together with Stubbins, needed to inform Citicorp as soon as possible. Only Stubbins had ever dealt directly with Citicorp's chairman, Walter B. Wriston, and he was flying home that same day from California and still didn't know his building was flawed. That evening, LeMessurier took the shuttle to Boston, went to Stubbins' house in Cambridge, and broke the news. "He winced, I must admit—here was his masterpiece," LeMessurier says. "But he's a man of enormous resilience, a very grown man, and fortunately we had a lifelong relationship of trust."

The next morning, August 2nd, Stubbins and LeMessurier flew to New York, went to LeMessurier's office at 515 Madison Avenue, put in a call to Wriston, but failed to penetrate the layers of secretaries and assistants that insulated Citicorp's chairman from the outside world. They were no more successful in reaching the bank's president, William I. Spencer, but Stubbins finally managed to get an appointment with Citicorp's executive vice-president, John S. Reed, the man who has now succeeded Wriston as chairman. LeMessurier and Stubbins went to see Reed at the bank's ornate executive offices, in an older building on Lexington Avenue, across the street from Citicorp Center. LeMessurier began by saying, "I have a real problem for you, sir."

Reed was well equipped to understand the problem. He had an engineering background, and he had been involved in the design and construction of Citicorp Center, the company had called him in when it was considering the tuned mass damper. Reed listened impassively as LeMessurier detailed the structural defect and how he thought it could be fixed. LeMessurier says, "I'd already conceived that you could build a little plywood house around each of the connections that were critical, and a welder could work inside it without damaging the tenants' space. You might have to take up the carpet, take down the sheetrock, and work at night, but all this could be done. But the real message I conveyed to him was 'I need your help—at once.' "

When Reed asked how much the repairs would cost, LeMessurier offered an estimate of a million dollars. At the end of the meeting, which lasted half an hour, Reed thanked the two men courteously, though noncommittally, and told them to go back to their office and await further instructions. They did so, but after waiting for more than an hour they decided to go out to lunch. As they were finishing their meal, a secretary from LeMessurier's office called to say that John Reed would be in the office in ten minutes with Walter Wriston.

In the late nineteen-seventies, when Citicorp began its expansion into global banking, Wriston was one of the most influential bankers in the country. A tall man of piercing intelligence, he was not known for effusiveness in the best of circumstances, and LeMessurier expected none now, what with Citicorp Center—and his own career—literally hanging in the balance. But the bank's chairman was genuinely proud of the building, and he offered his support in getting it fixed.

"Wriston was fantastic," LeMessurier says. "He said, 'I guess my job is to handle the public relations of this, so I'll have to start drafting a press release.' " But he didn't have anything to write on, so someone handed him a yellow pad. That made him laugh. According to LeMessurier, " 'All wars,' Wriston said, 'are won by generals writing on yellow pads.' " In fact, Wriston simply took notes; the press release would not go out for six days. But his laughter put the others at ease. Citicorp's general was on their side.

Within hours of Wriston's visit, LeMessurier's office arranged for emergency generators for the tower's tuned mass damper. The bank issued beepers to LeMessurier and his key engineers, assuring them that Reed and other top managers could be reached by phone at any hour of the day or night. Citicorp also assigned two vice-presidents, Henry DeFord III and Robert Dexter, to manage the repairs; both had overseen the building's construction and knew it well.

The next morning, Thursday, August 3rd, LeMessurier, Robertson, and four of LeMessurier's associates met with DeFord and Dexter in a conference room on the thirtieth floor of Citicorp Center. (The decision to hold the initial meeting near the structure's weakest point was purely coincidental.) LeMessurier outlined his plan to fix the wind braces by welding two-inch-thick steel plates over each of more than two hundred bolted joints. The plan was tentatively approved, pending actual examination of a typical joint, but putting

it into effect depended on the availability of a contractor and on an adequate supply of steel plate. Since Bethlehem Steel had dropped out of the business of fabricating and erecting skyscraper structures, Robertson suggested Karl Koch Erecting, a New Jersey-based firm that had put up the World Trade Center.

“I called them,” Robertson says, “and got, ‘Well, we’re a little busy right now,’ and I said, ‘Hey, you don’t understand what we’re talking about here.’” A few hours later, two Koch engineers joined the meeting. LeMessurier and Robertson took them to an unoccupied floor of the building, and there workmen tore apart enough sheetrock to expose a diagonal connection. Comparing the original drawings of the joints with the nuts-and-bolts reality before their eyes, the engineers concluded that LeMessurier’s plan was indeed feasible. Koch also happened to have all the necessary steel plate on hand, so Citicorp negotiated a contract for welding to begin as soon as LeMessurier’s office could issue new drawings.

Two more contracts were drawn up before the end of the following day. One of them went out to MTS Systems Corporation, the Minneapolis firm that had manufactured the tuned mass damper. MTS was asked to provide full-time technical support—in effect, around-the-dock nurses—to keep its machine in perfect health. The company flew one of its technicians to New York that night. Four days later, in a letter of agreement, MTS asked Citicorp to provide a long list of materials and spare parts, which included three buckets, a grease gun, rags, cleaning solvent, and “1 Radio with weather band.”

The other contract engaged a California firm, also recommended by Robertson, to fit the building with a number of instruments called strain gauges—pieces of tape with zigzag wires running through them. The gauges would be affixed to individual structural members, and electrical impulses from them would be funneled to an improvised communications center in Robertson’s office, eight blocks away, at 230 Park Avenue; like a patient in intensive care, the tower would have every shiver and twitch monitored. But this required new telephone lines, and the phone company refused to budge on its leisurely installation schedule. When Robertson voiced his frustration about this during a late-night meeting in Walter Wriston’s office, Wriston picked up the phone on his desk and called his friend Charles Brown, the president and chief operating officer of AT&T. The new lines went in the next morning.

A different problem-solving approach was taken by Robertson during another nighttime meeting in Citicorp’s executive suite. Wriston wanted copies of some documents that Robertson had shown him, but all the secretaries had gone home—the only people’ on the floor were Wriston, Robertson, and John Reed—and every copying machine was locked. “I’m an engineer,” Robertson says, “so I kneeled down, tipped the door off one of the machines, and we made our copies. I looked up at them a little apologetically, but, what the hell—fixing the door was a few hundred bucks, and these guys had a hundred-and-seventy-five-million-dollar building in trouble across the street.”

Robertson also assembled an advisory group of weather experts from academia and the government’s Brookhaven National Laboratory, on Long Island, and hired two independent weather forecasters to provide wind predictions four times a day. “What worried us more than hurricanes, which give you hours and days to anticipate, were unpredictable events,” Robertson says. “From time to time, we’ve had small tornadoes in this area, and there was a worry that a much bigger one would come down and take hold.” Then Robertson raised an issue that LeMessurier had dreaded discussing. In a meeting on Friday that included LeMessurier, Robertson told Citicorp’s representatives, DeFord and Dexter, that they needed to plan for evacuating Citicorp Center and a large area around it in the event of a high-wind alert.

During the first week of August, discussions had involved only a small circle of company officials and engineers. But the circle widened on Monday, August 7th, when final drawings for the steel plates went out to Arthur Nusbaum, the veteran project manager of HRH Construction, which was the original contractor for Citicorp Center, and Nusbaum, in turn, provided them to Koch Erecting. And it would widen again, because work could not go forward, as Robertson reminded the officials, without consulting the city’s Department of Buildings. Citicorp faced a public-relations debacle unless it came up with a plausible explanation of why its brand-new skyscraper needed fixing.

That night, DeFord and Dexter, following Robertson's advice, met with Mike Reilly, the American Red Cross's director of disaster services for the New York metropolitan area. "They laid out the dilemma, and it was clearly an ominous event," Reilly recalls. From that first meeting, which was attended by Robertson but not by LeMessurier, and from half a dozen subsequent working sessions with other disaster agencies, came plans for joint action by the police and the mayor's Office of Emergency Management, along with the Red Cross. In the event of a wind alert, the police and the mayor's emergency forces would evacuate the building and the surrounding neighborhood, and the Red Cross would mobilize between twelve hundred and two thousand workers to provide food and temporary shelter. "Hal DeFord was the bank's point man for all this," Reilly says. "The anxiety was so heavy on him that we weren't sure if he was going to make it."

On Tuesday morning, August 8th, the public-affairs department of Citibank, Citicorp's chief subsidiary, put out the long delayed press release. In language as bland as a loan officer's wardrobe, the three-paragraph document said unnamed "engineers who designed the building" had recommended that "certain of the connections in Citicorp Center's wind bracing system be strengthened through additional welding." The engineers, the press release added, "have assured us that there is no danger." When DeFord expanded on the handout in interviews, he portrayed the bank as a corporate citizen of exemplary caution—"We wear both belts and suspenders here," he told a reporter for the *News*—that had decided on the welds as soon as it learned of new data based on dynamic-wind tests conducted at the University of Western Ontario.

There was some truth in all this. During LeMessurier's recent trip to Canada, one of Alan Davenport's assistants had mentioned to him that probable wind velocities might be slightly higher, on a statistical basis, than predicted in 1973, during the original tests for Citicorp Center. At the time, LeMessurier viewed this piece of information as one more nail in the coffin of his career, but later, recognizing it as a blessing in disguise, he passed it on to Citicorp as the possible basis of a cover story for the press and for tenants in the building.

On Tuesday afternoon at a meeting in Robertson's office, LeMessurier told the whole truth to New York City's Acting Building Commissioner and nine other senior city officials. For more than an hour, he spoke about the effect of diagonal winds on the Citicorp tower, about the failure of his own office to perceive and communicate the danger, and about the intended repairs.

In the discussion that followed, the city officials asked a few technical questions, and Arthur Nusbaum expressed concern over a shortage of certified welders who had passed the city's structural-welding test. That would not be a problem, the representatives from the Department of Buildings replied; one of the area's most trusted steel inspectors, Neil Moreton, would have the power to test and immediately certify any welder that Citicorp's repair project required. Nusbaum recalls, "Once they said that, I knew we were O.K., because there were steamfitter welders all over the place who could do a fantastic job."

Before the city officials left, they commended LeMessurier for his courage and candor, and expressed a desire to be kept informed as the repair work progressed. Given the urgency of the situation, that was all they could reasonably do. "It wasn't a case of 'We caught you, you skunk,' " Nusbaum says. "It started with a guy who stood up and said, 'I got a problem, I made the problem, let's fix the problem.' If you're gonna kill a guy like LeMessurier, why should anybody ever talk?"

Meanwhile, Robertson's switchboard was besieged by calls. "Every reporter in town wanted to know how come all these people were in our office," Robertson says. Once the meeting ended, the Building Commissioner returned the reporters' calls and, hewing to Citicorp's line, reassured them that the structural work was only a prudent response to new meteorological data.

As a result, press coverage in New York City the next day was as uninformative as the handout: a short piece in the *Wall Street Journal*, which raised no questions about the nature of the new data, and one in the *News*, which dutifully quoted DeFord's remark about belts and suspenders. But when LeMessurier went back to his hotel room, at about 5 P.M. on Wednesday, he learned from his wife, who had come down from

Cambridge to join him, that a reporter from the *Times* had been trying to reach him all afternoon. That worried him greatly; being candid with city officials was one thing, but being interrogated by the *Times* was another. Before returning the call, LeMessurier phoned his friend Carl Sapers, the Boston attorney who represented Hugh Stubbins, and mixed himself a Martini. Sapers understood the need for secrecy, but he saw no real choice; talk to them, he said, and do the best you can. Two minutes after six o'clock, LeMessurier called the *Times* switchboard. As he braced himself for an unpleasant conversation, he heard a recording. The *Times*, along with all the other major papers in the city, had just been shut down by a strike.

Welders started work almost immediately, their torches a dazzlement in the night sky. The weather was sticky, as it had been since the beginning of the month—New Jersey's tomato crop was rotting from too much rained forecasts called for temperatures in the mid-eighties the next day, with no wind; in other words, a perfect day for Citicorp Center.

Yet tropical storms were already churning the Caribbean. Citicorp pushed for repair work around the clock, but Nusbaum refused to allow welding during office hours, for fear that clouds of acrid smoke would cause panic among the tenants and set off every smoke detector in the building. Instead, he brought in drywall crews and carpenters to work from 5 P.M. to 8 P.M., putting up plywood enclosures around the chevrons and tearing down Sheetrock; welders to weld from 8 P.M. until 4 A.M., with the building's fire-alarm system shut off; and then laborers to clean up the epic mess before the first secretaries arrived.

The welders worked seven days a week. Sometimes they worked on unoccupied floors; sometimes they invaded lavish offices. But decor, or the lack of it, had no bearing on their priorities, which were set by LeMessurier. "It was a tense time for the whole month," he says. "I was constantly calculating which joint to fix next, which level of the building was more critical, and I developed charts and graphs of all the consequences: if you fix this, then the rarity of the storm that will cause any trouble lengthens to that."

At Robertson's office, a steady stream of data poured in from the weather forecasters and from the building itself. Occasionally, the strain-gage readings jumped, like spikes on an electrocardiogram, when the technicians from MTS Systems exercised their tuned mass damper to make sure it was working properly. One time, the readings went off the chart, then stopped. This provoked more bafflement than fear, since it seemed unlikely that a hurricane raging on Lexington and Fifty-third Street would go otherwise unnoticed at Forty-sixth and Park. The cause proved to be straightforward enough: When the instrumentation experts from California installed their strain gauges, they had neglected to hire union electricians. "Someone heard about it," LeMessurier says, "went up there in the middle of the night, and snipped all the wires."

For most of August, the weather smiled on Citicorp, or at least held its breath, and the welders made steady progress. LeMessurier felt confident enough to fly off with his wife for a weekend in Maine. As their return flight was coming in for a landing at LaGuardia Airport Sunday night, they looked out across the East River and saw a pillar of fire on the Manhattan skyline. "The welders were working up and down the building, fixing the joints," LeMessurier recalls. "It was an absolutely marvelous thing to see. I said to Dorothy, 'Isn't this wonderful? Nobody knows what's going on, but we know and we can see it **lighting up** the sky.'"

A great deal of work remained. Robertson was insisting on a complete reevaluation of the Citicorp tower: not just the sensitivity of the chevrons to quartering winds but the strength of other skeletal members, the adequacy of braces that kept the supporting columns in plumb, and the rigidity of the building's corrugated metal-and-concrete floors, which Robertson feared might be compromised by trenches carrying electrical connections.

His insistence was proper—settling for less would have compromised Robertson's own position. It amounted to a post-construction autopsy by teams of forensic engineers. For LeMessurier, the reevaluation was harrowing in the extreme; every new doubt about his design for Citicorp Center reflected on him.

In one instance, Robertson's fears were unwarranted: tests showed that the tower floors were entirely sound—the trenches were not a source of weakness. In another, Robertson, assuming the worst about construction tolerances, decided that the columns might be slightly, even though undetectably, out of plumb, and therefore he ordered the installation of supplemental bracing above the fourteenth floor.

Shortly before dawn on Friday, September 1st, weather services carried the news that everyone had been dreading—a major storm, Hurricane Ella, was off Cape Hatteras and heading for New York. At 6:30 A.M., an emergency-planning group convened at the command center in Robertson's office. "Nobody said, 'We're probably going to press the panic button,'" LeMessurier recalls. "Nobody dared say that. But everybody was sweating blood."

As the storm bore down on the city, the bank's representatives, DeFord and Dexter, asked LeMessurier for a report on the status of repairs. He told them that the most critical joints had already been fixed and that the building, with its tuned mass damper operating, could now withstand a two-hundred-year storm. It didn't have to, however. A few hours later, Hurricane Ella veered from its northwesterly course and began moving out to sea.

LeMessurier spent the following night in Manhattan, having canceled plans to spend the Labor Day weekend with his family in Maine. But the hurricane kept moving eastward, and daybreak dispelled any lingering thoughts of evacuation. "Saturday was the most beautiful day that the world's ever seen," LeMessurier says, "with all the humidity drawn away and the skies sunny and crystal clear." Alone in the city, he gave himself a treat he'd been thinking about for years—his first visit to the Cloisters, where he basked in an ineffable calm.

The weather watch ended on September 13th. That same day, Robertson recommended terminating the evacuation plans, too. Welding was completed in October, several weeks before most of the city's newspapers resumed publication. No further stories on the subject appeared in the wake of the strike. The building, in fact, was now strong enough to withstand a seven-hundred-year storm even without the damper, which made it one of the safest structures ever built—and rebuilt—by the hand of man.

Throughout the summer, Citicorp's top management team had concentrated on facilitating repairs, while keeping the lawyers on the sidelines. That changed on September 13th, when Citicorp served notice on LeMessurier and Hugh Stubbins, whose firm held the primary contract, of its intention to seek indemnification for all costs. Their estimate of the costs, according to LeMessurier, amounted to \$4.3 million, including management fees. A much higher total was suggested by Arthur Nusbaum, who recalled that his firm, HRH Construction, spent eight million dollars on structural repairs alone. Citicorp has declined to provide its own figure.

Whatever the actual cost, Citicorp's effort to recoup it was remarkably free of the punitive impulse that often poisons such negotiations. When the terms of a settlement were first discussed—without lawyers—by LeMessurier, on one side, and DeFord and Dexter, on the other, LeMessurier spoke of two million dollars, which was the amount that his liability insurer, the Northbrook Insurance Company, had agreed to pay. "DeFord and Dexter said, 'Well, we've been deeply wounded here,' and they tried to play hardball," LeMessurier says. "But they didn't do it with much conviction." After a second meeting, which included a Northbrook lawyer, the bank agreed to hold Stubbins' firm harmless and to accept the two-million-dollar payment from LeMessurier and his joint-venture partners; no litigation ever ensued. Eight years ago, Citicorp turned the building into a condominium, retaining the land and the shops but selling all the office space, to Japanese buyers, at a handsome profit.

The crisis at Citicorp Center was noteworthy in another respect. It produced heroes, but no villains; everyone connected with the repairs behaved in exemplary fashion, from Walter Wriston and his Citicorp management team to the officials at the city's Department of Buildings. The most striking example, of course, was set by LeMessurier, who emerged with his reputation not merely unscathed but enhanced. When Robertson speaks of him, he says, "I have a lot of admiration for Bill, because he was very

forthcoming. While we say that all engineers would behave as he did, I carry in my mind some skepticism about that.”

In the last few years, LeMessurier has been talking about the summer of 1978 to his classes at Harvard. The tale, as he tells it, is by turns painful, self-deprecating, and self-dramatizing—an engineer who did the right thing. But it also speaks to the larger question of how professional people should behave. “You have a social obligation,” LeMessurier reminds his students. “In return for getting a license and being regarded with respect, you’re supposed to be self-sacrificing and look beyond the interests of yourself and your client to society as a whole. And the most wonderful part of my story is that when I did it nothing bad happened.”

ACM Code of Ethics and Professional Conduct

Adopted by ACM Council 10/16/92

Preamble

Commitment to ethical professional conduct is expected of every member (voting members, associate members, and student members) of the Association for Computing Machinery (ACM).

This Code, consisting of 24 imperatives formulated as statements of personal responsibility, identifies the elements of such a commitment. It contains many, but not all, issues professionals are likely to face. Section 1 outlines fundamental ethical considerations, while Section 2 addresses additional, more specific considerations of professional conduct. Statements in Section 3 pertain more specifically to individuals who have a leadership role, whether in the workplace or in a volunteer capacity such as with organizations like ACM. Principles involving compliance with this Code are given in Section 4.

The Code shall be supplemented by a set of Guidelines, which provide explanation to assist members in dealing with the various issues contained in the Code. It is expected that the Guidelines will be changed more frequently than the Code.

The Code and its supplemented Guidelines are intended to serve as a basis for ethical decision making in the conduct of professional work. Secondly, they may serve as a basis for judging the merit of a formal complaint pertaining to violation of professional ethical standards.

It should be noted that although computing is not mentioned in the imperatives of Section 1, the Code is concerned with how these fundamental imperatives apply to one's conduct as a computing professional. These imperatives are expressed in a general form to emphasize that ethical principles which apply to computer ethics are derived from more general ethical principles.

It is understood that some words and phrases in a code of ethics are subject to varying interpretations, and that any ethical principle may conflict with other ethical principles in specific situations. Questions related to ethical conflicts can best be answered by thoughtful consideration of fundamental principles, rather than reliance on detailed regulations.

Contents & Guidelines

1. General Moral Imperatives.
2. More Specific Professional Responsibilities.
3. Organization Leadership Imperatives.
4. Compliance with the Code.
5. Acknowledgments.

1. GENERAL MORAL IMPERATIVES.

As an ACM member I will

1.1 Contribute to society and human well-being.

This principle concerning the quality of life of all people affirms an obligation to protect fundamental human rights and to respect the diversity of all cultures. An essential aim of computing professionals is to minimize negative consequences of computing systems, including threats to health and safety. When designing or implementing systems, computing professionals must attempt to ensure that the products of their efforts will be used in socially responsible ways, will meet social needs, and will avoid harmful effects to health and welfare.

In addition to a safe social environment, human well-being includes a safe natural environment. Therefore, computing professionals who design and develop systems must be alert to, and make others aware of, any potential damage to the local or global environment.

1.2 Avoid harm to others.

“Harm” means injury or negative consequences, such as undesirable loss of information, loss of property, property damage, or unwanted environmental impacts. This principle prohibits use of computing technology in ways that result in harm to any of the following: users, the general public, employees, employers. Harmful actions include intentional destruction or modification of files and programs leading to serious loss of resources or unnecessary expenditure of human resources such as the time and effort required to purge systems of “computer viruses.”

Well-intended actions, including those that accomplish assigned duties, may lead to harm unexpectedly. In such an event the responsible person or persons are obligated to undo or mitigate the negative consequences as much as possible. One way to avoid unintentional harm is to carefully consider potential impacts on all those affected by decisions made during design and implementation.

To minimize the possibility of indirectly harming others, computing professionals must minimize malfunctions by following generally accepted standards for system design and testing. Furthermore, it is often necessary to assess the social consequences of systems to project the likelihood of any serious harm to others. If system features are misrepresented to users, coworkers, or supervisors, the individual computing professional is responsible for any resulting injury.

In the work environment the computing professional has the additional obligation to report any signs of system dangers that might result in serious personal or social damage. If one's superiors do not act to curtail or mitigate such dangers, it may be necessary to "blow the whistle" to help correct the problem or reduce the risk. However, capricious or misguided reporting of violations can, itself, be harmful. Before reporting violations, all relevant aspects of the incident must be thoroughly assessed. In particular, the assessment of risk and responsibility must be credible. It is suggested that advice be sought from other computing professionals. See principle 2.5 regarding thorough evaluations.

1.3 Be honest and trustworthy.

Honesty is an essential component of trust. Without trust an organization cannot function effectively. The honest computing professional will not make deliberately false or deceptive claims about a system or system design, but will instead provide full disclosure of all pertinent system limitations and problems.

A computer professional has a duty to be honest about his or her own qualifications, and about any circumstances that might lead to conflicts of interest.

Membership in volunteer organizations such as ACM may at times place individuals in situations where their statements or actions could be interpreted as carrying the "weight" of a larger group of professionals. An ACM member will exercise care to not misrepresent ACM or positions and policies of ACM or any ACM units.

1.4 Be fair and take action not to discriminate.

The values of equality, tolerance, respect for others, and the principles of equal justice govern this imperative. Discrimination on the basis of race, sex, religion, age, disability, national origin, or other such factors is an explicit violation of ACM policy and will not be tolerated.

Inequities between different groups of people may result from the use or misuse of information and technology. In a fair society, all individuals would have equal opportunity to participate in, or benefit from, the use of computer resources regardless of race, sex, religion, age, disability, national origin or other such similar factors. However, these ideals do not justify unauthorized use of computer resources nor do they provide an adequate basis for violation of any other ethical imperatives of this code.

1.5 Honor property rights including copyrights and patent.

Violation of copyrights, patents, trade secrets and the terms of license agreements is prohibited by law in most circumstances. Even when software is not so protected, such violations are contrary to professional behavior. Copies of software should be made only with proper authorization. Unauthorized duplication of materials must not be condoned.

1.6 Give proper credit for intellectual property.

Computing professionals are obligated to protect the integrity of intellectual property. Specifically, one must not take credit for other's ideas or work, even in cases where the work has not been explicitly protected by copyright, patent, etc.

1.7 Respect the privacy of others.

Computing and communication technology enables the collection and exchange of personal information on a scale unprecedented in the history of civilization. Thus there is increased potential for violating the privacy of individuals and groups. It is the responsibility of professionals to maintain the privacy and integrity of data describing individuals. This includes taking precautions to ensure the accuracy of data, as well as protecting it from unauthorized access or accidental disclosure to inappropriate individuals. Furthermore, procedures must be established to allow individuals to review their records and correct inaccuracies.

This imperative implies that only the necessary amount of personal information be collected in a system, that retention and disposal periods for that information be clearly defined and enforced, and that personal information gathered for a specific purpose not be used for other purposes without consent of the individual(s). These principles apply to electronic communications, including electronic mail, and prohibit procedures that capture or monitor electronic user data, including messages, without the permission of users or bona fide authorization related to system operation and maintenance. User data observed during the normal duties of system operation and maintenance must be treated with strictest confidentiality, except in cases where it is evidence for the violation of law, organizational regulations, or this Code. In these cases, the nature or contents of that information must be disclosed only to proper authorities.

1.8 Honor confidentiality.

The principle of honesty extends to issues of confidentiality of information whenever one has made an explicit promise to honor confidentiality or, implicitly, when private information not directly related to the performance of one's duties becomes available. The ethical concern is to respect all obligations of confidentiality to employers, clients, and users unless discharged from such obligations by requirements of the law or other principles of this Code.

2. MORE SPECIFIC PROFESSIONAL RESPONSIBILITIES.

As an ACM computing professional I will

2.1 Strive to achieve the highest quality, effectiveness and dignity in both the process and products of professional work.

Excellence is perhaps the most important obligation of a professional. The computing professional must strive to achieve quality and to be cognizant of the serious negative consequences that may result from poor quality in a system.

2.2 Acquire and maintain professional competence.

Excellence depends on individuals who take responsibility for acquiring and maintaining professional competence. A professional must participate in setting standards for appropriate levels of competence, and strive to achieve those standards. Upgrading technical knowledge and competence can be achieved in several ways: doing independent study; attending seminars, conferences, or courses; and being involved in professional organizations.

2.3 Know and respect existing laws pertaining to professional work.

ACM members must obey existing local, state, province, national, and international laws unless there is a compelling ethical basis not to do so. Policies and procedures of the organizations in which one participates must also be obeyed. But compliance must be balanced with the recognition that sometimes existing laws and rules may be immoral or inappropriate and, therefore, must be challenged. Violation of a law or regulation may be ethical when that law or rule has inadequate moral basis or when it conflicts with another law judged to be more important. If one decides to violate a law or rule because it is viewed as unethical, or for any other reason, one must fully accept responsibility for one's actions and for the consequences.

2.4 Accept and provide appropriate professional review.

Quality professional work, especially in the computing profession, depends on professional reviewing and critiquing. Whenever appropriate, individual members should seek and utilize peer review as well as provide critical review of the work of others.

2.5 Give comprehensive and thorough evaluations of computer systems and their impacts, including analysis of possible risks.

Computer professionals must strive to be perceptive, thorough, and objective when evaluating, recommending, and presenting system descriptions and alternatives. Computer professionals are in a position of special trust, and therefore have a special responsibility to provide objective, credible evaluations to employers, clients, users, and the public. When providing evaluations the professional must also identify any relevant conflicts of interest, as stated in imperative 1.3.

As noted in the discussion of principle 1.2 on avoiding harm, any signs of danger from systems must be reported to those who have opportunity and/or responsibility to resolve them. See the guidelines for imperative 1.2 for more details concerning harm, including the reporting of professional violations.

2.6 Honor contracts, agreements, and assigned responsibilities.

Honoring one's commitments is a matter of integrity and honesty. For the computer professional this includes ensuring that system elements perform as intended. Also, when one contracts for work with another party, one has an obligation to keep that party properly informed about progress toward completing that work.

A computing professional has a responsibility to request a change in any assignment that he or she feels cannot be completed as defined. Only after serious consideration and with full disclosure of risks and concerns to the employer or client, should one accept the assignment. The major underlying principle here is the obligation to accept personal accountability for professional work. On some occasions other ethical principles may take greater priority.

A judgment that a specific assignment should not be performed may not be accepted. Having clearly identified one's concerns and reasons for that judgment, but failing to procure a change in that assignment, one may yet be obligated, by contract or by law, to proceed as directed. The computing professional's ethical judgment should be the final guide in deciding whether or not to proceed. Regardless of the decision, one must accept the responsibility for the consequences.

However, performing assignments "against one's own judgment" does not relieve the professional of responsibility for any negative consequences.

2.7 Improve public understanding of computing and its consequences.

Computing professionals have a responsibility to share technical knowledge with the public by encouraging understanding of computing, including the impacts of computer systems and their limitations. This imperative implies an obligation to counter any false views related to computing.

2.8 Access computing and communication resources only when authorized to do so.

Theft or destruction of tangible and electronic property is prohibited by imperative 1.2 – "Avoid harm to others." Trespassing and unauthorized use of a computer or communication system is addressed by this imperative. Trespassing includes accessing communication networks and computer systems, or accounts and/or files associated with those systems, without explicit authorization to do so. Individuals and organizations have the right to restrict access to their systems so long as they do not violate the discrimination principle (see 1.4). No one should enter or use another's computer system, software, or data files without permission. One must always have appropriate approval before using system resources, including communication ports, file space, other system peripherals, and computer time.

3. ORGANIZATIONAL LEADERSHIP IMPERATIVES.

As an ACM member and an organizational leader, I will

BACKGROUND NOTE: This section draws extensively from the draft IFIP Code of Ethics, especially its sections on organizational ethics and international concerns. The ethical obligations of organizations tend to be neglected in most codes of professional conduct, perhaps because these codes are written from the perspective of the individual member. This dilemma is addressed by stating these imperatives from the perspective of the organizational leader. In this context "leader" is viewed as any organizational member who has leadership or educational responsibilities. These imperatives generally may apply to organizations as well as their leaders. In this context "organizations" are corporations, government agencies, and other "employers," as well as volunteer professional organizations.

3.1 Articulate social responsibilities of members of an organizational unit and encourage full acceptance of those responsibilities.

Because organizations of all kinds have impacts on the public, they must accept responsibilities to society. Organizational procedures and attitudes oriented toward quality and the welfare of society will reduce harm to members of the public, thereby serving public interest and fulfilling social responsibility. Therefore, organizational leaders must encourage full participation in meeting social responsibilities as well as quality performance.

3.2 Manage personnel and resources to design and build information systems that enhance the quality of working life.

Organizational leaders are responsible for ensuring that computer systems enhance, not degrade, the quality of working life. When implementing a computer system, organizations must consider the personal and

professional development, physical safety, and human dignity of all workers. Appropriate human-computer ergonomic standards should be considered in system design and in the workplace.

3.3 Acknowledge and support proper and authorized uses of an organization's computing and communication resources.

Because computer systems can become tools to harm as well as to benefit an organization, the leadership has the responsibility to clearly define appropriate and inappropriate uses of organizational computing resources. While the number and scope of such rules should be minimal, they should be fully enforced when established.

3.4 Ensure that users and those who will be affected by a system have their needs clearly articulated during the assessment and design of requirements; later the system must be validated to meet requirements.

Current system users, potential users and other persons whose lives may be affected by a system must have their needs assessed and incorporated in the statement of requirements. System validation should ensure compliance with those requirements.

3.5 Articulate and support policies that protect the dignity of users and others affected by a computing system.

Designing or implementing systems that deliberately or inadvertently demean individuals or groups is ethically unacceptable. Computer professionals who are in decision making positions should verify that systems are designed and implemented to protect personal privacy and enhance personal dignity.

3.6 Create opportunities for members of the organization to learn the principles and limitations of computer systems.

This complements the imperative on public understanding (2.7). Educational opportunities are essential to facilitate optimal participation of all organizational members. Opportunities must be available to all members to help them improve their knowledge and skills in computing, including courses that familiarize them with the consequences and limitations of particular types of systems. In particular, professionals must be made aware of the dangers of building systems around oversimplified models, the improbability of anticipating and designing for every possible operating condition, and other issues related to the complexity of this profession.

4. COMPLIANCE WITH THE CODE.

As an ACM member I will

4.1 Uphold and promote the principles of this Code.

The future of the computing profession depends on both technical and ethical excellence. Not only is it important for ACM computing professionals to adhere to the principles expressed in this Code, each member should encourage and support adherence by other members.

4.2 Treat violations of this code as inconsistent with membership in the ACM.

Adherence of professionals to a code of ethics is largely a voluntary matter. However, if a member does not follow this code by engaging in gross misconduct, membership in ACM may be terminated.

5. ACKNOWLEDGMENTS.

This Code and the supplemental Guidelines were developed by the Task Force for the Revision of the ACM Code of Ethics and Professional Conduct: Ronald E. Anderson, Chair, Gerald Engel, Donald Gotterbarn, Grace C. Hertlein, Alex Hoffman, Bruce Jawer, Deborah G. Johnson, Doris K. Lidtke, Joyce Currie Little, Dianne Martin, Donn B. Parker, Judith A. Perrolle, and Richard S. Rosenberg. The Task Force was organized by ACM/SIGCAS and funding was provided by the ACM SIG Discretionary Fund. This Code and the supplemental Guidelines were adopted by the ACM Council on October 16, 1992.

This Code may be published without permission as long as it is not changed in any way and it carries the copyright notice. Copyright ©1997, Association for Computing Machinery, Inc.

the ACM



IEEE CODE OF ETHICS

WE, THE MEMBERS OF THE IEEE, in recognition of the importance of our technologies in affecting the quality of life throughout the world and in accepting a personal obligation to our profession, its members and the communities we serve, do hereby commit ourselves to the highest ethical and professional conduct and agree:

1. to accept responsibility in making decisions consistent with the safety, health and welfare of the public, and to disclose promptly factors that might endanger the public or the environment;
2. to avoid real or perceived conflicts of interest whenever possible, and to disclose them to affected parties when they do exist;
3. to be honest and realistic in stating claims or estimates based on available data;
4. to reject bribery in all its forms;
5. to improve the understanding of technology, its appropriate application, and potential consequences;
6. to maintain and improve our technical competence and to undertake technological tasks for others only if qualified by training or experience, or after full disclosure of pertinent limitations;
7. to seek, accept, and offer honest criticism of technical work, to acknowledge and correct errors, and to credit properly the contributions of others;
8. to treat fairly all persons regardless of such factors as race, religion, gender, disability, age, or national origin;
9. to avoid injuring others, their property, reputation, or employment by false or malicious action;
10. to assist colleagues and co-workers in their professional development and to support them in following this code of ethics.



Computer Society and ACM Approve Software Engineering Code of Ethics

Don Gotterbarn, Keith Miller, Simon Rogerson
Executive Committee, IEEE-CS/ACM Joint Task Force
on Software Engineering Ethics and Professional Practices

Software engineering has evolved over the past several years from an activity of computer engineering to a discipline in its own right. With an eye toward formalizing the field, the IEEE Computer Society has engaged in several activities to advance the professionalism of software engineering, such as establishing certification requirements for software developers. To complement this work, a joint task force of the Computer Society and the ACM has recently established another linchpin of professionalism for software engineering: a code of ethics.

After an extensive review process, version 5.2 of the Software Engineering Code of Ethics and Professional Practice, recommended last year by the IEEE-CS/ACM Joint Task Force on Software Engineering Ethics and Professional Practices, was adopted by both the IEEE Computer Society and the ACM.

PURPOSE

The Software Engineering Code of Ethics and Professional Practice, intended as a standard for teaching and practicing software engineering, documents the ethical and professional obligations of software engineers. The code should instruct practitioners about the standards society

expects them to meet, about what their peers strive for, and about what to expect of one another. In addition, the code should inform the public about the responsibilities that are important to the profession.

Adopted by the Computer Society and the ACM—two leading international computing societies—the code of ethics is intended as a guide for members of the evolving software engineering profession. The code was developed by a multinational task force with additional input from other professionals from industry, government posts, military installations, and educational professions.

CHANGES TO THE CODE

Major revisions were made between version 3.0—widely distributed through *Computer* (Don Gotterbarn, Keith Miller, and Simon Rogerson, “Software Engineering Code of Ethics, Version 3.0,” November 1997, pp. 88-92) and *Communications of the ACM*—and version 5.2, the recently approved version. The preamble was significantly revised to include specific standards that can help professionals make ethical decisions. To facilitate a quick review of the principles, a shortened version of the code was added to the front of the full version. This shortened version is not intended to be a stand-alone abbreviated code. The details of the full version are necessary to provide clear guidance for the practical application of these ethical principles.

In addition to these changes, the eight principles were reordered to reflect the order in which software professionals should consider their ethical obligations: Version 3.0’s first principle concerned the product, while version 5.2 begins with the public. The primacy of well-being and quality of life of the public in all decisions related to software engineering is emphasized throughout the code. This obligation is the final arbiter in all decisions: “In all these judgements concern for the health, safety and welfare of the

About the Joint Task Force

This Code of Ethics was developed by the IEEE-CS/ACM Joint Task Force on Software Engineering Ethics and Professional Practices. Members are

Executive Committee

Donald Gotterbarn (Chair)
 Keith Miller
 Simon Rogerson

Tom Jewett
 Mark Kanko
 Ernie Kallman
 Duncan Langford
 Joyce Currie Little
 Ed Mechler
 Manuel J. Norman
 Douglas Phillips
 Peter Ron Prinzivalli
 Patrick Sullivan
 John Weckert
 Vivian Weil
 S. Weisband
 Laurie Honour Werth

Members

Steve Barber
 Peter Barnes
 Ilene Burnstein
 Michael Davis
 Amr El-Kadi
 N. Ben Fairweather
 Milton Fulghum
 N. Jayaram

Editor: Mary-Louise G. Piner, *Computer*,
 10662 Los Vaqueros Circle, PO Box 3014,
 Los Alamitos, CA 90720-1314; mpiner@
 computer.org

public is primary; that is, the ‘Public Interest’ is central to this Code.” For example, the whistle-blowing clauses (6.11-6.13) describe a software engineer’s obligations when public safety is threatened by defective software development and describe steps to meet those obligations.

The code now contains an open-ended clause (8.07) against using prejudices or bias in any decision making, written broadly enough to include consideration of new social concerns.

Finally, the code includes specific language about the importance of ethical behavior during the maintenance phase of software development. The new text reflects the amount of time a computer professional spends modifying and improving existing software and also makes clear that we need to treat main-

The primacy of well-being and quality of life of the public is emphasized throughout the code.

tenance with the same professionalism as new development. The quality of maintenance depends upon the professionalism of the software engineer, because maintenance is more likely to be scrutinized only locally, whereas new development is generally reviewed at a broader corporate level.

In the same spirit that created the code of ethics, the Computer Society and the ACM continue to support the software engineering profession through the Software Engineering Professionalism and Ethics Project (<http://computer.org/tab/swecc/Sepec.htm>). This project will help make the code an effective practical tool by publishing case studies, supporting further corporate adoption of the code, developing curriculum material, running workshops, and collaborating with licensing bodies and professional societies.

Software Engineering Code of Ethics and Professional Practice

SHORT VERSION: PREAMBLE

The short version of the code summarizes aspirations at a high level of abstraction. The clauses that are included in the full version give examples and details of how these aspirations change the way we act as software engineering professionals. Without the aspirations, the details can become legalistic and tedious; without the details, the aspirations can become high-sounding but empty; together, the aspirations and the details form a cohesive code.

Software engineers shall commit themselves to making the analysis, specification, design, development, testing, and maintenance of software a beneficial and respected profession. In accordance with their commitment to the health, safety, and welfare of the public, software engineers shall adhere to the following eight Principles:

1. Public. Software engineers shall act consistently with the public interest.
2. Client and employer. Software engineers shall act in a manner that is in the best interests of their client and employer, consistent with the public interest.
3. Product. Software engineers shall ensure that their products and related modifications meet the highest professional standards possible.
4. Judgment. Software engineers shall maintain integrity and independence in their professional judgment.
5. Management. Software engineering managers and leaders shall subscribe to and promote an ethical approach to the management of software development and maintenance.
6. Profession. Software engineers shall advance the integrity and reputation of the profession consistent with the public interest.
7. Colleagues. Software engineers shall be fair to and supportive of their colleagues.
8. Self. Software engineers shall participate in lifelong learning regarding

the practice of their profession and shall promote an ethical approach to the practice of the profession.

FULL VERSION: PREAMBLE

Computers have a central and growing role in commerce, industry, government, medicine, education, entertainment, and society at large. Software engineers are those who contribute, by direct participation or by teaching, to the analysis, specification, design, development, certification, maintenance, and testing of software systems. Because of their roles in developing software systems, software engineers have significant opportunities to do good or cause harm, to enable others to do good or cause harm, or to influence others to do good or cause harm. To ensure, as much as possible, that their efforts will be used for good, software engineers must commit themselves to making software engineering a beneficial and respected profession. In accordance with that commitment, software engineers shall adhere to the following Code of Ethics and Professional Practice.

The Code contains eight Principles related to the behavior of and decisions made by professional software engineers, including practitioners, educators, managers, supervisors, and policy makers, as well as trainees and students of the profession. The Principles identify the ethically responsible relationships in which individuals, groups, and organizations participate and the primary obligations within these relationships. The Clauses of each Principle are illustrations of some of the obligations included in these relationships. These obligations are founded in the software engineer’s humanity, in special care owed to people affected by the work of software engineers, and in the unique elements of the practice of soft-

Software Engineering Code of Ethics and Professional Practice, copyright © 1999 by the Institute of Electrical and Electronics Engineers, Inc., and the Association for Computing Machinery, Inc.

Computer Society Connection

ware engineering. The Code prescribes these as obligations of anyone claiming to be or aspiring to be a software engineer.

It is not intended that the individual parts of the Code be used in isolation to justify errors of omission or commission. The list of Principles and Clauses is not exhaustive. The Clauses should not be read as separating the acceptable from the unacceptable in professional conduct in all practical situations. The Code is not a simple ethical algorithm that generates ethical decisions. In some situations, standards may be in tension with each other or with standards from other sources. These situations require the software engineer to use ethical judgment to act in a manner that is most consistent with the spirit of the Code of Ethics and Professional Practice, given the circumstances.

Ethical tensions can best be addressed by thoughtful consideration of fundamental principles, rather than blind reliance on detailed regulations. These Principles should influence software engineers to consider broadly who is affected by their work; to examine if they and their colleagues are treating other human beings with due respect; to consider how the public, if reasonably well informed, would view their decisions; to analyze how the least empowered will be affected by their decisions; and to consider whether their acts would be judged worthy of the ideal professional working as a software engineer. In all these judgments concern for the health, safety and welfare of the public is primary; that is, the “Public Interest” is central to this Code.

The dynamic and demanding context of software engineering requires a code that is adaptable and relevant to new situations as they occur. However, even in this generality, the Code provides support for software engineers and managers of software engineers who need to take positive action in a specific case by documenting the ethical stance of the profession. The Code provides an ethical foundation to which individuals within teams and the team as a whole can appeal. The Code helps to define those actions that are ethically improper to request of a software engineer or teams of software engineers.

The Code is not simply for adjudicat-

ing the nature of questionable acts; it also has an important educational function. As this Code expresses the consensus of the profession on ethical issues, it is a means to educate both the public and aspiring professionals about the ethical obligations of all software engineers.

PRINCIPLES

Principle 1: Public

Software engineers shall act consistently with the public interest. In particular, software engineers shall, as appropriate:

- 1.01. Accept full responsibility for their own work.
- 1.02. Moderate the interests of the software engineer, the employer, the client, and the users with the public good.
- 1.03. Approve software only if they have a well-founded belief that it is safe, meets specifications, passes appropriate tests, and does not diminish quality of life, diminish privacy, or harm the environment. The ultimate effect of the work should be to the public good.
- 1.04. Disclose to appropriate persons or authorities any actual or potential danger to the user, the public, or the environment, that they reasonably believe to be associated with software or related documents.
- 1.05. Cooperate in efforts to address matters of grave public concern caused by software, its installation, maintenance, support, or documentation.
- 1.06. Be fair and avoid deception in all statements, particularly public ones, concerning software or related documents, methods, and tools.
- 1.07. Consider issues of physical disabilities, allocation of resources, economic disadvantage, and other factors that can diminish access to the benefits of software.
- 1.08. Be encouraged to volunteer professional skills to good causes and to contribute to public education concerning the discipline.

Principle 2: Client and employer

Software engineers shall act in a manner that is in the best interests of their client and employer, consistent with the public interest. In particular, software engineers shall, as appropriate:

- 2.01. Provide service in their areas of competence, being honest and forthright about any limitations of their experience and education.
- 2.02. Not knowingly use software that is obtained or retained either illegally or unethically.
- 2.03. Use the property of a client or employer only in ways properly authorized, and with the client's or employer's knowledge and consent.
- 2.04. Ensure that any document upon which they rely has been approved, when required, by someone authorized to approve it.
- 2.05. Keep private any confidential information gained in their professional work, where such confidentiality is consistent with the public interest and consistent with the law.
- 2.06. Identify, document, collect evidence, and report to the client or the employer promptly if, in their opinion, a project is likely to fail, to prove too expensive, to violate intellectual property law, or otherwise to be problematic.
- 2.07. Identify, document, and report significant issues of social concern, of which they are aware, in software or related documents, to the employer or the client.
- 2.08. Accept no outside work detrimental to the work they perform for their primary employer.
- 2.09. Promote no interest adverse to their employer or client, unless a higher ethical concern is being compromised; in that case, inform the employer or another appropriate authority of the ethical concern.

Principle 3: Product

Software engineers shall ensure that their products and related modifications meet the highest professional standards

possible. In particular, software engineers shall, as appropriate:

- 3.01. Strive for high quality, acceptable cost, and a reasonable schedule, ensuring significant tradeoffs are clear to and accepted by the employer and the client, and are available for consideration by the user and the public.
- 3.02. Ensure proper and achievable goals and objectives for any project on which they work or propose.
- 3.03. Identify, define, and address ethical, economic, cultural, legal, and environmental issues related to work projects.
- 3.04. Ensure that they are qualified for any project on which they work or propose to work, by an appropriate combination of education, training, and experience.
- 3.05. Ensure that an appropriate method is used for any project on which they work or propose to work.
- 3.06. Work to follow professional standards, when available, that are most appropriate for the task at hand, departing from these only when ethically or technically justified.
- 3.07. Strive to fully understand the specifications for software on which they work.
- 3.08. Ensure that specifications for software on which they work have been well documented, satisfy the user's requirements, and have the appropriate approvals.
- 3.09. Ensure realistic quantitative estimates of cost, scheduling, personnel, quality, and outcomes on any project on which they work or propose to work and provide an uncertainty assessment of these estimates.
- 3.10. Ensure adequate testing, debugging, and review of software and related documents on which they work.
- 3.11. Ensure adequate documentation, including significant problems discovered and solutions adopted, for any project on which they work.
- 3.12. Work to develop software and related documents that respect the privacy of those who will be affected by that software.
- 3.13. Be careful to use only accurate data derived by ethical and lawful means, and use it only in ways properly authorized.
- 3.14. Maintain the integrity of data, being sensitive to outdated or flawed occurrences.
- 3.15. Treat all forms of software maintenance with the same professionalism as new development.

Principle 4: Judgment

Software engineers shall maintain integrity and independence in their professional judgment. In particular, software engineers shall, as appropriate:

- 4.01. Temper all technical judgments by the need to support and maintain human values.
- 4.02. Only endorse documents either prepared under their supervision or within their areas of competence and with which they are in agreement.
- 4.03. Maintain professional objectivity with respect to any software or related documents they are asked to evaluate.
- 4.04. Not engage in deceptive financial practices such as bribery, double billing, or other improper financial practices.
- 4.05. Disclose to all concerned parties those conflicts of interest that cannot reasonably be avoided or escaped.
- 4.06. Refuse to participate, as members or advisors, in a private, governmental, or professional body concerned with software-related issues in which they, their employers, or their clients have undisclosed potential conflicts of interest.

Principle 5: Management

Software engineering managers and leaders shall subscribe to and promote an ethical approach to the management of software development and maintenance. In particular, those managing or leading software engineers shall, as appropriate:

- 5.01. Ensure good management for any project on which they work, including effective procedures for promotion of quality and reduction of risk.
- 5.02. Ensure that software engineers are informed of standards before being held to them.
- 5.03. Ensure that software engineers know the employer's policies and procedures for protecting passwords, files, and information that is confidential to the employer or confidential to others.
- 5.04. Assign work only after taking into account appropriate contributions of education and experience tempered with a desire to further that education and experience.
- 5.05. Ensure realistic quantitative estimates of cost, scheduling, personnel, quality, and outcomes on any project on which they work or propose to work, and provide an uncertainty assessment of these estimates.
- 5.06. Attract potential software engineers only by full and accurate description of the conditions of employment.
- 5.07. Offer fair and just remuneration.
- 5.08. Not unjustly prevent someone from taking a position for which that person is suitably qualified.
- 5.09. Ensure that there is a fair agreement concerning ownership of any software, processes, research, writing, or other intellectual property to which a software engineer has contributed.
- 5.10. Provide for due process in hearing charges of violation of an employer's policy or of this Code.
- 5.11. Not ask a software engineer to do anything inconsistent with this Code.
- 5.12. Not punish anyone for expressing ethical concerns about a project.

Principle 6: Profession

Software engineers shall advance the integrity and reputation of the profession consistent with the public interest. In particular, software engineers shall, as appropriate:

Computer Society Connection

- 6.01. Help develop an organizational environment favorable to acting ethically.
- 6.02. Promote public knowledge of software engineering.
- 6.03. Extend software engineering knowledge by appropriate participation in professional organizations, meetings, and publications.
- 6.04. Support, as members of a profession, other software engineers striving to follow this Code.
- 6.05. Not promote their own interest at the expense of the profession, client, or employer.
- 6.06. Obey all laws governing their work, unless, in exceptional circumstances, such compliance is inconsistent with the public interest.
- 6.07. Be accurate in stating the characteristics of software on which they work, avoiding not only false claims but also claims that might reasonably be supposed to be speculative, vacuous, deceptive, misleading, or doubtful.
- 6.08. Take responsibility for detecting, correcting, and reporting errors in software and associated documents on which they work.
- 6.09. Ensure that clients, employers, and supervisors know of the software engineer's commitment to this Code of Ethics, and the subsequent ramifications of such commitment.
- 6.10. Avoid associations with businesses and organizations which are in conflict with this Code.
- 6.11. Recognize that violations of this Code are inconsistent with being a professional software engineer.
- 6.12. Express concerns to the people involved when significant violations of this Code are detected unless this is impossible, counterproductive, or dangerous.
- 6.13. Report significant violations of this Code to appropriate authorities when it is clear that consultation with people involved in these significant violations is impossible, counterproductive, or dangerous.

Principle 7: Colleagues

Software engineers shall be fair to and supportive of their colleagues. In particular, software engineers shall, as appropriate:

- 7.01. Encourage colleagues to adhere to this Code.
- 7.02. Assist colleagues in professional development.
- 7.03. Credit fully the work of others and refrain from taking undue credit.
- 7.04. Review the work of others in an objective, candid, and properly-documented way.
- 7.05. Give a fair hearing to the opinions, concerns, or complaints of a colleague.
- 7.06. Assist colleagues in being fully aware of current standard work practices including policies and procedures for protecting passwords, files, and other confidential information, and security measures in general.
- 7.07. Not unfairly intervene in the career of any colleague; however, concern for the employer, the client, or public interest may compel software engineers, in good faith, to question the competence of a colleague.
- 7.08. In situations outside of their own areas of competence, call upon the opinions of other professionals who have competence in those areas.

Principle 8: Self

Software engineers shall participate in lifelong learning regarding the practice of their profession and shall promote an ethical approach to the practice of the profession. In particular, software engineers shall continually endeavor to:

- 8.01. Further their knowledge of developments in the analysis, specification, design, development, maintenance, and testing of software and related documents, together with the management of the development process.
- 8.02. Improve their ability to create safe, reliable, and useful quality

software at reasonable cost and within a reasonable time.

- 8.03. Improve their ability to produce accurate, informative, and well-written documentation.
- 8.04. Improve their understanding of the software and related documents on which they work and of the environment in which they will be used.
- 8.05. Improve their knowledge of relevant standards and the law governing the software and related documents on which they work.
- 8.06. Improve their knowledge of this Code, its interpretation, and its application to their work.
- 8.07. Not give unfair treatment to anyone because of any irrelevant prejudices.
- 8.08. Not influence others to undertake any action that involves a breach of this Code.
- 8.09. Recognize that personal violations of this Code are inconsistent with being a professional software engineer.

Renew Your Membership Online

Renewing your Computer Society membership just got easier. As an IEEE Computer Society member or an affiliate member, you can renew your membership for 2000 on the Web. Starting 1 October, members can access a secure online application form and submit renewal payments via credit card. The service is open to anyone with an IEEE Web account, which full and affiliate members can create online with their member numbers. IEEE Web accounts also allow access to members-only products and services.

To ensure a continuation of services, members in the US and Canada must renew by the end of February. Members in other regions must renew by late April. For complete instructions on renewing online, visit <http://iee.orgmembership/index.html>.

Unlocking the Clubhouse: Women in Computing

By

Jane Margolis and Allan Fisher

The MIT Press, 2002

Chapter 3, Computing with a Purpose

Students' motivation to study computer science varies by gender. For most women students, the technical aspects of computing are interesting, but the study of computer science is made meaningful by its connections to other fields. Men are more likely to view their decision to study computer science as a "no-brainer," an extension of their hobby and lifelong passion for computing.

The women in our study were the survivors of the "boys' club" of high school computing. Some of them self-identified as "girl geeks." We interviewed the former president of her high school's computer club and a student who took pride in being the computer "genius" in her family. Most of these women decided to major in computer science because they did well in a high school class, they found computing came easy to them, and they derived pleasure from it. Almost every woman in our sample is a self-described "math and -science person" who enjoys problem solving, doing puzzles, and doing logical thinking tasks. They talk about "the rush in having my program run" and about being at the cutting edge of technology invention. Yet even among these computing-oriented women, we heard about values and preferences that were distinct from those of most male computer science students. This distinction and its ramifications play a key role in the experiences and perceptions of women in computing.

The Decision to Major: The Passionate and the Rational

We have found that women decide to major in computer science based on a broad set of criteria. The simple enjoyment of computing is a leading factor for women, but other factors also weigh heavily in their decisions. They value the versatility of computing, its relation to their interests in math and science, its career path to safe and secure employment, the exciting and changing nature of the field, and the encouragement they received from parents or teachers.

For many male students, in contrast, the decision to major in computer science barely reaches the level of conscious consideration; it is a natural extension of their lifelong passion for computing. Aaron's response to our query about his decision to the major is quite typical: "I don't ever recall really deciding.... As long as I've known that there was such a thing as majoring in computer science, I just basically assumed that I was going to be majoring in computer science." When the interviewer asked him if he had ever considered another major, he replied: "Not really." Like Aaron, many of the men in our study are convinced of a perfect fit, while more women describe their decision to major in computer science as "checking it out."

The differences between the choices made by men and women emerge strikingly from our tabulation of their reasons for majoring in computer science, shown in figures 3.1 and 3.2. These graphs tabulate the percentage of each group who mentioned the associated factor as a motivation in their decision. Most women take a large number of factors into account: five of the seven categories we tabulated were mentioned by at least 30 percent of the women. In contrast, the only motivation listed by at least 30 percent of the men (in fact, by 70 percent) is the enjoyment of computing.

Versatility of the Field

The versatility of computing is a big draw for women students. As Karina says, "You can do anything with a computer science degree," and "almost any field is computer-related now." She wants to learn skills that will help her get a job and pay back some of her college loans. Then if she likes computers, she can stay with them. Contrasting her job possibilities in biology with those in computers, she calculates:

There is not very much money in biology on the undergraduate level. If you ever want to go to grad school, you could get somewhere from there possibly. But computers are so versatile that basically you could do anything with a computer science degree.




Figure 3.1
Factors in women's decision to major ($n = 32$)

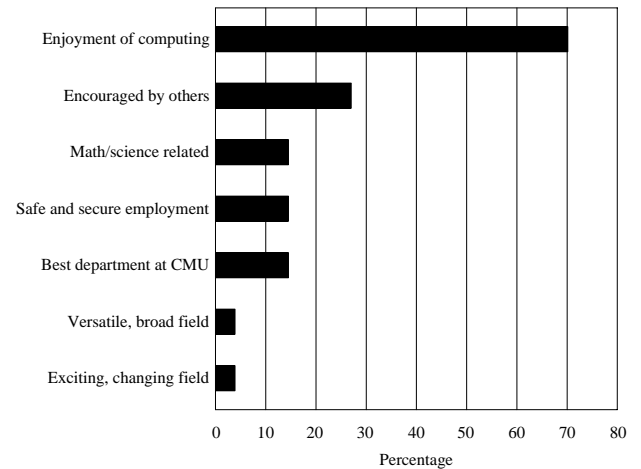


Figure 3.3
Factors in men's decision to major ($n = 27$)

Another Carnegie Mellon student, Andrea, worked with Intel computers in high school. She enjoyed that experience and “caught on very quickly to what was going on.” But what really influenced her decision to major was “thinking of the future and what I wanted to do, and it seems that computers are one of the most influential pieces of technology in our world at this time ... And I thought that if I get into computers, then I’ll definitely get a job.” Andrea is not exactly sure what she wants to do when she gets out of college. Her heart is in film, but she knows that it is a risky business and wants something enjoyable to “fall back on.” She figures that she “can use computer graphics and incorporate that in my film production.”

Computing with a Purpose

Besides having a broader set of criteria for majoring in computer science, many women have interests in computing that go beyond the technical aspects. Connecting computing to other fields and working within its human and social contexts makes the study of computer science compelling and meaningful for them. For instance, Deborah wants to use computing to study diseases to “solve the problems of science”:

I think with all this newest technology there is so much we can do with it to connect it with the science field, and that [studying diseases] is kind of what I want to do ... —use all this technology and use it to solve the problems of science we have, the mysteries.

Laurette is interested in the links between computing and “the most efficient way” to use computers in education. Phyllis, a first-year woman student, describes a difference between herself and most of the male students. She says her male peers are focused on “building bigger and bigger computers.” “That’s fine,” she says, “and I’d like to be involved with that, too, but in the long run I want to use computers for what they are now and just use them to help people.” She has been inspired by a recent Carnegie Mellon lecture about a robot car that promises to save lives by reducing the number of accidents and deaths caused by human error. Phyllis is determined to not let herself get detached from society but instead to connect computer science to real-world problems:

The idea is that you can save lives, and that’s not detaching yourself from society. That’s actually being a part of it. That’s actually helping. Because I have this thing in me that wants to help. I felt the only problem I had in computer science was that I would be detaching myself from society a lot, that I wouldn’t be helping—that there would be people in third world countries that I couldn’t do anything about. ... I would like to find a way that I could help. That’s where I would like to go with computer science.

Another first-year student, Louise, describes a difference she felt between herself and her male peers when she saw her peers' nonchalant response to a lecture on ways that computers can be nonproductive in society:

Everyone just said how boring it was: "Who cares that computers did not benefit anyone? We like computers! We love computers! We know computers! And who cares about the rest of the world"

She describes herself, on the other hand, as someone who scrutinizes the worth of each computing project in terms of what it is doing to change and help the world:

And it you're trying to make something that's going to change the world, that's going to help the world, you have to have some sort of concern about what's your long-term goal. Not just to produce Word 8 ... or Excel ... whatever. How is this helping? Or is it helping? Go see if that stuff is doing anything.

Jessica, a woman student who has always done well in math and science, feels deeply that computer science must "make-a contribution": "just ... making video games" is not "worth the energy and talent that it takes," She relates her interest in computer science to her concern for her grandmother's medical condition:

I don't think science—just for making video games—is worth the energy and talent that it takes, but I think it's important if it makes a contribution. So part of that would be a contribution in medicine. My Grandma had a pacemaker, a renal dialysis machine. ... I've seen the contribution in my family in my life. ... Medicine has always fascinated me, so I just always wanted to apply my sciences there. And I see the opportunities now, with the computer technology to apply there and that's what I want to do.

Forty-four percent of the women we interviewed and nine percent of the men link their interest in computing to other arenas. We refer to this orientation as "computing with a purpose." This is not to say that the men lack interests outside computing. Sam, for instance, is interested in music and computing. But when he describes his interest in computing, he says "it is the code itself that is interesting, even more so than the actual effect it has." He compares music to programming this way:

Music as an art form is similar to programming as an art form; it's something you can sit down and within a day you can be doing something which has an essence of beauty to it... . That process I still find very interesting—how to say what you want to say—so that the code in itself is what is interesting, even more so than the actual effect that it has.

Christopher, a first-year student, relates his attraction to computing to his feeling that "when you write a program, you don't have to worry too much about outside factors influencing how your program is going to run." He contrasts computing to physics experiments, where in physics "you have to take into account gravity and air friction and all these things." For Christopher, computing is "an environment where you can do what you want to do, and you have more control over how you're going to do it." Furthermore, Christopher notes with delight that "you get to do cool things and play around with it and it's fun."

Women's Counternarratives

Women students' descriptions of why they are majoring in computer science are a "counternarrative" to the stereotype of computer scientists who are narrowly focused on their machines and are hacking for hacking's sake. Instead, these women tell us about their multiple interests and their desire to link computer science to social concerns and caring for people. These women may or may not qualify as "people people" on a psychological inventory exam to the same degree as those involved in nursing, social work, or child care, but they need their computing to be useful for society.

Our finding that women students bring contextual concerns to the study of computer science resonates with what feminist psychologists write about women's development. Jean Baker Miller (1976), in her landmark study of women's psychology, *Toward a New Psychology of Women*, writes:

It is true that women, like everyone, are motivated out of the well-springs of their own being. In that sense, we all, at bottom, act on what is moving us individually. It is also true, however, that women feel compelled to find a way to translate their own motivations into a means of serving others and work at this all their lives. If they can keep finding ways to do this, they are often comfortable and satisfied—and they do thereby serve others. This translation of

motivation accomplishes an integration that is significantly different from the integration that society encourages in men. In fact, our society specifically discourages men from even attempting anything like this. (p. 63)

Our findings also align with much research on gender and math and science that has found women students bringing contextual concerns to their study. A metaanalysis of research on gender and science by Marcia Linn and Janet Hyde (1989) concluded that a major sex difference in interests in math and science is its perceived usefulness. Schofield's (1995) ethnographic study of the introduction of computers into a high school found that a number of male teachers reported doing things such as building computers for fun or deciding to teach computer science out of a fascination with the subject that led them to switch fields. Not a single female teacher, on the other hand, was fascinated with the computer in this same way. Rather, the female teachers who responded positively to them "tended to speak about their actual or potential usefulness" (p. 161). University of Michigan researcher Jacquelynne Eccles (1994) reports that the Michigan Study of Adolescent Life Transitions, a longitudinal study of approximately 1,000 adolescents from southeastern Michigan, found that "women select the occupation that best fits their hierarchy of occupationally relevant values" and that helping others and doing something worthwhile for society is high in that hierarchy (Eccles, 1994, 600).

Researchers Honey et al. (1991) at the Bank Street College Center for Children and Technology asked twenty-four adult technology experts and eighty early adolescents, approximately equally divided by gender, to describe a science fiction story they would like to write about a perfect computer. They found that the women experts designed their machines to be people connectors, for communication and collaboration. They placed the technology in the context of human relationships. The men experts, in contrast, envisioned technology as "extensions of their power over the physical universe. Their fantasies were often about absolute control, tremendous speed, and unlimited knowledge" (p. 3).

The girls designed household helpers, machines that offered companionship, or devices that they could use to broaden their social and personal networks. The boys, on the other hand, fantasized about extensions of power, often imagining technology that could overpower natural constraints. Besides noting women's concern with relationships and men's concern with power, Cornelia Brunner (1997), a member of the research team and a longtime investigator of the relationship of gender and technology, observed that "The feminine take on technology looks right through the machine to its social function, while the masculine view is more likely to be focused on the machine itself" (p. 55).

Which Orientation Does the Curriculum Support?

In most computer science programs, the early semesters are narrowly focused on the technical aspects of programming, and applications and multidisciplinary projects are deferred to the end. This, unfortunately, gives beginning students, male and female, the false message that computer science is only "programming, programming, programming" and removed from real-world context and concerns. In this context, writes the American Association of University Women Technology Commission (2000), the "cultural emphasis on technical capacity, speed, and efficiency" dominates the scene, and this culture "estranges a broad array of learners," especially women (p. 7). Traditional computer science curriculum and programming assignments often lack the larger interdisciplinary framework that women find important.

Feminist educator Sue Rosser (1990) argues that "insuring science and technology are considered in their social context with assessment of their benefits for the environment and human beings may be the most important change that can be made in science teaching for all people, both male and female" (p. 72). Computer science professor Dianne Martin (1992), in her article "In Search of Gender-Free Paradigms for Computer Science Education," speculates that an integrated approach to computer science would attract more female students. She suggests that "greater attention to values, human issues, and social impact as well as to the mathematical and theoretical foundations of computer science" would redress the balance (p. 1).

Computer scientist Frances Grundy (1998), in her article "Mathematics and Computing: A Help or Hindrance for Women?," explores what lies behind the traditional orientation of computer science. She argues that "pure" computing (such as analysis of algorithms and complexity theory) has historically been considered more prestigious than applied computing because male theoreticians, who are the "inner circle" of computer science, define what is "real computer science." She believes that while "in fact, everything done on a computer requires some abstraction ... abstraction by itself is not enough; we must be able to set

the results of our task back into the real world” (p. 5). Since the majority male “insider” group finds particular value in abstraction and in the joy of playing at the computer, these become dominant parts of the computing curriculum and culture. Thus, many women are left questioning whether computer science has a place for them and whether their orientation will allow for success and a sense of belonging in the discipline.

A Broader View of Programming

Building on the joys that women find in computing is an important task for narrowing the gender gap in the field. Historically, computer science has been associated with number-crunching and quantitative skills to the neglect of more varied thinking and creative activities. In their article “Mismeasuring Women: A Critique of Research on Computer Ability and Avoidance,” authors Pamela Kramer and Sheila Lehman (1990) argue that much of the research on the gender gap in computing has mimicked the research on women and math (since computers have been linked to math). In doing so, this research has neglected to look at the larger educational context of how computers are used and how computer learning is different from math learning:

As the use of computers expands in educational and workplace settings, the contexts and applications of their use are rapidly changing so that the presumed closeness of the domains of computing and mathematics knowledge constitutes an increasingly inaccurate portrayal of what experienced and highly skilled computer users describe as being the most advanced types and forms of creative computer related work. (p. 170) .

They argue that, instead, “creative computing now relies at least as much upon language, visual design, problem definition, and organizational skills as upon quantitative analysis” (p. 171).

We, too, found that both women and men find pleasure in areas beyond the traditional quantitative programming lexicon. Here are some of the women’s answers to our question “What do you enjoy about programming?”

It seems like solving a giant puzzle. It feels like someone’s giving you a Rubik’s cube and you have to figure it out. It’s something you can think about, play with, and do weird things that no one ever thought of.... It’s still a challenge that is not in another field that I’ve found, with possibly the exception of genetics, which is also like a puzzle.

It gives me such a great feeling to have my own little creation. That’s the thing about CS that I do really enjoy. It’s kind of like an art in itself ... like it’s an art of thinking. And everyone has their own little way of doing an algorithm or something. So that’s what I enjoy.

I just like being interactive with a computer. You’re almost playing with the computer and making it do things. You can actually see progress before your eyes ... when you get it to work. ... It’s interaction.

And then at some point I found myself just enjoying the feeling of getting something to work ... of communicating with this machine and taking my thought, which was not necessarily organized, and turning it into ... an abstract, organized mathematical language that something as stupid as a machine could understand. It’s kind of like you’re given a problem and you have to sit down and solve it. That’s your mission, and you do it, and it works, and when it works it’s the most amazing feeling in the world!

In a similar vein, here are some answers from men:

Programming? There’s that really good sense of accomplishment, the most positive feeling you get when it works. When your program finally works, and you throw your hands up in the air and say, “YES, IT WORKS!” And I just really enjoy problem solving in general and the sort of logical approach to it and the way of thinking that’s behind programming.

The part I enjoy about computer science is building things and watching them run. ... It’s just such a great feeling to have created something. ... In computer science I get a very complete feeling about the wonder of something I’ve created out of thin air. All my ideas are put together into one thing.

Whether or not people think so, I think that programming is a way of expressing yourself. ... It’s not just a tool. ... A person that’s knowledgeable in computers can do a lot more with computers than you would normally think is possible.

Love of puzzles, creating something from nothing, the art of thinking, interaction, communication: all are facets of the computing endeavor but not always part of the traditional lexicon.

A tally of the factors that students like about programming (table 3.1) shows a few quantitative differences between women and men. Women are more likely than men to cite the satisfaction of getting their program to work, the problem-solving aspects of programming, and the challenge. Despite Kramer and Lehman's (1990) critique, many also report their enjoyment of programming involves its mathematical aspects. The area that men cite significantly more often than women is the creativity of the process. It is not clear to us whether this reflects an independent gender difference, is linked to a primarily male sense of play, or is an effect of the men's overall greater experience. A larger sample of women exhibiting substantial experience or a play orientation toward computing would be needed to draw a conclusion.

Table 3.1
Attractions of programming

	Men (<i>n</i> =35)	Women (<i>n</i> =37)
Satisfaction in success	40%	57%
Creativity	34	20
Control	31	26
Problem solving	20	49
Challenge	9	43
Ownership	9	14
Logic	9	14
Math related	6	29

Programming is fundamental to computer science. Listening to students explain what they enjoy about programming raises the possibility of a wider range of approaches than are currently used to explain programming. It could be that a more expansive set of metaphors could help to refute the stereotype of programming as a dry, mathematical endeavor and engage the interest of a wider range of students.

Summary

IBM (2000) ran a series of magazine advertisements in an effort to recruit a more diverse workforce. One advertisement describes the job held by Grace Liu, educational strategist: "Work with schools to develop and implement tech-based solutions that improve the way students learn." Her experience is listed as "Currently helping San Francisco's school district design a case management tool that prevents students in need from getting lost in the shuffle." This ad stands in sharp contrast to more common computer-related ads that feature young white males in a state of game bliss, focused only on what is on the screen.

Can a creative person, a "people person," care about the world and people and be happy in computer science? While the stereotype says no, a broader vision of what the field is and how it is best taught answers in the affirmative. Computing can be taught in an interdisciplinary setting, honoring the goal of "solving the world's problems." Furthermore, this does not require devaluing the single-minded pursuit of technical virtuosity that marks some of the best computer science students. Instead, it establishes multiple standards of excellence, which together can yield a stronger community of computing professionals than any one by itself. The perspective that computer science can make itself stronger by incorporating the values typical of women in the field changes the question from "How can women change to fit into computer science?" to "How can computer science change to attract more women?"

Chapter 6, Persistence and Resistance: Staying in Computer Science

During our research we were often surprised by which students stayed in the program and which left. Especially in the first two years, many women ride an emotional roller coaster of certainty and doubt from term to term, indeed from week to week, and whether they decide to finish the ride or get off before it ends

is unpredictable. Although we interviewed students each semester, students' decisions to leave the program or to stay surprised us more than once.

Paula, for example, began the program excited, enthusiastic, and confident. She had completed a summer internship at one of the local computing labs and was enthusiastic about majoring in computer science. But not long after her arrival, she began to have doubts about her interest and abilities and started talking about leaving. The following semester she told us she had decided to stay, was happy in the program, and was sure she would continue. In her third semester, she told us she had decided to transfer out because "it just isn't worth it" any more.

As often as we were unprepared when women who seemed happy left, we were also sometimes surprised to find them staying. In *Talking About Leaving: Why Undergraduates Leave the Sciences*, E. Seymour and N. Hewitt also refer to this back-and-forth dynamic of students' decision-making process. The one thing that did become predictable was timing: students would most likely leave in the sophomore year, the time when most students, across all majors, do their switching.

What determines whether a woman chooses to stay in or leave computer science? In this chapter we look at what we call the pillars of persistence—the qualities, experiences, and strengths that allowed the women we interviewed to persist despite doubt and uncertainty. We are particularly intrigued with the counterintuitive stories of some of the women students. While a segment of the female persisters resembles the majority of men in certain ways, the portraits of many successful majors run contrary to expectations and assumptions about who can and will succeed.

The Expected: "I Have Always Been Around Computers"

One may intuit that women who persist are likely to come from backgrounds similar to many of the males: computer-intensive families, lots of parental support, a fair share of hands-on experience, fascination with computers. But one of the most surprising findings of our research is that the backgrounds of the women persisters varied wildly. Brenda is someone whose background is similar to what we've described, except that her family includes female role models. She describes her family as "basically a whole family of nerds." Brenda has had computers in her house since she was in kindergarten. The whole family used them, and they often had several going at once. As a result, she says computers and her interest in them are "natural" to her. Brenda's dad is an engineer, her mom is a librarian, and her sister is studying computer science at MIT:

So I've always been around computers, and it's just ... natural to me. Even when we first had an Apple, they'd [parents] encourage me to just pick up stuff and try around. ... We'd do it cold—do it without a disk—and I started programming in Apple Basic, just very simple stuff, and it got me interested in it. So everything else later just came naturally that I wanted to learn about.

Brenda's family didn't watch TV much, and computer games were her entertainment. She "dabbled a bit in Apple Basic to see what fun stuff I could do," learned word processing, and did her science projects on the computer. Her parents have lots of computer-literate friends, and when they visited, they would all play computer games together. In junior high, Brenda started getting involved in the Internet through her sister and mother. She helped run bulletin boards. She also had friends who used computers, though not as much as she.

Brenda's "family of nerds" helped her sense of fit and belonging in computer science. Computers were part of her furniture; they became "natural" to her. Perhaps unsurprisingly, Brenda describes her decision to major in CS as a "kind of a default." She had a wide variety of interests, from music to math to writing, "so it was kind of a toss-up of what I really wanted to do." But she decided that she was "probably the most comfortable around computers in general." She adds, "I'm not sure exactly what area I want to go into. I only know ... I like computers. So that's a good place to start." Even though she found many of her classes very challenging, Brenda is satisfied with how she did in all of them. She enjoys learning to write code. She says, "I know how to think like a programmer." But she adds, "I'm also not a super-genius or anything."

Family make-up emerged in our study as worthy of further investigation. As in Brenda's case, we repeatedly heard women with no brothers attribute their interest in computer science to this fact. While we do not have enough data to draw a firm conclusion, we heard many reports of boys claiming the title of "family computer wizard," with this spot seemingly opening up for a girl in families with no brothers.

The careers and interests of a student's parents also have a major influence on whether a woman pursues an interest in science or engineering. Not only do women with parents in technical occupations pick up language and concepts around the dinner table, but the intimidation factor decreases, and parental mentoring and encouragement increase. The impact of parents is documented by Paula Rayman and Belle Brett's (1993) Pathways Project, a longitudinal research effort at Wellesley College that investigated the experiences of young women in science and mathematics during their undergraduate, graduate, and early career years. Rayman, Brett, and their colleagues found that parental support is one of the pivotal factors that distinguish women who go on to science careers from those who do not.

Coming from a computing or engineering family certainly provides important emotional and intellectual stepping stones for majoring in computer science, but our research shows that it is not required. Forty-eight percent of the persisters we interviewed did not come from "computing families." These students' stories provide us with an opportunity to find other stars in the constellation of persistence.

The Counterintuitive Persisters

Some of the most fascinating stories of persistence were told to us by women students who had absolutely no computing experience in their family background. These were mostly international students, raised and educated primarily in countries other than the United States. (Approximately 30 percent of the female computer science majors at Carnegie Mellon during the course of our study have been international women—primarily from Asia and Eastern Europe.) Their motivations for choosing computer science, along with their lack of computer experience, make them the antithesis of the "computer-obsessed since childhood" stereotype. In fact, many of these women were only marginally interested in the field when they began the program.

From their experiences, we learned that despite the tremendous range of computing experience among students, women who are complete novices are no less likely to persist than the most experienced women. Their stories show us that prior computer experience does not make the critical difference. The portraits of these students fly in the face of expectations and assumptions about who can and will succeed in a competitive computer science program.

Little Experience and "No Choice"

Kanitha was a junior from Thailand. As one of ten children, her parents could not afford for her to attend university in Thailand. She came to the United States for high school, where she took her first computing class. Her decision to major in computer science was not based on prior experience or love of computing. She told us about her completely pragmatic and in some ways very uninformed, decision to major in computer science:

Actually, I came from Thailand, and basically I hadn't dealt with any computer at all before I came. And after that I got a scholarship to study computer science, but I didn't know anything about computer science. And then I went to high school here, and then I started taking a course about computer programming, and it was kind of interesting. But then I mean, I have no choice, so that is why I am doing computer science.

Kanitha came to Carnegie Mellon on a corporate scholarship, which requires her to return to Thailand after graduation and work for her scholarship sponsor. She is very clear that the chance to study abroad is most important to her; what she studies is secondary. She eventually decided to choose computer science as a major over electrical engineering because the best scholarship offered was from the Bank of Thailand, which wanted computer science majors. ... When asked, "How did you end up getting a scholarship to study computer science with no computer background?" she answers, "I just want to study abroad, so anything is fine with me." Kanitha has been an extremely successful student at Carnegie Mellon and is considering graduate school in computer science.

"You Have This Bridge You Have to Walk Over, and You Just Don't Look Down"

In another set of accounts, we hear how the pressing need of many international students to become breadwinners for their families leads them to pursue economic opportunity over personal interest. Concern for their families motivates them to stick it out and work hard despite doubt and lack of confidence.

Larissa, for example, moved to the United States with her family from Russia two years prior to attending Carnegie Mellon. She learned English while attending an American high school for two years. While Larissa had more prior computing experience than did Kanitha (she used to play computer games with her dad), she had little experience in comparison to either men or women from the United States. Throughout her four years at Carnegie Mellon, Larissa consistently ranked at the top of her class. She was thoughtful in reflecting on her experiences learning to live with the computer culture, accepting how little she knew compared to the peers around her.

Larissa described her first two years as walking over an “abyss.” It was very difficult for her, and she frequently doubted herself:

You have this bridge you have to walk over, and you just don’t look down. ... There were cases when I started looking down, and it was really scary. I’d think, “WHY am I putting myself through this?” ... But I have to do this, anyway, because I have to.

Larissa felt there was no option for failure, since her entire family was counting on her for financial support. Her father had been a research scientist in Russia, but in the United States has been managing a small restaurant. Her brother’s ability to go to college depends on the money she will make after graduation. She has no financial safety net beneath her and feels she must persist. She believes that “you cannot have everyone doing what they want to do,” that there is “supply and demand with jobs and what needs to be done,” and that “basically, we have to do good to stay here.” And she adds, “It’s just a matter that if I’m doing something, I have to be good at it, so ... you just work hard.”

Degrees of Freedom

Motivations like these can boost persistence of students, even in less than ideal circumstances. Seymour and Hewitt (1997), in *Talking About Leaving*, speculate that “gender differences in perceived degrees of freedom to choose and to change direction” lead more women than men to leave the sciences (p. 278). They suggest that especially among students from socially and economically advantaged backgrounds, women choose disciplines “largely by the degree of personal satisfaction they offer” and “pay less regard to their economic viability” (p. 279). The result is that when the math-science tightrope becomes culturally or academically uncomfortable, women with safety nets may jump: “Reports of relatively easy release from initial commitment to a science, math, or engineering major were most common among women from economically advantaged families” (p. 278). On the other hand, Seymour and Hewitt found that black women, older women returning to school, and women from working-class families did not feel the same degree of freedom. We found this also to be the case with many of the international women students.

We do not advocate that women forgo personal happiness and sacrifice academic pleasure in the interest of expediency or financial incentives; rather, we are pointing to ways that motivations can affect persistence. But what also is required is a strong sense of self-efficacy. From interviews with these counterintuitive persisters we were able to identify several “pillars of persistence” that help boost student’s sense of self-esteem.

Attributional Beliefs about Intelligence and Talent

Research on learning motivation based on U.S. students has found that students generally hold one of two opposing views on intelligence. One view is that intelligence is a fixed trait—as in “you are born with the talents that you have, and nothing you do can change them.” Students who hold this view tend to focus on performance issues such as grades and other forms of external approval. The other view holds that intelligence is a malleable quality—as in “if you work hard and practice, you will improve.” These students tend to orient toward learning goals such as improvement and developing mastery.

Which of these dueling views a woman in computer science holds can make a difference in her sense of self-efficacy and persistence. The research of psychologist Carol Dweck (1986), who studies learning motivation, shows that “a focus on ability judgments can result in a tendency to avoid and withdraw from challenge, whereas a focus on progress through effort creates a tendency to seek and be energized by challenge” (p. 1041).

Believing in the link between effort, hard work, and success seems to be the mantra for many of the international women students. A woman from Thailand, in describing her first-year experiences, credits

hard work for her success:

I know it's hard, it's really hard, because I remember my freshman year. I want to give it up because it's hard. But then I thought, "That's a loser's talk." So then I should try it and work hard. I think I can do it.

An Indian student attributes her persistence to "lagan," a Hindi term akin to "putting your nose to the grindstone." Using an example from Indian math education and its routine disciplined drills, she connects her cultural and educational training to her success in computer science:

But that routineness, I think, is something that isn't taught enough here. ... And so people here have, from my experience with my classmates, I see they have a lot of insight, a lot of intelligence. ... You know, they [snaps finger] pick things up as quickly, but they don't have the grit to sit down with something for, say, six hours and say, "All right, I'm going to get this done no matter what."

When we ask Larissa what factor she feels contributed most to her success, she tells us, in no uncertain terms, that it was "hard work." She believes that despite knowing less than other students, she will catch up and succeed by working hard.

Culturally Inscribed Attributions of Success

Psychology professors Harold Stevenson and James Stigler (1992) have conducted a cross-cultural examination of beliefs about achievement. Their research aimed to figure out why American children seem to be forever losing educational ground compared to children in some Asian countries. In their book *The Learning Gap: Why Our Schools Are Failing and What We Can Learn from Japanese and Chinese Education*, they examine the organization of schooling and the practice and profession of teaching. They also look at attributions of success and show how these beliefs are culturally inscribed.

Stevenson and Stigler (1992) consider the prevailing philosophies in Asian cultures and note that Confucian philosophy promotes the belief that lack of achievement is due to insufficient effort rather than to a lack of ability or to personal or environmental obstacles. In other words, a person who works hard will master a task. Many Asian students grow up hearing adages like those of Chinese philosopher Hsun Tzu: "Achievement consists of never giving up. ... If there is no dark and dogged will, there will be no shining accomplishment; if there is no dull and determined effort, there will be no brilliant achievement" (p. 97).

In elementary schools throughout China, young children hear parables instructing them to work hard, put in the effort, and learn. One such tale is about Li Po, a poet who walks by a small stream and sees a white-haired old woman who has made a needle from a rock. The woman advises Po: "All you need is perseverance. If you have a strong will and do not fear hardship, a piece of iron can be ground into a needle." Other sayings and mottos convey the belief in hard work and effort, such as "The rock can be transformed into a gem only through daily polishing," and "the slow bird must start out early" (Stevenson and Stigler 1992, p. 98).

Suzuki, the early childhood educator who introduced a now world-famous method of teaching the violin to very young children, had a similar philosophy about children's learning. Teaching violin to young children is not a question of seeking out the naturally talented. Suzuki (1978) believed that all children, with daily practice and hard work, could learn to play the violin. A boy or girl does not have to be a child prodigy to learn to play very young. Suzuki's teaching model compares violin playing to language acquisition: it happens through regular practice and repetition at a very young age.

Jane has read her daughter the story of Lilia, the 1996 Olympic gymnastic gold medalist from the Ukraine. In the official version of the Ukrainian gymnastic federation, Lilia is not a "natural" gymnast. Her hands are too small for the bars, and her back is weak. But Lilia's coaches recognize her determination—"a will to win and work exceptionally hard." Almost every section of the book repeats this refrain. The book also describes how it takes a team effort of Lilia, her coach, and her choreographer to win the medal. None of them could do it alone. Rather than the single famous star, the book is about a team that works hard until it wins.

Hard Work Versus the "Computer Gene" Theory

When faced with difficult course work, American women also work hard—very hard. Yet too quickly they

hit bottom, concluding that they lack the “natural and innate talent” with which the men seem to be born. Lily, a U.S. student who was full of enthusiasm when she began a year ago, in her last interview questions whether she should be in the program:

I don't really feel like I should be in the department. What am I doing here? So many other people know so much more than me, and this just comes so easy to some people. ... It's just like there are so many people that are so good at this, without even trying. Why am I here? Do I want to work my butt off for four years, when there are so many people that it comes naturally to? Should I be here for the sake of the field even? You know, someone who doesn't really know what she is doing?

Lily ultimately despairs, concluding that no amount of practice or time spent on a task could improve her mastery of the material. As another female student says:

There are people who are born to do this, and I am not one of them. And it's definitely not one of those things that, like, “Oh, with practice, you will become one who is born to do it.” ... I think a lot of people are just born with it. You just gotta be like, “Computers! Yeah! They are awesome!! They are my life!” You know, a lot of computer scientists, that's all they do.

We continued to hear this refrain, as women looked around and experienced their male peers knowing more and doing the work with greater ease. We have found too many American women fall victim to the “computer gene theory,” even if unconsciously.

Gender and the Entity View of Intelligence

In her article “Motivational Processes Affecting Learning,” Carol Dweck (1986) suggests girls may be more likely than boys to subscribe to an “entity” view of intelligence—seeing ability as a fixed, static trait—and therefore exhibit a tendency toward low expectations, challenge avoidance, and debilitation under failure. She describes a series of studies by Leggett who assigned a novel “concept formation task” to bright junior high school students (Leggett 1985). Researchers observed a greater tendency of those girls who subscribed to the “entity” view to avoid challenge.

The entity view of intelligence can take its toll even on a student who works extremely hard. We witnessed how a student who attributes her math success to hard work rather than ability can have low expectations for future success precisely because she thinks her future courses will be even more difficult and demanding than the ones in which she is currently enrolled. A top student in her class reasoned that her A's were the result of hard work, not ability; in her view, others got A's without working so hard. Despite her 4.0 average, she ended up leaving the major, convinced that she was ill-suited for the field because she put in so much effort.

Cultural Resistance

In chapter 4, we discussed how the male hacker has become the cultural norm in computer science, the standard to which women students begin to compare themselves. We have found that women who persist are those who find a way to get grades they are satisfied with and who can develop a personalized view of computing and their place in it. Women who accept the prevailing culture as the norm and who continuously compare themselves to this norm and find themselves coming up short are the ones who suffer the most.

The majority of women struggle to find a place where they can feel comfortable in the prevailing culture. One female student told us how she has refused to conform to the image of the myopically focused “computer geek.” ... When the interviewer asks her if she feels any need to conform to the culture around her, she answers:

I refuse to. I was worried what if I don't. Will I need to conform to that? Will I need to read books on computers all of my free time or something to survive here? And I feel so far I haven't. I'm getting really good grades without that ... without changing myself. So I feel much more confident now that I don't have to. It's kind of nice. I can prove them wrong or something.

Ironically, it is in this area of relationship to the culture that the international women may have an

edge. The international women do not as readily use the U.S. male hacker as their reference group. Since they are not fully part of this culture, their reference group is elsewhere. Many international students have alternative success norms and social bonds that protect them. Other priorities are dominant, and with these come other scales for self-evaluation.

It is important to note that some women students do feel the prevailing culture is a relatively good fit for their interests and personality. They take pleasure learning to walk the walk and talk the talk; becoming part of this culture helps them persist. An American female student talks of a sense of mastery when she became familiar with computer science (CS) jargon: “It kind of feels like becoming part of a club—CS club.” She observes that her new adopted lingo may not be required but that “it is what you grow into:”

I’ve had several friends who are walking along the sidewalk and make a joke and say it in code. It’s something that non-CS people or maybe an arts person would just think is totally stupid, but we think it’s funny. It comes naturally.

Another woman reports, *à la* Star Trek, that “resistance is futile” and takes pleasure in the thought.

Breaking the Isolation and Building Support

“Surround yourself with supportive people!” is the mantra of a current American graduate student who attended Carnegie Mellon as an undergraduate. She attributes her undergraduate survival to the support she received from her family and friends. She recently tells of being the only woman in her lab in graduate school. She didn’t mind that except that there was a “guy in the lab who was a sexist pig, to put it nicely.” She describes the support she got from the other students in the lab:

But the best part of it all was that any remarks he made would be stifled by the other men in the lab. I had good friends! They were shocked at this guy, and he shut up (and thankfully left school) eventually.

Rebecca, a junior, tells us that her boyfriend, “can’t really help me with my assignments, but he’s good moral support.” She describes him as “one of those people who, when I am saying ‘I can’t do this assignment anymore!’ he’s like, ‘Yes, you can. I know you can. I’ve seen you do these things before!’ ”

Vera, a junior, talks about the support she received from a computer science women’s dinner. She begins by describing her earlier social isolation, being one of a minority of women in the midst of male bonding:

Being female is scary in this program. First you feel alone, and you don’t know who to go to, and you don’t know who to talk to. You just feel weird because you see the immediate bonding between other people, just male bonding ... just showing off and talking. ... I can still get intimidated easily. And you just feel like you’re in a minority. It’s just a weird feeling.

She then describes how her self-doubt turned around when she attended her first dinner for computer science women students. She realized that she was not alone:

I had all those feelings, and I didn’t think that anyone shared those. I remember we had a CS dinner with the women in grad school. And it helped me a lot because I wasn’t talking, but I was listening, and I thought everybody was saying the exact same things that I was feeling ... like everybody was talking about them. And it was a big relief for me to realize that actually other people, other females were feeling the same way. And I just felt so much better. I remember after feeling ... it was such a big relief.

Supportive Learning Communities

Salina grew up in Malaysia and has ten brothers and sisters. Both her father, a forester, and her mother, a housewife, were computer illiterate. She attended a boarding school and was in the “science track.” When she arrived at Carnegie Mellon, she “knew a bit about Basic, and I had never really done any hard programming work at the time.” She rather her preparedness at the time of beginning Carnegie Mellon as two on a scale of five and had low confidence. By her junior year, she rated her preparedness as a four and her ability as a three.

Salina describes her first year as a “really hard year for me.” Her confidence was low, and “I see

all these other students just grasping the concept in less time than I could.” She sat in class, feeling lost and “in shock,” feeling that maybe she couldn’t make it. She says, “I was just totally scared at the time.” But she says, “just by working harder I eventually caught up with the whole class, and I ended up getting an A in the class.”

Salina attributes her success partly to the support she received from friends. She said that everybody was just helping each other out. In her second semester she took 15-211, the course with a reputation of being a major hurdle:

I was really just baffled in that class because I just couldn’t understand anything, so my confidence went down again at that point, plus I didn’t know anybody in that class. So I dropped the class because I didn’t have any confidence in doing that. ... I took the course again in sophomore year, and things started to get clearer for me. Understand things better, plus at the time I made a lot of friends in the major. And you know, it is just the feeling that you have people going through the same thing with you. So it makes it better.

Former University of California calculus professor Uri Treisman (1992) believes that a supportive learning community is critically important for the success of minority students in math and science. Seeking answers to the high failure rate of African American students studying calculus at the University of California at Berkeley, Treisman observed that Asian American students formed social communities in which they helped each other with math, competed at mastering the material, and generally supported each others’ learning ...

He also found that most African American math students he studied were highly motivated, worked hard, and studied long hours but that even the best-prepared among them were failing. What stood out between the Asian and African American students was not a difference in motivation, preparation, or family support but in integrating studying and learning into social lives. African Americans were academically isolated and did not congregate into learning social communities the way the Asian students did. Instead, their academic interests and social interests were separate while they worked hard (and somewhat unproductively) on their own.

Observing the extra boost that comes from living and engaging with the material, Treisman has formed communities for African Americans similar to those created by Asian American students. These communities are built around intellectual interests (in this case calculus), provide well chosen problem sets that drive group interaction, and foster a supportive learning environment. Currently, Treisman-inspired Emerging Scholars programs operate in numerous colleges and universities and achieve high rates of retention in calculus courses among African American and Hispanic students.

Computer Science as an Acquired Taste

Studying the life arcs of women students in computer science over a four-year period has revealed to us some patterns of persistence. If students are able to stick it out through the second year, get grades they are pleased with, and reconcile their relationship to the culture, then their initial level of confidence often returns, accompanied by an increase in interest.

Interviews with persisters often reveal a key moment of success or achievement that keeps them going. For one senior, this moment was in her third semester, when she got over the hump of the data structures course (211) and began taking more advanced classes. She says that she had no confidence after 15-211 and “thought I would flunk out or get kicked out of CS.” But then she ended up getting an A in the course that immediately follows in sequence (15-212). She is in awe that she mastered the more advanced material. And the fact that she did it on her own became very important for her. In 211 she frequently needed to consult her teaching assistant, but in 212 she “was able to go right through the course without help.” That was her confidence builder:

In 211 I was constantly going to the TA, and I was like, “I don’t know how to do this!” And I felt like he was practically writing my programs for me because every time I’d have a bug or something, I’d be going to my TA two or three times for each program, at least. Then in 212 I was able to go right through the course without help or anything. It was just a great feeling for me, and I feel I learned a lot. And it was just a big transition for me. It was a lot of big “Ah-hah! So that’s what we were learning before!” All of a sudden things started clicking. It was just like a really big transition for me.

While this feeling of self-sufficiency may seem contradictory to the confidence gained from working with a supportive group, one way or another students have to internalize a sense that they can do it. If students persist for a sufficient amount of time (at least through the sophomore year), the odds are that they will regain confidence in themselves. Brenda, a sophomore student, talks of this confidence:

But it's kind of like if you're running, and you get to this big hill, and you're like, "Oh man, I'll never be able to run up that." And you do, and then you get to the next big hill. So it's like you're not exactly dreading it because once you get to the top, you feel really good about yourself. I guess I used to be afraid of a lot of things, but as I keep getting over and over these courses that I never thought I could pass, I think I'm ready to do the next step. And I don't know how I'm ever going to do senior-year courses, but I'll know when I get there.

We have found that if students get through the first two years, that a sense of mystery about computing turns into a sense of mastery. Asked if her interest in CS had increased or decreased, one junior provides an example of an upward spiral of confidence and interest:

I think partly it's increased just because I put so much work into it. It's like when you invest this much time in something, you want to do good in it. And also, I think the more I learn the more I think, "I can do this thing!" I just need to work really hard at it. But yeah, I think I've gotten more interested in it.

A Malaysian woman describes the satisfaction she felt in sticking it out:

It's like an acquired taste for me. ... At first it was very hard. ... After a couple of years, I realized it's kind of late to back out. I sort of went through with it, and along the process I'm beginning to think I like it more and more. So at the end, I just went along with it, and it's pretty exciting, now that that I learn more about it.

Conclusion

Despite doubts and uncertainties, women tend to persist in computer science when they reject and find alternatives to the dominant culture of the field. A larger question, though, is what institutions can and should do to eliminate the negative factors that lead students to leave computing programs.

The Future of Our Profession

There's more to being a good engineer than a high level of technical competence.

An organization like the ACM, with a variety of publications, newsletters, interest groups, and conferences, provides an important forum for discussing the nature and changes of the computing profession. Such discussions should of course go on continuously, but there will be particular periods of rapid social and technological change in which they are more necessary. A number of articles in *Communications* over the past few years [1, 3, 4, 10] suggest this is such a time.

The purpose of this article is to contribute to this discussion from a Scandinavian perspective. We begin by presenting a framework for understanding our profession and for discussing possible changes. The framework offers competing conceptions in terms of professional focus, approach, and role. We then use the framework in an examination of the current discussion within the ACM. The cited articles—in spite of having different motivation and purpose—all identify a need for professional changes in education, research, and practice. But, we argue, the changes suggested are insufficient to

meet the demands placed on our profession by the ongoing changes in technology and society. As a consequence, we propose a more radical position on the future of our profession, with *technology in use* as its foundation, and we outline a new curriculum for the education of computer professionals.

Competing Conceptions

Discussions of professional identity tend to be internal affairs, dealing with details of professional competence. If we expand our field of vision a bit, we will often find that parallel discussions occur in slightly different terms in other professions. It should come as no surprise that the issue under discussion here, the question of our professional identity, is an example of a much more general question that has been a recurring topic of debate during the process of modernization. Engineers have played a leading role in that process, turning a society of farmers and craftsmen into an industrial, high-tech society. How they think of themselves and their roles, echoes the two major, competing views of that process: the mechanistic and the romantic. (See the sidebar “The Mechanistic and the Romantic Views.”)

To put it bluntly: if your focus is on machines, and if you use the machine as a model for how you think of the world, society, and people, then you are a *mechanist*. If your focus is on people, on how they interact and change, on the processes of developing and using machines rather than on machines as such, then you are a *romantic*.

In our book *Computers in Context* [2], we analyze our profession using the mechanistic/romantic dichotomy as a dimension to compare and discuss different, and competing, conceptions of our professional identities. We have developed a conceptual framework distinguishing among three positions on what is the focus, approach, and role of a computer professional. Our framework is simple (perhaps deceptively so), and is summarized in Table 1 and is further elaborated in three subsequent sidebars on professional focus, approaches, and roles.

This framework is used in teaching throughout Scandinavia, both at engineering schools and at business schools. We have used it ourselves in courses for professionals at the computing centers of major Scandinavian companies, and we have seen how a framework like this can enlighten experienced professionals as to the identity of their professional roles.

A simple use of this framework is to distinguish the identities of three different professions within

Table 1. Competing conceptions of the computing profession

Slogan	I build things	I help people	I change things
Focus	Artifact	Culture	Power
Approach	Construction	Evolution	Intervention
Role	Engineer	Facilitator	Emancipator

computing: *programmers* in the software industry are engineers who build things; *support personnel* in computer departments rely on an evolutionary approach to help people use computers; and management and trade union *consultants* focus on power when they use computer technology to change things. It is more interesting, however, to use the framework to question such simple labeling, crossing the column boundaries, perhaps arguing that engineers ought to take a more active interest in questions of power, and that consultants should know more about artifacts. This is how we will use the framework here.

In our analysis of the recent discussions within ACM, we read the articles using our framework as an organizing principle, trying to determine what each of the texts have to say about professional focus, approach, and role. Our findings are presented accordingly in the following three sections, para-

The Mechanistic and the Romantic Views

The *mechanistic* world view was developed in the 17th century by great philosopher-scientists like Descartes, Newton, and Leibniz. Turning the idea of the clockwork mechanism, the machine, into a powerful conception of the universe, they used this world view as a foundation for a new science of nature, a new definition of man, and a new political program. This world view played a central role in the Enlightenment, and thus in the modernization of Europe. It contained all the elements that we have come to associate with a scientific attitude: truth as the mapping of the world, the use of mathematics as an exact language of representation, the idea of methods, of formalization and rules for calculation and deduction, the view of thinking as the competent manipulation of symbols, a penchant for rational means-ends thinking, for planning, analysis, and final solutions. Applied to society, this world view gave us such powerful principles of organization as bureaucracy and liberal democracy.

The *romantic* world view grew out of a reaction against mechanistic thinking, and was formulated toward the end of the 18th century, primarily by German philosopher-artists like Herder and Hegel. This reaction made much of the difference between organisms and machines, wanting to defend nature and everything natural against machines and everything artificial. The romantic philosophers were not interested in taking the universe apart like a machine, in analyzing it into its smallest atoms. They wanted to contemplate, understand, interpret, feel, and see through the world to its meaning, as with a poem or a painting. If the mechanistic philosophers of the 17th and 18th centuries tended to think horizontally, mapping the causal sequences of controllable machines, the romantics developed a vertical way of looking at things. Where the mechanists saw structures and systems, the romantics saw processes and change. Changes are the expressions of hidden, bottled-up forces, of eruptions resulting from power struggles and the unleashing of mounting tension. The romantics wanted to change our world by changing our experience of it. The world we human beings experience, the world of phenomena, is construed by concepts that vary from culture to culture and from time to time.

phrasing the actual texts to support our interpretations. We then evaluate the positions taken in the articles using our framework as a basis for criticism, summarizing and paraphrasing to substantiate our points.

Professional Focus

The conventional view of the computing profession within academia is that it is a set of experts who make their livings by solving problems in hardware and software, yielding faster, cheaper, smaller, and more reliable information-processing systems. In his program for “Educating a New Engineer,” Denning argues that this view no longer reflects the practice of computing.

Recording, processing, and communicating information have become an enduring concern of all human beings for effective coordination of work and action, and computer professionals are seen as the people who take care of other people’s concerns

The Professional Focus of Computing

Artifact focus: The computing profession is concerned with technical, computer-based artifacts. The professional focus is on tools and techniques for development of such artifacts for individuals, organizations, and markets. Questions of quality address the artifacts themselves, and primarily their technical functionality.

Culture focus: The computing profession is concerned with computer-based artifacts in the practical context of their use. The professional focus is on adapting artifacts to the practice of individuals and the different cultures of organizations. Questions of quality concern quality in use, the way artifacts fit organizational contexts, the way they influence and are influenced by, individual practice and organizational culture.

Power focus: The computing profession is concerned with the role of computing in changing society and the lives of people. The professional focus is on moral and political issues related to when and how to use computer-based artifacts. Questions of quality concern the impact of artifacts on the distribution of power, autonomy, integrity, and democracy [2].

in this domain. The computer professional is no longer an expert providing solutions to general information-processing problems. Instead, the computer professional has become an expert partner with clients in other domains [3].

This change in perceptions of the computing profession implies a change in our understanding of innovation, the *raison d’être* of the profession. We have always, according to Denning, placed a high value on innovation. But the traditional understanding is that innovation means introduction of a new computer-based artifact that makes a set of actions more efficient. A new understanding is emerging, according to which an innovation is a shift of the standard practices of a community to help them achieve their purposes more effectively [3].

Signs of similar changes in professional focus can be found in the other texts. Turner et al.—concerned with the design of computing curricula—claim that computer professionals must be able to anticipate the impact of introducing computer-based artifacts into a given environment, and they must understand the responsibility they will bear in doing so, as well as the possible consequences of failure. Students of computing need to understand the history of the discipline, they should appreciate the philosophical questions, technical problems, and aesthetic values that play important roles in the development of the discipline, and they should discuss serious questions about the social impact of computing [10].

Hartmanis et al.—assessing the scope and direction of computing research and education—believe the field and society will benefit if a broadening course is taken rather than if efforts at the core are increased. They suggest that members of the academic field of computing must intensify their intellectual interchange with other disciplines, focusing more on applications of computing in areas of economic, commercial, and social significance, and realizing that substantive research problems often emerge from such applications [4].

The “ACM Code of Professional Conduct” [1] exhorts computing professionals to give comprehensive and thorough evaluations of computer systems and their impacts, including analyses of possible risks. They should manage personnel and resources in such a way as to design and build information systems that enhance the quality of working life. Moreover, they should ensure that those who will be affected by a system have their needs clearly articulated during the assessment and design of require-

ments and that the system subsequently is validated to meet these requirements [1].

Analyzing these positions using our framework (see the sidebar “The Professional Focus of Computing”) we see a change from a narrow, traditional artifact perspective to a broader focus in which the cultural context of computer-based artifacts plays an important role and in which moral and even political issues are included.

A recommendation to broaden the professional focus to include the cultural context is clearly expressed in all the writings: the computer professional should help clients in other domains to achieve their purposes more effectively through a shift of the practices of the communities in question [3]; computer professionals must be able to anticipate the impact of introducing computer-based artifacts into a given environment, and they must understand the responsibility they have in doing so, as well as the possible consequences of failure [10]; research efforts should be broadened rather than redoubled at the core [4]; and, finally, computer professionals should provide comprehensive evaluations of computer systems and their impacts, including analyses of possible risks [1].

The idea of focusing on power is more vaguely expressed, and with important differences between the writings. There are only a few and rather vague hints at a power focus in Denning and Hartmanis et al., while Turner et al. address such a focus more directly. The latter emphasize that computer professionals have a responsibility in anticipating the impact, and possible consequences of failure, of introducing computer-based artifacts, and that students of computing therefore should learn to appreciate the philosophical views and aesthetic values that play an important part in the development of the discipline. Anderson et al. focus directly on both moral and political issues.

Professional Approach

North American employers and business executives, says Denning, are dissatisfied because computing graduates lack practical competence. They cannot build useful systems, formulate or defend a proposal, write memos, draft a simple project budget, prepare an agenda for a meeting, work in teams, or bounce back from adversity; they lack a passion for learning. The current concept-oriented curriculum is well suited for preparing research engineers, but not the practice-oriented engineer on which competition increasingly depends [3].

The environment of computing has changed, and we need to change our perception of what it means to work as a computer professional. Denning suggests the conventional understanding of professional work as a set of tasks by which a group of people carry out a mission has become inadequate. Instead, we should perceive work as a closed-loop process by which a performer completes action leading to the satisfaction of a request by a customer or client. Computer professionals are not performing tasks to meet abstract objectives. They work for, or with, other people striving to meet their changing needs and requirements, and many of the skills the students lack are in the areas of communication and collaboration, rather than in technologies [3].

In the same vein, Turner et al. identify three fundamental processes: theory, abstraction, and design. The new curriculum attempts to integrate these fundamental processes and the social context of computing [10]. The conventional emphasis on theory is complemented with new perspectives on the importance of laboratories in the curriculum to strengthen abstraction and experimentation. The curriculum also develops communication skills and includes significant design experiences, such as working in teams.

Hartmanis et al.—being primarily concerned with the overall policies of computing—are rather vague when it comes to the actual practices of computer professionals. One of their important points, however, is that the strong connections between research and computing practices imply that the traditional separation of basic research, applied research, and development is somewhat dubious [4]. They recommend increased interaction between academia, industry, and society at large to enhance the cross-fertilization of ideas within computing between theoretical underpinnings and experimental experience [4].

The ethics code provides abstract and general rules for professional conduct; it does not prescribe guidelines. It does, however, imply a certain repertoire of skills and techniques that a computer professional must master. To give one example, computer professionals should give comprehensive and thorough evaluations of computer systems and their impacts, including analyses of possible risks. As a consequence, they are in a position of special trust, and have a special responsibility to provide objective, credible evaluations to employers, clients, users, and the public [1].

In terms of the three professional approaches of computing (see the sidebar “Professional Approaches”) the targeted articles express positions

Professional Approaches

Construction approach: The task is to develop a technical artifact in response to a given and well-defined problem. To cope effectively with the complexity of this task, a rational approach is taken in which a sequence of specifications is transformed from overall requirements, through various levels of design, to a final, optimal solution.

Evolution approach: The task is to develop a technical artifact for a client or user with more or less clear and stable requirements. To cope effectively with the uncertainty of this task, an experimental approach is taken in which various models, prototypes, and versions are tried to reach a satisfactory solution.

Intervention approach: The task is to change a problematic situation in an organization through design and implementation of a computer-based artifact. The major challenge is to transcend traditions while at the same time protecting the good qualities of established work practices. This is done through an iterative learning process in which various solutions are designed, tested, and negotiated [2, 8].

in which an evolutionary approach is placed on equal footing with a more conventional construction approach. In fact, evolution is seen as a key approach to computing: practice is a closed-loop process by which a computer professional completes actions leading to the satisfaction of a request by a customer or client [3]; the fundamental processes of professional education are abstractions based on experiments and design, in addition to theory, and to facilitate these processes we should use laboratories in the curricula [10]; we need to increase interaction between academia, industry, and society at large, and to enhance the cross-fertilization of ideas within computing between theoretical underpinnings and experimental experience [4].

The intervention approach seems to play no role—or at least only a minor role. The implication of much of the rhetoric is, however, that an intervention perspective is needed to comprehend modern computing needs and practices. It takes the appreciation and skill of an intervention approach to develop computer-based artifacts that enhance the quality of working life, to thoroughly evaluate possible impacts and risks, and to ensure that the

requirements and needs of different interest groups are taken into account and eventually met [1]. Similarly, Denning's program includes many aspects of organizational intervention. In addition to being competent in engineering, the computer professional should be a skilled listener for concerns of customers or clients, be rigorous in managing commitments and achieving customer satisfaction, and be organized for ongoing learning [3].

Professional Role

Turner et al. provide new perspectives on a number of key issues, one being the importance of social, ethical, and professional issues in computing curricula. They identify a body of subject matter representing the social and professional context of the discipline that is considered essential for every undergraduate program. Students of computing should develop the ability to ask serious questions about the social impact of computing and to evaluate proposed answers to those questions, and they must be able to anticipate the impact of introducing a given product into a given environment. To do so, they need additional experiences that will help them develop the capacity for critical thinking, problem solving, research methods, and professional development [10].

Similarly, Denning suggests that we should recognize a second kind of knowledge besides procedural knowledge. This second kind of knowledge includes knowing how to listen, to design, to persuade, to be organized for new learning, to be professional, and even to be trustworthy and honest. Both kinds of knowledge are essential for an engineer. Moreover, in emphasizing the important role of innovation, Denning sees an innovator as a person or organization that articulates a change, offers the means to do it, and mobilizes people to adopt the new practices. Innovation is an organizational phenomenon, not merely an individual one [3].

Hartmanis et al. are rather vague when it comes to specific views on the roles of computer professionals, but one can infer underlying changes in direction of a broader, more interdisciplinary professional profile. The code of professional conduct is, in contrast, quite explicit in its perspective on this aspect of our professional identity. The specific imperatives cited earlier all express the view of a socially concerned and responsible professional who actively attempts to enhance the quality of working life. This view is further elaborated in the general moral imperatives.

Computer professionals should contribute to society when designing or implementing systems. Com-

puter professionals must attempt to ensure that the products of their efforts will be used in socially responsible ways, will meet social needs, and will avoid harmful effects on health and welfare. They also have a duty to be honest and trustworthy about their qualifications, and about any circumstances that might lead to conflicts of interest [1].

In our framework there are three different roles that computer professionals can play. None of these roles need to be less moral than the others, or less socially concerned. The difference between them lies rather in what is considered the most important factor to attend to if we want to improve the world: wealth, understanding, or equality.

The writings in *Communications* express positions in which the traditional role of the engineer is being supplemented with other values and responsibilities, most often emphasizing the role of facilitator but to some extent also the role of emancipator.

All texts clearly identify the facilitator role, and emphasize its importance and relevance to the computing profession. Computer professionals should, according to Denning, not only see themselves as traditional engineers who introduce new computer artifacts to make a set of actions more effective, but, in addition, as innovators working in teams to help organizations change their standard practices. Anderson et al. formulate the ideal of a computer professional who helps users express their needs and ensures that requirements are clearly articulated and implemented [1].

The role of emancipator is only vaguely expressed in the writings, except in the code of professional conduct. Here we find the image of a computer professional who is socially concerned, aware of the conflicts of interest related to the use of computers, strives to improve the quality of working life, and who generally contributes to society and human well-being. The computing professional who lives up to the new code of ethics and conduct is actively protecting and emancipating people from oppressive use of computing technologies.

Computers in Use

We can summarize the message of the recent discussion in the ACM thus: in addition to the traditional technical competence of an engineer, a computing professional today needs more social, organizational, and communication skills. A vision of dual competence is presented, in which each engineer is both engineer and manager, salesperson and organizational expert, an ethically, socially responsible engineer, and so on.

This is fine. We have no quarrels with these sug-

Professional Roles

Engineer role: Engineers try to improve the world by developing better computer artifacts. They are mainly interested in technical knowledge that gives them superior control over the processes of computing. Engineers want to increase the efficiency of computing and computer use.

Facilitator role: Facilitators try to increase the competence of users and clients, handing over to them the responsibility of acting. They are mainly interested in contributing to a better understanding of the use of computing technology. Facilitators strive to increase our understanding of how technology could be made to serve people rather than the other way around.

Emancipator role: Emancipators try to use computing technology as an opportunity to advance society and social organizations. They are mainly interested in protecting and emancipating people from oppressive use of computing technologies. They worry about the role of computing technology in supporting injustice and unequal distribution of power [2, 5].

gestions. Broadening professional competence is necessary, and it is already a fact in the sense that engineers have performed as managers, in sales and marketing, working with organizational change, facing social responsibilities, and so on, even if they often have not been prepared by their education for these tasks. One can debate what is the best way to prepare a student for these tasks. Given a certain time frame, should students concentrate on core engineering skills, making sure that they become confident in their professional identities, or should students spend more time preparing for their actual tasks? This discussion echoes a more general dilemma confronted by all education in a complex society, and often the solution is compromise and we will have to continue debating, and changing, the proportions in that compromise.

We don't think, however, that the suggestions in the considered articles are sufficient. Turner et al. and Hartmanis et al. remain purely within mainstream thinking when it comes to their specific recommendations; Denning offers a critical analysis, but his recommendations are mainly supplements to, or modifications of, established traditions (broaden our research agenda, reformulate our curriculum

around exhibitions of competent performance, reformulate our means and measures of assessment, provide institutional support for faculty development, and provide a more modern information infrastructure [3]. The new code of professional conduct implies a more radical view of our profession, but the imperatives are (naturally) so general as to give very little information on how to live up to them.

We must ask ourselves whether the proposed recommendations are effective responses to the requirements that seem to be imposed on our profession. What changes are needed and what should we as a profession do to facilitate these changes? Simply put, we argue that rather than envisioning a new engineer with social skills in addition to technical skills, we have to change what we consider technical competence.

In traditional engineering we concentrate on how artifacts function, and on how to make them function. The new engineering, we envision, will take a different view, attending to the use of artifacts, to the roles they play in our lives and how they play those roles. Such a perspective will revolutionize engineering, embedding it in a social context, making artifacts *in use*, rather than artifacts, its subject matter.

Sometimes, attempts at introducing more social and humanistic themes into the engineering curriculum seem to be an expression of an ambition to reduce the power of technology. This is not our vision. Technology is an outstanding social force and no amount of humanism will change that. But just because it is such a force, is it important for us, who have the ability to control it, to take an interest in its use.

A quick illustration of what we mean is the impact that a higher level of environmental awareness is beginning to have on engineering professions working with technologies, which have direct effects on the physical environment and on people as physical organisms.

Initially, the typical engineering response was to think of environmental issues as consequences of technology, leaving those consequences to politicians and other decision makers to deal with. As public interest in environmental issues grew, those issues began to interfere with engineering work, causing irritation. Gradually, however, environmental issues have entered engineering curricula, influencing the

nature of technical competence itself in more and more branches of engineering. From being a source of irritation, pollution has become an area of engineering expertise.

Analogously, we argue that as technology becomes more and more important and pervasive

Table 2. Key concepts of an ACM Computing Curricula [10].
The numbers in parentheses indicate the proposed lecture hours per subject area.

Fundamental Processes	theory; abstraction; design
Recurring Concepts	binding; complexity of large problems; conceptual and formal models; consistency and completeness; efficiency; evolution; levels of abstraction; ordering in space; ordering in time; reuse; security; trade-offs and consequences
Subject Areas	algorithms and data structures (47); architecture (59); artificial intelligence and robotics (9); database and information retrieval (9); human-computer communication (8); numerical and symbolic computation (7); operating systems (31); programming languages (46); software methodology and engineering (44); social, ethical and professional issues (11)

in all aspects of modern life, engineers will have to include more and more social aspects in their technical problem solving and into the very core of the technical curriculum. It would be much easier and much less dangerous to introduce such aspects as additional, but distinct, skills, of course; but it would also be much more inefficient and contrived to attempt to do so.

Our main argument is simple. When technology played a less pervasive role in society, as long as engineers were engaged primarily in military affairs, or in heavy industry; as long as they were not engaged in the everyday affairs of everyone, it was possible to carry on as if technology was somehow different from society. It made sense to speak of the social consequences of technology, and engineers became experts designing, constructing, maintaining, and repairing, technology while knowing next to nothing about the actual details of its use. But all this is changing now that technology is interwoven into everyday life.

The efficiency and productivity of modern society is based on the division of labor. Without division of labor, there is no hope of expertise. It has often been a long and arduous task—the different branches of engineering are good examples—to define and

purify the particular areas of expertise. But in a changing society, lines of division will have to be redrawn, and even such “natural” dividing lines as those defining technical competence may have to be adjusted.

Computer technology, or as it is nowadays often called, information technology, is a particularly striking example of the need for change. The name “information technology” makes some of us a bit nervous. As computer professionals we work with computer technology. It is the users, politicians, and media that talk of information technology. How do the two relate to one another? When we develop computer technology, do we also develop information technology?

Really, that question puts the finger on the transition we are advocating. As long as our professional task was restricted to the machine itself, we were doing fine as traditional engineers. As long as computers were used as automata, it was their independence from human beings that made them so powerful. The power of information technology, on the contrary, lies in its dependence on human beings, in the many ways—as tools, networks, media, and the like—in which it involves and enhances human actions and interactions. It is this power of information technology to infiltrate our lives and our minds that places new demands on our profession.

Consequences for Computing Curricula

These are big issues, and in order to result in something more than hot air, visions will have to be turned into specific programs for action. We can begin by taking a look at our curricula from this perspective and compare what we would do to the suggestions given by Turner et al.

Turner et al. propose to change the computing curriculum. As summarized in Table 2, they identify, using their own terminology, three processes and twelve recurring concepts that should be fundamental to the computing discipline as a whole, and they identify ten subject areas that should comprise the subject matter of the discipline. The subject areas listed in Table 2 appear to be the result of a bargaining process between established traditions. As a consequence, it is terribly conservative and skewed (half of the subject areas—algorithms and data structures, architecture, operating systems, programming languages, and software methodology and engineering—cover 84% of the lecture hours). Nontraditional engineering subjects are added as isolated subjects (e.g., “social, ethical, and professional issues”) and little weight is assigned to such subjects (4% of the lecture hours).

A number of questions can be raised regarding the proposed subject areas: Why is programming not treated as a subject area or a fundamental process in its own right (instead, programming occurs in all of the first nine subject areas)? Why are operating systems, as a subject area in its own right, and compilers, as an important part of the programming language subject area, still the only types of applications of computing that are given in-depth treatment? Why does the subject area “programming languages” not include specification languages in general as well as development environments? Why not include subject areas like programming and modeling, languages and environments, application paradigms, and computers as media?

Two of the recurring concepts of computing, proposed by Turner et al., go beyond a traditional approach to computing. One is the concept of “evolution,” needed to deal with evolutionary changes and their implications for all levels of computing. Another is “trade-offs and consequences,” to be used in handling trade-offs in computing, the consequences of such trade-offs, and the technical, economic, cultural and other effects of selecting one design alternative rather than another [10]. The traditional orientation of the recurring concepts is, however, illustrated by questions like: Why “complexity,” but not uncertainty or risk? Why “efficiency,” but not effectiveness? Why is quality not a recurring concept? Why are important dichotomies like data and information, and process and structure not included?

The three fundamental processes proposed by Turner et al. are characterized in the following way: *theory* is a process rooted in mathematics and used to develop coherent formal theories; *abstraction* is a process rooted in the experimental sciences and used to develop conceptual understanding; and *design* is a process that is rooted in engineering and used to develop computer-based artifacts to solve given problems [10]. These concepts raise a number of fundamental questions: Why is “theory” reserved for formal theories rooted in mathematics? Why is “abstraction” considered to be distinct from “theory”? Why are experiments only related to “abstraction”? Why is “design” only rooted in engineering? Why is “design” restricted to solving given problems? Based on such questions, a different conception of the fundamental processes of computing emerges:

Theory: This process is rooted in scientific disciplines, such as mathematics and organizational

behavior, that are fundamental to computing. We use this process to develop theories and conceptual frameworks to understand, design, and evaluate computer-based artifacts in use.

Design: This process is rooted in design disciplines, such as architecture and industrial design, that share with computing an interest in artifacts in use. We use this process to develop specific design skills and the ability to organize

types of computer applications.

Emphasize interpretation: Introduce the students to a spectrum of (qualitative and quantitative, informal and formal) approaches to describe, evaluate, and communicate about the design and use of computer systems; make them conscious of the role of the observer and the perspectives underlying different approaches.

Integrate perspectives: Encourage the students to use and integrate different perspectives, stressing the importance of being able to view the technology from the user's point of view; address natural and formal languages together to give a general understanding of languages as related to understanding, design and use of computer systems and to help students appreciate the relative strengths and weaknesses of different forms of expression and communication; distinguish between support and control, drawing attention to the role of computer applications as instruments of control.

Change priorities: Focus on use, and on dichotomies like information and data, process and structure, complexity and uncertainty, efficiency and effectiveness, rather than on a traditional, one-dimensional frame-

work based on data, algorithms, architecture, complexity, and efficiency.

Emphasize quality: Make the notion of quality an integrating, recurring concept, encouraging the students to develop a critical, but constructive attitude to the design and use of computer systems, and inviting them to reflect on their profession from a practical, more holistic point of view across traditional subject areas and academic disciplines.

and manage experiments.

Interpretation: This process is rooted in the humanities, in anthropology, and history. We use this process to understand and evaluate artifacts in use and problematic situations in computing practices.

In this discussion of the subject areas, recurring concepts, and fundamental processes of the "Computing Curricula 1991" we have begun to formulate a new program for the education of computer professionals, beginning with the use of computers, with an interest to improve both technology and its use. In more specific terms, such a program would:

Focus on use: Give the students a chance to develop a conscious and critical attitude as users of computer systems; teach them how to evaluate and compare the use of different types of systems; experiment with various forms of collaboration between users and computer professionals; give them a deep understanding of the differences and similarities between major

Using the schema introduced by Turner et al., these considerations result in a curriculum such as the one outlined in Table 3, where the subject areas to a large extent are defined in terms of the various ways in which computers can be used.

On a more general level, there is—as pointed out by Denning [3]—a growing doubt about the effectiveness of educational systems, and as companies loosen their bureaucratic ties in favor of networking, people are beginning to speak about "the vir-

Table 3. Key concepts of an alternative curriculum

Fundamental Processes	theory; design; interpretation
Recurring Concepts	data and information; quality; complexity and uncertainty; conceptual and formal models; consistency and completeness; efficiency and effectiveness; binding; evolution; levels of abstraction; ordering in space; ordering in time; reuse
Subject Areas	computing machines, architectures and networks; information systems, database management and information retrieval; control systems and operating systems; personal computing, human-computer interaction and interface design; artificial intelligence and robotics; programming, algorithms and data structures; software methodology and engineering; programming languages and computer linguistics; multimedia, intelligent agents, and Internet technologies.

tual university." In Scandinavia, we are now experimenting with extensive programs for "life-long learning" and "student projects," introducing a more practically oriented curriculum running parallel with the theoretical one, "professional networks," involving professionals and professors in schemes of job rotation and joint seminars, and "collaborative, virtual research centers," bringing universities and software organizations together in joint projects. Nothing about this is new, of course, as pointed out recently by Norman et al. in their discussion of learner-centered education [9]. But the attitude is different. The transcendence of traditional boundaries between education and work indicate that the industrial-age educational system, with the school as just another type of factory, is ready for revolution, after all.

Conclusion

The computing profession has grown out of mathematics and engineering departments, and we continue to seek our identity in these departments. We see ourselves as experts on the mechanism itself, on computer architectures, operating systems, compilers, programming languages, and database management systems. Over the years, the attention of our profession has shifted from numerical analysis to programming to software engineering to human-computer interaction to networking. These shifts have added new elements to our professional competence, slightly changed the center of gravity of our profession, but otherwise left many of us undisturbed.

The change that we have advocated here is more radical. It means introducing a romantic use perspective into every aspect of teaching and practicing our profession. Mechanistic thinking, so powerful in producing and characterizing the machine, may actually hamper us when we are trying to put it to good use. Even if acquiring a more romantic perspective may make no immediate difference to one's work habits as a programmer in the software industry, it is important for one's general role as a computer professional. Whoever is working with a technology with a great impact on people risks having to address questions concerning that impact.

We have illustrated the effects of the change we are advocating by outlining an alternative to the "Computing Curricula 1991" [10], and by suggesting a change of the educational system as such, exemplified by some experiments currently under way in Scandinavia and designated by concepts like learner-centered education.

Rob Kling has proposed an approach to educating computing professionals in which organizational informatics plays an important role [6, 7]. Kling argues that students should acquire reliable knowledge about the social dimensions of systems development and use and such educational efforts should build upon both the traditional technological foundations of computer science and the social sciences. Kling's proposal goes further in responding to the emerging requirements on our profession, but it shares with the "Computing Curricula 1991" the idea of adding new features to existing disciplines. We have indicated a different option: to redefine our notion of technical competence based on an interest in the use of technology—not with the purpose of rejecting useful technical knowledge, but with the ambition to challenge and eventually change the very basis of our profession. **G**

ACKNOWLEDGMENTS

We are grateful to Sten Henriksson and Finn Verner Jensen for valuable and useful comments on previous versions of this article.

REFERENCES

1. Anderson, R.E., Johnson, D.G., Gotterbarn, D. and Perrolle, J. Using the ACM Code of Ethics in decision making. *Commun. ACM* 36, 2 (Feb. 1993), 98–107.
2. Dahlbom, B. and Mathiassen, L. *Computers in Context: The Philosophy and Practice of Systems Design*. Blackwell, Cambridge, Mass., 1993.
3. Denning, P.J. Educating a new engineer. *Commun. ACM* 35, 12 (Dec. 1992), 83–97.
4. Hartmanis, J., et al. Computing the future. *Commun. ACM* 35, 11 (Nov. 1992), 30–40.
5. Hirschheim, R. and Klein, H.K. Four paradigms of information systems development. *Commun. ACM* 32, 10 (Oct. 1989), 1199–1216.
6. Kling, R. Broadening computer science. *Commun. ACM* 36, 2 (Feb. 1993), 15–17.
7. Kling, R. Organizational analysis in computer science. *The Information Society* 9, 2 (Mar.–June 1993), 71–87.
8. Lanzara, G.F. The design process: Frames, metaphors and games. In Briefs, U., Ciborra, C. and Schneider, L., Eds. *Systems Design For, With and By the Users*. North Holland, Amsterdam, 1983.
9. Norman, D.A. and Spohrer, J.C. Learner-centered education. *Commun. ACM* 39, 4 (Apr. 1996), 24–27.
10. Turner, A.J., et al. Computing curricula 1991. *Commun. ACM* 34, 6 (June 1991), 69–84.

BO DAHLBOM (dahlbom@adb.gu.se) is a professor of Informatics at Göteborg University, University of Aalborg, Aalborg, Denmark.

LARS MATHIASSEN (larsm@iesd.auc.dk) is a professor of Electronic Systems at the University of Aalborg, Denmark.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

Disciplined Minds

A Critical Look a Salaried Professionals and the Soul-Battering System that Shapes Their Lives

by

Jeff Schmidt

Rowman & Littlefield Publishes, Inc, 2000

TIMID PROFESSIONALS

(Excerpt from Chapter 1)

“No two people are allowed to read the same thing,” I said above the noise, gesturing toward the other passengers on the crowded subway car. My out-of-town visitor glanced around the clattering train. Indeed, the commuters hurtling toward their jobs in Manhattan’s office buildings, restaurants, shops and other workplaces were reading such a wide variety of material that my joke almost held up. That typical weekday morning found riders engrossed in all kinds of magazines, paperback books, the *Daily News*, the *Post*, the *Times*, office documents, a software instruction book and, yes, the Bible. Those who weren’t reading were listening to headphones, talking to others or, apparently, just thinking.

Seeing this every day on the subway set me up for a surprise one moaning when I went to catch a suburban commuter train to Manhattan. I had stayed overnight in Westchester County, an upscale New York City suburb where many executives and professionals live. I would be riding into the city with lawyers headed for big corporate law firms, financial analysts going to investment banks, editors bound for publishing conglomerates, as well as accountants, journalists, doctors, architects, engineers, public relations specialists and a host of other professionals. Boarding the train felt something like entering a library. There were no conversations even though nearly all the seats were occupied. Almost everyone was reading. But the dozens of passengers were reading only two things: the *New York Times* and the *Wall Street Journal*. I could have formulated another joke about allowed reading matter, but the scene was too spooky, like the aftermath of an invasion of the body snatchers; everyone dressed the same, in suits, sitting silently in neat rows and columns, each holding up a large newspaper, absorbing the same information.

A herd of independent minds?¹ Something seemed very wrong with this picture. It was obvious that when the subway riders and the suburban train riders converged at the workplace, the people who showed the greatest diversity in their dress, behavior and thought—the nonprofessionals—would be asked to do the least creative work, while the most regimented people would be assigned the creative tasks. This seemed just the opposite of what one might expect. And even more disturbingly, it indicated that people who do creative work are not necessarily independent thinkers.

Evidence that professionals are not independent thinkers has been around for a long time but has generally been ignored, in part because people don’t know how to make sense of it. The Vietnam War produced some revealing examples, which are worth looking back at.

On 12 January 1971, the federal government indicted Philip Berrigan and other East Coast antiwar activists on felony charges of plotting to impede the Vietnam War through violent action. The activists’ agenda supposedly included blowing up underground heating pipes in Washington to shut down government buildings, kidnapping presidential adviser Henry Kissinger to ransom him for concessions on the war and raiding draft boards to destroy records and slow down the draft.

The Justice Department prosecutors chose to hold the conspiracy trial in Harrisburg, Pennsylvania, a conservative area where a randomly chosen jury would be heavily against the defendants. However, before the jury was selected at what came to be known as the Harrisburg-7 trial, a group of left-leaning social scientists supporting the defendants interviewed a large number of registered voters in the area to try to figure out how to get a sympathetic jury there. They discovered, among other things, that college-

¹ The phrase comes from an essay title: Harold Rosenberg, “The Herd of Independent Minds,” *Commentary*, vol. 6, (September 1948), pp. 242–252. *Note: for concision, I have eliminated almost all of the numerous, subsequent footnotes found in the original reading. Please get Schmidt’s book if you are interested in further exploring the themes of this reading. –PR*

educated people were more likely than others to be conservative and to trust the government. Thus, in court, during the three weeks that it took to examine 465 potential jurors and pick a panel of 12, lawyers for the defense quietly favored skilled blue-collar workers and white-collar workers without a lot of formal education—nonprofessionals, although the sociologists and lawyers apparently never used that term.

The lawyers were uneasy doing this, however, because it went against their intuition. The notion of closed-minded hard hats and open-minded intellectuals is widespread and is reinforced by mass-media characters like loading-dock worker Archie Bunker and his college-student son-in-laws, “pinko” Mike. In fact, *All in the Family* made its television debut the very day of the Harrisburg indictments, 12 January 1971; by the time the trial and jury selection started, it had been on the air for a year.

Ignoring these false stereotypes paid off. The government put on a month long, \$2 million extravaganza featuring 64 witnesses, including 21 FBI agents and 9 police officers. The defense called no one to the witness stand. After seven days of deliberation, the jury was not able to reach a unanimous decision and the judge declared a mistrial; but with 10 of the 12 carefully selected jurors arguing for a not-guilty verdict, the government dropped the case.

Blue-collar skeptics? Loyal intellectuals? Was the Harrisburg survey a regional fluke? Look at what the nationwide polls showed at the time. On 15 February 1970 the *New York Times* reported the results of a Gallup poll on the war in Vietnam. Gallup had found that the number of people in sharp disagreement with the government over the war had increased but still constituted a minority. While this increase in opposition was important news, what were particularly intriguing were the data on the opinions of subgroups of the population. These numbers announced with striking clarity that those with the most schooling were the most reluctant to criticize the government’s stand in Vietnam. There was a simple correlation (although only in part a cause-and-effect relationship): The further people had gone before leaving school, the less likely they were to break with the government over the war. (See note 3 for the results of the poll.)

During the war in Vietnam, nearly everyone seemed to have one or another gripe about the U.S. government’s effort, but few took positions that dissented fundamentally from the government’s goals. Some said they were for negotiations, some said they were for an end to the bombing and some simply said they were “for peace.” Gallup’s survey cut to the bottom line by posing what was always the most incisive question on the war. It asked people whether they would favor or oppose the immediate withdrawal of all U.S. troops from Vietnam.

Age didn’t affect the answers much. The ratio of those in favor to those opposed was about the same for young adults as it was for older people. But dramatic differences appeared according to formal education. Those with college educations opposed immediate withdrawal by more than two to one, whereas those not formally schooled beyond the elementary grades were evenly divided on the question. And high school graduates were in between.

Polls taken earlier and later in the Vietnam War and polls taken during other wars—Korea, for example—show the same correlation with formal education.

Gallup was not the only one to find this connection between attitude and formal education. In a study entitled *A Degree and What Else? Correlates and Consequences of a College Education*, sponsored by the Carnegie Commission on Higher Education, researchers found college graduates to be “more supportive, or ‘hawkish,’ than the rest of the population.” Even in 1968, a year of rising antiwar sentiment and militant actions against the war, people who had been to college remained less likely than others to criticize the U.S. intervention in Vietnam, the Carnegie study found.

People’s reluctance to criticize the war was not simply the result of their careful analysis of an isolated issue. Rather, the position people took on the war followed almost mechanically from their overall political outlook (although some had their overall political outlook radicalized by what they experienced when they acted to do something about the war). With Americans being killed every day, almost anything one said about the U.S. intervention in Vietnam was heard as a statement on the U.S. political, economic and social system itself, and rightly so. Thus a narrow statement against the war could elicit a broad response such as “If you don’t like it here, go to Russia!” Few now seem to remember that throughout most of the war, those who called for the immediate withdrawal of all U.S. troops were seen as radicals—as critics of a lot more than the war. This explains, in part, the disparity between opposition and activism—why many opponents of the war didn’t speak out publicly. More students than workers were antiwar activists, even though workers who had antiwar sentiments far outnumbered students of all persuasions. Workers organizing publicly to get the United States out of Vietnam risked a lot more—namely, their jobs—because their employers were likely to see them as radicals and therefore a threat to the tranquility of

the local workforce.

The correlation between attitude and formal education is important for a book about professionals, because professionals typically have large amounts of schooling. Indeed, people in Gallup's occupational category "professional and business workers" have attitudes similar to those of people in the top education category. (Unfortunately, Gallup has no category for professionals alone.)

The relatively uncritical stand of professionals on the issue of war is just the tip of the iceberg, for it is *on the job* that professionals have the greatest number of opportunities to display their ideological caution. Anyone who has ever had a job that involved interacting with professionals, or who has had to deal with doctors, lawyers, bankers or the like, has surely encountered individuals with what we might call the "professional attitude"—confident and assertive individuals who exude the feeling that they are very much at home playing by the rules and that there is no pressing need to question the social structure in which they do their work. In many individuals such identification with the system shows up in the negative: Their confidence immediately melts into fear at any suggestion of not playing by the system's rules. (By "the system" I mean the hierarchical organization of production—the system of bosses and employees—and the social, economic and political practices that go along with it. Here and throughout the book my emphasis is on the hierarchical structure: "The system" may be read as "the hierarchy.") And in fewer but more memorable individuals, this conservatism takes the form of elitism or pompousness, seemingly critical postures that cover for personal insecurity but involve no risk, because they compliment the system by implying that it is too egalitarian, too democratic. Whether you are a professional or a nonprofessional, you encounter the professional attitude most frequently at work—and on matters of work—not only because it is in the workplace that you are most often thrust into contact with professionals, but also because it is on the job that professionals are most sure to act like professionals.

Most importantly, it is at work that the attitude of professionals has its greatest impact, both on you as an individual and on society as a whole. Whether a given professional designs buildings, writes newspaper articles, teaches courses or develops investment strategies, she makes important decisions that affect many people. Outside of work, however, the professional's attitude has relatively little effect on society (unless the professional makes a deliberate effort to the contrary). If, for example, you were given the power to dictate the outlook that governs the day-in day-out decision-making of a professional at work, and I were given the power to dictate the outlook that governs what that professional does inside the voting booth once every four years, then your power to shape society would be vastly greater than mine. ...

Public opinion pollsters report that professionals are more liberal than nonprofessionals on many social issues, such as civil liberties, personal morality and cultural issues. Liberal professionals smugly conclude from this that they are a force for social progress and that nonprofessionals are a conservative force in society. But the polls do not justify such a conclusion, for two reasons.

First of all, although professionals may be liberal on this or that question of the day, they tend to be very conservative on a long-standing issue of much greater importance to society: democracy. Discuss politics with a liberal professional and you will not hear a word in favor of a more democratic distribution of power in society, perhaps because in the professional's view ignorant nonprofessionals make up the large majority of the population. Even the most liberal professionals tend toward authoritarianism in their social visions.

The second reason the polls don't demonstrate that professionals are a more progressive force in society than are nonprofessionals is that the surveys focus on broad social questions and not on attitudes in the workplace, where both professionals and nonprofessionals exert their greatest influence on society. The nonprofessional who is conservative off the job is often not at all conservative on workplace issues and therefore is not necessarily a net conservative force in society. Similarly, the professional who is liberal off the job is often very conservative on work issues and therefore is a net conservative force in society.

Indeed, there is an enormous gap between the opinions of professionals and their professional opinions—the opinions that guide their work. When their opinions count, most professionals are conservative. Thus the engineer who believes that corruption is common among politicians in the United States freely offers that opinion. The political scientist, however, fears being quoted as saying any such thing, even though few people would find it shocking. Ask the nuclear engineer whether the nuclear industry influences reactor safety estimates, something that has long been obvious even to nonexperts, and you may get a lecture on the objectivity of mathematical calculations. And the liberal doctor who offers a cocktail party opinion against authoritarian police practices? Go to that doctor's office with a medical problem and see her lean toward the traditional authoritarian doctor-patient relationship. Professionals are liberal on distant social issues, issues over which they have no authority at work and no influence outside of

work.

Note that developments that raise doubts about the social or economic system itself are never distant issues, even when they are geographically distant and not direct issues at work. As we saw, the Vietnam War, which involved the state's forcing people to make the highest possible sacrifice for debatable reasons, was such a crisis of legitimacy for the system. Such national crises and other anxiety-producing situations or events are immediate issues and so tend to elicit from professionals the same politically timid outlooks that guide their work.

Not surprisingly, while professionals are tolerant of distant social criticism, they have little tolerance for anyone who tries to provoke a debate about the politics that guide their own work. Noam Chomsky, a Massachusetts Institute of Technology professor and an outspoken critic of the state and the intellectuals who serve it, sees this firsthand when he takes a short trip down Massachusetts Avenue to Harvard University. Stepping out of the domain of conservative engineers and into the world of liberal theorists of the state and state policy, Chomsky feels a marked change in the level of tolerance for his radical democratic views. He described this to me in a letter:

By conventional measures, the Harvard faculty is much more liberal, in fact left-liberal. MIT faculty are very conservative often, even reactionary. I get along fine with the MIT faculty, even when we disagree about everything (which is the usual case). If I show up at the Harvard faculty club, you can feel the chill settle; it's as if Satan himself entered the room.²

In this book I want to examine the outlook of professionals where it matters the most, which is on immediate issues—that is, on issues where what professionals do or say affects society directly. All workplace issues are immediate, as are a few outside of work. Thus, when I speak of professionals as uncritical and ideologically obedient, I am referring not to their opinions on distant social issues, but rather to the attitudes they display at work and in their work, where their conservatism shows up in its biggest and most socially significant way. And I am referring to their attitudes toward immediate nonworkplace issues, which are issues that raise questions about the merit or strength of the larger system—questions that professionals are usually quick to play down.

I don't mean to imply that all professionals are conservative when it counts. Some professionals do make trouble for the establishment. Although relatively few in number, such activist professionals help maintain an influential oppositional subculture in their workplaces, in their disciplines and in society. This subculture provides inspiration, encouragement and a vital safe haven for individuals whose thought deviates from the mainstream. And it gives its members the support they need to challenge their employers and others with power and to push for reform. Oppositional professionals have become increasingly influential since the 1960s, in part because of the battles fought at that time. The civil rights and antiwar movements, by successfully challenging the powers that be, helped make speech freer and the population more skeptical, conditions favorable for the opposition.

However, contrary to common belief, the number of oppositional professionals has remained relatively small. Consider, for example, college professors, who are among the most left-leaning of all professionals. Today, only about 5% of the 550,000 full-time college faculty members in the United States consider themselves to be to the left of the conservative-to-liberal mainstream. This 1-in-20 proportion of leftists hasn't fluctuated much in at least 30 years.¹ If the proportion seems higher than this, that may be because people who break away from the mainstream establish a presence way beyond their numbers and because radicals are speaking out more openly inside and outside of the classroom. Also, in a few disciplines in the humanities, leftists really have increased their proportion significantly—a fact that conservatives have misrepresented to make widely publicized claims that leftists have taken over higher education in the United States. The bottom line is that while the vast majority of professionals continue to share the views of corporate business executives on most basic issues, the important minority that dares to disturb the status quo has grown in influence, if not in size.

For understanding the professional, the concept of "ideology" will emerge as much more useful than that of "skill." But what is ideology, exactly? Ideology is thought that justifies action, including routine day-to-day activity. It is your ideology that determines your gut reaction to something done, say, by the president (you feel it is right or wrong), by protesters (you feel it is justified or unjustified), by your boss (you feel it is fair or unfair), by a coworker (you feel it is reasonable or unreasonable) and so on. More

² Letter from Noam Chomsky, 12 October 1992.

importantly, your ideology justifies your own actions to yourself. Economics may bring you back to your employer day after day, but it is ideology that makes that activity feel like a reasonable or unreasonable way to spend your life.

Work in general is becoming more and more ideological, and so is the workforce that does it. As technology has made production easier, employment has shifted from factories to offices, where work revolves around inherently ideological activities, such as design, analysis, writing, accounting, marketing and other creative tasks. Of course, ideology has been a workplace issue all along: Employers have always scrutinized the attitudes and values of the people they hire, to protect themselves from unionists, radicals and others whose “bad attitude” would undermine workplace discipline. Today, however, for a relatively small but rapidly growing fraction of jobs, employers will carefully assess your attitude for an additional reason: its crucial role in the work itself. On these jobs, which are in every field, from journalism and architecture to education and commercial art, your view of the world threatens to affect not only the quantity and quality of what you produce, but also the very nature of the product. These jobs require strict adherence to an assigned point of view; and so a prerequisite for employment is the willingness and ability to exercise what I call ideological discipline.

This book is about the people who get these jobs and become members of the ideological workforce—that is, professionals. My thesis is that the criteria by which individuals are deemed qualified or unqualified to become professionals involve not just technical knowledge as is generally assumed, but also attitude—in particular, attitude toward working within an assigned political and ideological framework. I contend, for example, that all tests of technical knowledge, such as the Graduate Record Examinations (GRE) or the Law School Admission Test (LSAT), are at the same time tests of attitude and that the examinations used to assess professional qualification are no exception. I consider in detail how the neutral-looking technical questions on such examinations probe the candidate’s attitude. The qualifying attitude, I find, is an uncritical, subordinate one, which allows professionals to take their ideological lead from their employers and appropriately fine-tune the outlook that they bring to their work. The resulting professional is an obedient thinker, an intellectual property whom employers can trust to experiment, theorize, innovate and create safely within the confines of an assigned ideology. The political and intellectual timidity of today’s most highly educated employees is no accident.

As attitudes and values have come to play an increasingly important role in the production of goods and services, employers have faced a choice: either hire huge numbers of managers to direct every move of the large number of employees who now do politically sensitive work, or hire employees who can be trusted politically and merely check the results of their work. Employers have pursued both strategies simultaneously. But the first one is limited by its cost, and so today every country in the world, from the United States to China, has a growing cadre of people trusted to do work that requires making decisions based not on detailed instructions but on an assigned ideology.

A long episode of the Cold War drew attention to the Soviet cadre. Beginning in the late 1940s the U.S. government beamed Voice of America radio programs directly to the people of the Soviet Union. These short-wave broadcasts were in English, Russian and a dozen minority languages spoken in the USSR. On and off from 1948 to 1987 the Soviet government operated as many as 3,000 jamming transmitters, at a cost estimated at up to half a billion dollars a year, to drown out these programs—except for the ones in English.

Never in its four decades of jamming did the Soviet government censor English-language propaganda broadcasts aimed at its population. Why? Was it simply because the number of Soviet citizens who understood English was too small to worry about? That is certainly part of the answer, but it cannot be the whole story, because no group was too small for the Soviet government to worry about. English was a standard course in the Soviet schools, and at least some of the students who did well in school and were selected to become professionals eventually learned it. The number of Soviet citizens who could understand English-language broadcasts may have been small, but so was the number who could understand many of the minority languages that were jammed, at least six of which were each spoken by less than 1.5% of the population.

The Soviets never censored the English-language propaganda broadcasts because those who spoke English were a select group of people who were trusted to maintain ideological discipline in their work (even when they were not enthusiastic about the assigned ideology). As Robert C. Tucker, a longtime student of the Soviet Union, told me; “They were more likely to be establishment people, and not dangerous.” Many of these people, such as journalists, academics and foreign service professionals, were

not only trusted to hear the U.S. government's viewpoint, but were also expected to know it so that they could answer it and not get caught off guard by it. The Soviets apparently treated the English-language broadcasts as if they were an exclusive service for their country's ideological workforce, prepping its members to handle any dangerous viewpoints that made it through the jamming and reached ordinary working people.

As work has become increasingly ideological, professionals have made up a growing fraction of the workforce. In the United States in 1920, only 1 employed person in 20 was a professional. By 1940, this ratio had increased to 1 in 15; by 1960 it was 1 in 12; and by 1980 it had risen to 1 in 8. Today, at the beginning of the 21st century, the ratio is approaching 1 in 6 and growing rapidly. (The year 2000 began with the number of professionals approaching 22 million and total employment approaching 35 million.) ...

The traditional image of the professional as an independent practicing doctor, lawyer or clergyman is misleading not only because of the proliferation of other professions, but also because very few professionals are free practitioners. The overwhelming majority are salaried employees. This has been true for many decades and is increasingly the case today as even the traditionally independent doctors and lawyers are swept into the salariat. Of every 9 professionals today, 8 are salaried employees and 1 is a free practitioner. Hence, when I use the term professional, I have salaried employees in mind. ...

A system of production that divides its nonmanagement workforce into two distinct components—employees trusted to follow an assigned ideology in their work and employees not trusted to do so—clearly takes ideology very seriously. In fact, this system, now nothing less than a world system, gives questions of ideology *highest* priority. It must do so because of its increasing vulnerability in the face of a more and more politically sophisticated population, and it does so within each and every corporate or governmental division and at all levels of administration within these units. As a result, you cannot make sense of the system as a whole, the organization that employs you, or even your own job, just from a list of the goods and services being produced; understanding, now more than ever, means knowing the very carefully construed ideologies that are guiding the production and that are being advanced through it. ...

Furthermore, professionals are the role models of the society towards which we are heading, a society in which ideology trumps gender, race and class origin as the biggest factor underlying the individual's success or failure. The victories of the feminist, civil rights and union movements of the past century have moved us closer to such a society. Thus, employers, led by the big corporations, are striving to ensure the survival of their precious hierarchical system of production by making it an equal opportunity system, which means subjecting employees to ideological scrutiny without sexist, racist or elitist discrimination. In the process, the corporations reveal what is most important to them and draw attention to the essential characteristics of the people who pass the strictest version of their scrutiny—professionals.

THE POLITICS OF NOT GETTING POLITICAL

(Excerpt from Chapter 2)

Doctor, lawyer, teacher, scientist, psychologist, economist, engineer, professor: What makes an individual a professional?

Technical knowledge and skill come to mind immediately. But there must be more to it than that, because the worker who picks up technical knowledge and skill on the job does not get reclassified as a professional.

With few exceptions, the professional is a product of the schools. The fact that off-the-job schoolings is what makes the difference between the professional and the nonprofessional is curious, because professionals-in-training often complain that much of the prescribed study is "irrelevant" to the technical knowledge and skills they will actually need to do the job. Students feel frustrated by the numerous "extra" requirements that they must fulfill to be allowed to work and that seem to constitute an unnecessary obstacle course....

If the seemingly irrelevant material is, in fact, irrelevant, then employers would be foolish to insist on hiring people with paper credentials when they could hire equally skilled nonprofessionals at much lower salaries.

Of course, employers are not being foolish when they insist on credentials. Professionals do something for them that skilled nonprofessionals cannot do. As a look at some examples will illustrate, employers can trust professionals to uphold the right outlook in their creative work.

Consider the school teacher. Those who employ teachers see them as more than workers who present the official curriculum to the students. A computer or television system could make such a

presentation. An important role of the schools is socialization: the promulgation of an outlook, attitudes and values. For example, the schools prepare students for the labor force not just by teaching them arithmetic, English, history and so on, but also by teaching them to follow instructions, adhere to a rigid time schedule, respect authority and tolerate boredom. Lessons in this “hidden curriculum” are taught as much in the numerous school-student interactions not involving the official curriculum as in those interactions that do. The employer trusts the teaching professional to manage these interactions in such a way as to advance the proper values. The professional is one who can be trusted to extrapolate to new situations the ideology inherent in the official school curriculum that she teaches.

As it professional, the teacher is “objective” when presenting the school curriculum: She doesn’t “take sides,” or “get political.” However the ideology of the status quo is *built into* the curriculum. The professional’s objectivity, then, boils down to not challenging this built-in ideology.

It is revealing that teachers who do question the curriculum attract the attention of school administrators, while teachers who are simply incompetent at teaching it tend to be ignored. (Indeed, when teachers are fired it is rarely for not teaching well.) “Legitimate” professional questions for teachers concern not what they should be doing politically in the classroom—the professional has an internalized willingness and ability to be directed in this area—but how best to convey the material in the official curriculum. In this alone, teachers are expected to use their creativity, and the awards of the profession go to those who do best.

Consider the cop. Robots could conceivably enforce the “letter of the law” and keep extremely busy doing so because of the abundance of infractions that occur. However, mindless enforcement would achieve the law’s goals only very crudely, if at all, and that is why law enforcers must be professionals. Professionals are hired to enforce the “spirit of the law”—the spirit in which the letter is written. Only the professional is trusted to sense, for example, which of the multitude of minor violations of the “letter” are acts of defiance against the “spirit” and therefore call for a response.

It may not seem very radical to say that the spirit of the law is to defend the status quo. However, the police adamantly deny playing anything but a neutral role in society. Nothing reveals better the actual partisan role of the police and the priority they give to the law’s spirit over its letter than do the thousands of “attitude crimes” that draw punishment every day in this country. An attitude crime is behavior that violates the spirit of the law, whether or not it also violates the letter. Maintaining a discourteous or disrespectful manner when pulled over by the police, for example, is not illegal, but it can get you a traffic citation instead of a warning, because the spirit of the law says “respect authority.” Similarly, subservience can sometimes get you off with a warning even though you’ve violated the letter of the law by, say, loitering. But if you talk back to the cops, the very same loitering can lead to handcuffs and it night in jail, especially if you are black or Latino. Surely many of the estimated 20,000 instances of police brutality in the United States each year are “provoked” by the suspect’s less-than-deferential attitude.

In 1980, statistics came to light in San Diego County indicating as many as 700 “attitude arrests” there each month. This figure included only cases in which arrestees were released hours or days later with no charges filed. The figure would have been much higher had it included arrests in which the police filed contrived charges as well as arrests for minor violations in which the police filed additional or more serious charges because of the violator’s attitude, a practice known as “overbooking.”

One San Diegan, Edward Lawson, was repeatedly stopped, frisked and arrested, often violently, solely because of his attitude. Lawson enjoyed walking in pretty residential areas, but as a black man with dreadlocks strolling through wealthy white neighborhoods at odd hours, he would be stopped frequently for questioning by the police. Lawson would demand to know why he was being stopped, but the cops were not interested in giving explanations. When Lawson would press his demand, he often found himself thrown in the back of a squad car with his hands manacled behind him. While Lawson’s demand was not illegal, it violated the spirit of the law, which says “know your place.”

Punishment for attitude crimes is rampant today. In California alone, police in 1997 made over 85,000 arrests in which they released the arrested individuals without filing charges, mainly because of lack of evidence. An even larger number of cases were thrown out by prosecutors before trial. Sixty-one percent of the individuals given the arrest-and-release treatment by police were minorities.

From employment law to landlord/tenant law to tax law to property law, the spirit of the law is to maintain the privileges of the wealthy. Yet the letter of the law is seemingly neutral on the question. “The law, in its majestic equality,” observed Anatole France, “forbids rich and poor alike to sleep under bridges, to beg in the streets and to steal bread.” Nevertheless, those who enforce the law tend to see the wealthy as

“good guys” and tend to be suspicious of people without property. This is not because police are inherently biased people but because they have to take up the spirit of the law to do a professional job enforcing it. The professional’s “objective” enforcement of the law boils down to acting in accord with no ideology *other* than the one built into the law. A cop who challenges the law’s built-in bias in favor of the status quo would quickly attract the attention of higher-ups. But this is rarely a problem, because the law enforcement professional is tuned more to following orders than to grappling with moral questions. The police officer’s “legitimate” professional questions concern not the nature of the social hierarchy that the law defends, but how best to enforce the law that defends it.

Consider the shrink. Many mental problems originate not in diseases of the brain but in deficiencies of society. The arduousness of living with unfulfilling work, financial insecurity, arbitrary bosses, lack of solidarity and insufficient personal power, together with the anguish caused by racism, sexism, ageism, lookism, ableism and all the other oppressive hierarchies that plague this society, helps explain the fact that more than 10% of the population (and not counting those with substance abuse disorders) suffers from mental or emotional problems. There are enough troubled individuals in the United States to keep busy 100,000 psychiatrists and clinical psychologists and a much larger number of clinically trained social workers and other mental health professionals. People’s mental problems often appear as deviations from social or legal norms and therefore are problems for the status quo as well as for the deviant individuals.

The problems of both would be solved if troubled individuals abided by the values of the status quo, and of course the mainstream mental health system more often than not works to alter behavior in that direction. But attempting to adjust people to the unhealthy society that caused their problems in the first place may not always be the healthiest approach for either the individuals or society. A simple alternative would be to help some troubled individuals bring out, clarify and sharpen their implicit critique—to strengthen them for the struggle in which they are engaged instead of removing them from it, because the struggle can be both therapeutic for the individual and beneficial to society. But the institutions of mental health, such as hospitals that employ psychiatrists and clinical psychologists, are institutions of the status quo. They are not about to turn the troubled into troublemakers, no matter how healthful that might be. The mental health professional is someone that such an employer can trust to move confused people away from struggle with social norms and authority and toward a life in which they are “well adjusted” to their place in the socioeconomic hierarchy.

As professionals, psychotherapists are “nonpartisan” in their work: They just help ill people get better. But to declare extreme nonconformity an illness, as psychology professionals do, is a partisan act because of the down-on-the-victim therapeutic framework it rationalizes: “Treating ‘sick’ individuals” is a much more politically conservative framework than is “treating individuals troubled by a sick and oppressive society.” Evidently it is not the place of the clinicians to question the health of the society to which the patient must be adjusted. Their “legitimate” professional concern is how best to bring about the adjustment. In this alone, they are expected to use their creativity. The few who do raise questions are seen as “getting political,” even though it is hard to imagine how they could get any more political than mainstream clinical psychology itself, which often practices conservative social action disguised as medical treatment.

As the above examples illustrate, the failure of professionals to question the politics built into their work serves the interests of those who have power in society and helps maintain the social and economic status quo. But refraining from questioning doesn’t look like a political act, and so professionals give the appearance of being politically neutral in their work.

Nevertheless, the public is becoming increasingly savvy about at least one way in which professionals support the system through their work. People are beginning to understand that the intellectual worker’s “professional judgment” or “expert opinion” is not objective as it claims, but rather favors the interest of his or her employer. (Supporting one’s employer and supporting the larger system are not the satire thing, but because it is basically a corporate system, each boosts the other.) The public most easily recognizes the political tilt of professionals toward their employers when it is blatant. Thus, not many people today are surprised at the Johns-Manville Corporation doctors, whose medical findings for decades helped the asbestos producer suppress information on the deadly health hazard posed by the “miracle mineral.” Similarly, many people immediately questioned the scientific opinion of a group of distinguished physicists arguing in favor of nuclear power when it was revealed that the physicists were

connected to the nuclear industry and major corporations. And today people may be outraged, but they are no longer surprised, when an HMO medical director—a doctor in a business suit—hustles patients out of the hospital very soon after major surgery, even when common sense indicates further close monitoring.

Expert witnesses in big-money court cases draw further attention to reason's eager subordination to power. Today only the most naive observers are surprised when reputable experts from the same field contradict one another under oath. In high-stakes trials each side can afford the best experts money can buy, and these experts often turn out to be big names in their fields.

Finally, consider the university professors who have received research grants from tobacco companies to study the health effects of cigarettes. These independent medical researchers, whose names are often followed by the letters "MD, PhD," are typically well-respected, highly prolific scientists at prestigious institutions such as Harvard University. Many have served on presidential or other high-level government advisory committees. For decades these scientists have served their sponsors' interests by finding tobacco to be safe and nonaddictive, and by attacking studies that find otherwise. In scientific journals, at scientific conferences, in press releases to the mass media, at congressional hearings and as expert witnesses in court, these doctors have given their professional opinion that cigarettes are not dangerous. Their views fly in the face of estimates, by public health scientists with no connection to the tobacco industry, that smoking kills 1,200 people per day in the United States.

In addition to their grantees, the tobacco companies also have scientists working for them directly. Over the years, a very tiny minority of these researchers have pushed behind the scenes to make public some of their findings critical of tobacco. But they typically did not push very hard and did not leak their findings to warn the public.

It took a socially conscious nonprofessional to show what needed to be done. In 1994 a paralegal who worked for a law firm representing a tobacco company, acting under the name "Mr. Butts" ... sent 4,000 pages of secret tobacco company documents to an antismoking activist. These revealing "Cigarette papers" show how embarrassingly easy it is for a well-heeled organization to get what it wants from reputable scientists.

The strategy of the tobacco companies has been to use scientists to make the dangers of cigarettes look controversial. The companies depend upon the fact that many observers hearing the word "scientist" naively think "nonpartisan." Thus the head of the Council for Tobacco Research, which was the major health research organization of the tobacco industry, told Congress, "We are scientists and we seek scientific truth" However, as the public has grown more aware of the need to ask for whom experts are working, the tobacco industry has found it increasingly difficult—but not yet impossible—to use its contrarian scientists to get people to think "controversial" when they hear about research findings that implicate tobacco in disease and death.

When employers designate certain jobs "professional" and insist that employees have professional training—not just the technical skills that seem sufficient to do the work—they must have more in mind than efficiency. Hierarchical organizations need professionals, because through professionals those at the top control the political content of what is produced, and because professionals contribute to the bosses' control of the workforce itself. It is crucial for the functioning and survival of the institution—and the hierarchical system of production as a whole—that the employees who make decisions do so in the interest of the employer. As we will see, the employer's control of the political content of the professional's creative work is assured by the ideological discipline developed during professional training. And the employer's control of the workforce is maintained in part through the professional's elitism and support for hierarchy in the workplace. The preparation process develops, and the qualification process measures, the students' willingness and ability to accept ideological direction from future employers. The one who has met the requirements—the "qualified professional"—can be trusted to do what is "politically correct" when making decisions and creative choices at work.

Professionals sell to their employers more than their ordinary labor power, their ability to carry out instructions. They also sell their ideological labor power, their ability to extend those instructions to new situations. It is this sale that distinguishes them from nonprofessionals, who sell only their ordinary labor power. Those in charge can trust professionals to make some decisions that must be made ideologically; nonprofessionals are trusted to make only decisions that can be made mechanically. Professionals implement their employers' attitudes as well as their employers' lists of instructions; nonprofessionals are only required to implement the instructions.

Ideological workers are more expensive than non-ideological workers, because they require a

greater amount of formal education. The same economic forces that drive employers to replace nonprofessionals with machines (which initially bring higher profits) also drive them to reduce the discretion of professionals by standardizing the work procedure, or even by introducing “expert” computer systems. In each workplace the bosses push for more and more detailed job descriptions and work guidelines, which transform the employee’s decision-making into a routine or rote activity and tend to strip the work-result of any imprint of the employee’s own thinking.

In fact, nonprofessionals are often *forbidden* to be creative in their work. In many jobs, the more closely the employees follow set workplace procedures and any special instructions for the tasks at hand, the happier the bosses are. Nonprofessionals know that they risk getting in trouble when they innovate to get the job done.

Professionals, on the other hand, are required to be creative in their work but within strict political limits. Their creativity must serve their employers’ interests, which often are not the same as their own interests, the interests of clients or customers or the public interest. Thus the corporate PR specialist assigned to field questions about pollution, defective products, the treatment of employees and other sensitive issues creatively uses the truth to paint a pro-company picture.... Employers don’t have time to decide every minor issue that affects their political or economic interests, and so they seek to hire others who will do things as if they had done them themselves. Thus, professionals control the technical means but not the social goals of their creative work. The professional’s lack of control over the political content of his or her creative work is the hidden root of much career dissatisfaction.

To say that professionals are ideological workers is not to say that they formulate the ideology in the first place, for they do not. Professionals have no more control over the ideology they propagate than nonprofessionals have over the design of the products they produce. Professionals merely have an operational grasp of the ideology inherent in their occupation’s actual role in society. Employers trust them to use that ideology to extrapolate policy and handle new problems as they arise, and to do so without constant supervision. Professionals are licensed to think on the job, but they are obedient thinkers.

All professional work is in part creative. However, individuals are selected to do professional work not because they are more creative than others, but because they can be trusted to make sure every detail of what they create is politically correct front their employers’ points of view. As human beings, professionals are not more creative than nonprofessionals. In fact, professional training tends to kill off natural creativity: In the corporate headquarters building you can often find more creativity down in the mail room than upstairs in the office of a lawyer, systems engineer or financial analyst, but it is untamed. Employers will hire dull but politically disciplined individuals over those displaying any amount of politically undisciplined creativity.

Just as professionals engage in playpen creativity, innovating within the safe confines of an assigned ideology, so too they engage in playpen critical thinking. Their work involves judging whether or not the ideas of others are in line with the favored outlook, but does not involve developing their own, independent point of view. Hence professionals tend to be what might be called “book review” critical, which is intellectually and politically safe because it doesn’t involve developing or taking a stand for an independent outlook. Professionals generally avoid the risk inherent in real critical thinking and cannot properly be called critical thinkers. They are simply ideologically disciplined thinkers. Real critical thinking means uncovering and questioning social, political and moral assumptions; applying and refining a personally developed worldview; and calling for action that advances a personally created agenda. An approach that backs away from any of these three components lacks the critical spirit.

Ideologically disciplined thinkers, especially the more gong-ho ones, often give the appearance of being critical thinkers as they go around deftly applying the official ideology and confidently reporting their judgments. The fact that professionals are usually more well-informed than nonprofessionals contributes to the illusion that they are critical thinkers.

Perhaps no one draws more attention to the political component of professional work than does the lawyer. All professionals give highest priority to making sure the right interests are served. Most professionals do this political work quietly as they much more visibly exercise technical skills that the public sees as nonpolitical: treating illness, informing readers, catching criminals, teaching children how to add and subtract, doing scientific research, developing new technology, designing whatever. For lawyers, however, the perception is reversed, because watching out for the right interests is not only their highest priority but also an unusually large part of what they do. Lawyers exercise the professional’s basic ability to sense interests, but, unlike other professionals, they exercise no other skill more prominently.

Indeed, lawyering involves such a high ratio of political work to technical work that technical knowledge is something of an afterthought in the training of lawyers. Yes, law schools do organize instruction around memorizing and applying specific principles of law, but this is done primarily as an exercise to teach proficiency at adopting and working within assigned ideologies. As Talbot D'Alemberte, former president of the American Bar Association and a critic of legal education told me, in law school the law is just it vehicle to teach a way of thinking. Perhaps the most obvious evidence of this is that students fresh out of law school—even those graduating at the top of their classes—do not feel they have the technical knowledge necessary to pass the bar examination. They must spend a couple thousand dollars to take an intensive six- to eight-week “bar review” course to learn what they need to know to pass the test and get their permit to practice law.

The emphasis on ideological skills in law school is precisely what the powerful corporate law firms want, because the high-stakes legal battles they fight defending big business and wealthy clients are paramountly political. The simple ability to recite the law does not qualify one to do this work. Representing powerful clients requires lawyers who can make creative arguments about the intent of the law, who can find ways to argue that the public interest would be served by a favorable ruling, and who can sway public opinion in high-profile cases, where this opinion influences the outcome. Settling losing cases out of court is political work, too. Lawyers for the powerful must know, for example, to give high priority to negotiating clauses that keep the terms of the settlement secret, this to protect the corporation's or rich person's public image and to avoid setting a precedent that would help other wronged parties to obtain justice. Thus the big laws firms aren't primarily interested in the technical skills of the law school graduates they hire. Those skills are easily picked up on the job; an ideologically disciplined mind is not. Similarly, the firms don't care much about bar exam scores. In fact, the large firms typically hire law school graduates before they even take the bar exam.

Lawyers have a negative public image because, unlike other professionals, they don't exercise socially redeeming technical skills. For this reason, they are seen as people who take without giving—a nonproductive element of society. Lawyers themselves, especially those at the big firms, make little pretense of doing work that benefits the public at large. Thus, more than other professionals, they feel the need to reserve some of their time to work “pro bono publico”—for the public good. (Most social workers, teachers, journalists, sociologists, scientists and other professionals would feel insulted if you asked them whether they set aside any working time to help make the world a better place.) In the words of Judge Laurence H. Silberman of the U.S. Court of Appeals for the D.C. Circuit, “Lawyers really see pro bono services as the penance they pay for serving a capitalist system.”

Perhaps the most widely distributed between-the-lines writing is the handiwork of journalists. The news stories they write for the front page of your daily newspaper are chock full of subtext. This becomes clear when different reporters describe the same event, because then their descriptions differ in substance only by what they have written between the lines.

On 1 March 1993, for example, the lead stories in both the *New York Times* and the *Wall Street Journal* covered the same topic: what had been learned about the World Trade Center bombing, which had occurred three days earlier. The article in the *Times* began with these words: “The bomb that devastated the garage under the World Trade Center in lower Manhattan on Friday apparently was ...” The article in the *Journal* began like this: “The bomb blast that drilled a four-story hold in a primary symbol of American commerce was” Clearly, the words in the *Times* played down the effect of the bomb—it destroyed a garage. Why was the destruction of a garage the top news story in the world, three days after it happened? The *Journal* answers this question up front.

What the *Times* and *Journal* reporters wrote between the lines here was no accident, but adhered closely to each paper's editorial outlook. And in each case that outlook is just what one would expect. The *Times* is written for a readership of professionals, who need ideological direction and reassurance of the system's strength. The *Journal* is written for bosses—business owners and executives—who give direction and do not need to be reminded where their interests lie. Among *Journal* subscribers, managers, outnumber professionals more than three to one. Among *Times* readers, professionals outnumber managers three to two.

JURORS: PROFESSIONALS FOR A WEEK

(Excerpt from Chapter 3)

Every year, more than a million nonprofessionals get a taste of what it is like to be a professional—when they serve on jury duty. Jury work involves decision making, and so it should be no surprise that the

government gives potential jurors a quick, essentials-only version of the special processing described in the previous chapter. As a result, the juror's courtroom adventure bears an uncanny resemblance to professional selection, training and employment, with the whole process speeded up to such a degree that days represent years. A look at the familiar drill of jury duty reveals what those in charge want most in their decision-makers, and it sets the stage for understanding the conflicts that surround selection for professional school, professional training after selection and professional work itself.

The first order of business in a trial is probing the attitudes and values of potential jurors through questioning and demographic analysis. Based on the results of this ideological assessment, each candidate is either weeded out or selected to serve on the jury. Favored are people who are programmable but not already programmed. (Thus professionals, who are loaded with ideological baggage from their fields and jobs, are often excluded.) Those selected are then subjected to a whirlwind indoctrination in which the judge impresses upon them that they must accept the law as it is given to them and follow that law rather than their own sense of right and wrong. The judge exhorts jurors not to let their views about the merit of the law affect their work. In the most typical words of the court, the guest professionals are told to judge only the facts of the case, not the law." Jurors who favor decriminalization of marijuana, for example, are expected to vote for conviction anyway if, by their judgment, the defendant really was caught smoking the contraband, as charged. Thus, jurors are expected to exercise judgment, but within an assigned ideological framework that they are forbidden to question—just like professionals. For professionals, of course, their employers' ideologies play the role of the law.

However, there is an important difference between jurors and real employed professionals: Jurors have greater freedom to criticize the assigned framework of values and to act on their own sense of right and wrong. Unlike professional employees, jurors can follow their consciences without worrying about losing their jobs or losing the favor of people who have power over them. If jurors think that justice demands it, they have the right to violate the court's instructions and judge whether the law itself is unjust or misapplied: they are not held in contempt of court for doing that. In fact, such "nullification of the law" by juries has a long and glorious history. Before the Civil War, for example, some northern juries found both slaves and abolitionists "not guilty" of violating the fugitive slave laws, even though their violation of these laws was clear.

Today the government's approach to the "problem" of such independent juries is simply to try to prevent jurors from learning that they have the right to criticize the law. As a result, obedient jurors can sometimes be seen after trials apologizing to defendants whom they didn't really want to convict, saying they had no choice. But not all jurors are obedient. This is because the courthouse system of ideological weeding out and indoctrination doesn't work perfectly, mainly because it is so rushed. Thus, even when prosecutors have an airtight case, juries that are uncomfortable with the law or the way the law is being used don't always convict. Sometimes these juries openly criticize the law. But much more often they choose to convince themselves that there is reasonable doubt in the evidence, because they are ignorant of their right to question the law or timid about asserting that right. Each of these types of juror behavior—ranging from outright obedience to principled dissent, with a kind of place-knowing dissent in between—corresponds to an equivalent type of behavior by professionals.

Members of the Fully Informed Jury Association, a national organization with headquarters in Montana, are dedicated to educating people about their rights and responsibilities as trial jurors. These activists argue that jurors have a moral responsibility to judge the law and the way the law is being used, in the interest of social justice and as a check on those with power. I argue that for the same reasons, all professionals, not just temporary ones, must sit in judgment of the social goals they have been recruited to further.

SUBORDINATION

(Excerpt from Chapter 13)

"The (expletive deleted) computers." When an interviewer asked young nuclear weapons designers at Lawrence Livermore National Laboratory to name the worst thing about their profession, this gripe about uncooperative computers was a typical answer. The goddamn computers. They don't have enough capacity and they're always going down. What's a designer of weapons of mass destruction to do?

The physicists' startlingly narrow answers were not the result of any pressure to respond quickly, without giving careful consideration to the question—the interviewer reports that nearly all of them did think for a while before answering. Rather, a narrow focus comes naturally to such individuals, who, as good products of the system of professional training, give higher priority to carrying out their assignments

than to questioning them, and in any case are not prepared to second-guess the political and ideological framework that engenders and guides their technical work. In the final analysis, the physicists' narrow answers are a sad sign of their subordination, of their approval of a work life that will ultimately give them insufficient satisfaction: a work life in which their employers define the big picture and they innovate safely within it, and in which attempting to alter the picture is not a legitimate on-the-job activity.

If an individual professional did have an independent political agenda, it would undermine the ideological discipline and assignable curiosity that ensure that he works in his employers interest. This is why the system of professional training, and the examination that stands at the center of that system, favors the individual who sees himself as having a technical orientation rather than a political one. Of course, the technical is itself political—the technically best solution to a given problem is often one thing from the point of view of those with an interest in maintaining the hierarchy but something quite different from the point of view of those without power. However, the favored individual sees no class interest in his own work: Because he *internalizes* the requisite ideology, he doesn't see himself as following an ideology at all, but as simply doing what he judges to be technically best.

Having long ago purged himself of his political agenda, and having internalized the dominant ideology, the expert sees the problems of the world as fundamentally technical in nature (although certainly exacerbated by politics, but not the other way around). As a 28-year-old Livermore physicist working on third-generation nuclear weapons (which aim to knock out attacking nuclear weapons) said about the buildup of nuclear arsenals, "Why not find technical solutions to a technical problem?" Whatever the issue, the rebel and the expert stand out in sharp distinction to each other. In any discussion, the expert's lack of political independence—his loyalty—becomes apparent immediately, as he confines his thinking to technical solutions—making adjustments, fine-tuning the system. He may offer a multitude of ways to deal with a problem, but, as if by magic, not a single one would reduce the flow of profits or otherwise disturb the hierarchical distribution of power.

Professionalism—in particular the notion that experts should confine themselves to their "legitimate professional concerns" and not "politicize" their work—helps keep individual professionals in line by encouraging them to view their narrow technical orientation as a virtue, a sign of objectivity rather than of subordination. This doesn't mean that experts are forbidden to let independent political thoughts cross their minds. They can do so as citizens, of course, and they can even do so as experts, but then only in the "proper" places and in the "proper" way. The expert is probably a member of a professional association that has a "committee on social implications" or a "forum on the profession and society." Such a group may take up a political issue, but only after it takes a debilitating blow to power, usually in the form of a protracted debate in which those who want to take up the issue must succeed in repackaging it as a "legitimate professional concern," often as a technical issue. Members of the group can then take a position on the sanitized issue without "being political" in the sense of acting like they don't know their place. Politically timid professionals fear that their organization will look like part of a social movement and so they try to limit their organization's actions to those of a narrow special-interest group.

As part of their very identity, professionals subordinate themselves to power on ideological matters. Thus, professionals can't take a stand on an unsanitized issue without going through a genuine identity crisis. Indeed, they respond with great fear and trembling whenever anyone proposes that they take such a stand. Even on life-or-death issues, professional associations can rarely muster the courage to take a position that they think might displease employers. Professionals don't want anyone to think that their own views might affect their work, because that would be insubordinate and therefore unprofessional. So even off the job (in professional associations and elsewhere), independence of thought feels out of line. As a result, the typical professional doesn't stand for anything.

Thus, for example, it was sad but not surprising when the National Lesbian and Gay Journalists Association decided not to participate in the massive 25 April 1993 gay rights march on Washington, an event that drew several hundred thousand people, making it one of the largest civil rights demonstrations in American history. Leroy Aarons, the group's president, explained that members didn't want to endanger their "credibility in the industry." As good little professionals they adjust their very identity for their employers: Both on and off the job they act like journalists who happen to be gay, not like gays who happen to work as journalists.

Consider the behavior of the National Association of Black Journalists in the case of Mumia Abu-Jamal, a well-known journalist convicted of killing a Philadelphia police officer. A passionate voice for the black community, Abu-Jamal had worked as a radio and print journalist, doing news reports and

commentary for a number of radio stations and networks, including National Public Radio, which aired his pieces on *All Things Considered*. The incident that landed Abu-Jamal in prison occurred in 1981, while he was president of the Philadelphia chapter of the National Association of Black Journalists. Late one night Abu-Jamal happened upon the scene where a police officer had stopped Abu-Jamal's brother for it traffic violation. In the events that followed, the officer shot Abu-Jamal and was fatally shot himself. There is no agreement about who fired first or who shot the officer. Many people felt that Abu-Jamal, a radical and longtime activist with no criminal record, did not receive a fair trial. But years of appeals through the courts were fruitless, and on 1 June 1995 the governor of Pennsylvania ordered prison officials to kill Abu-Jamal by lethal injection at 10 p.m. on 17 August 1995.

This touched off a worldwide outcry involving hundreds of thousands of people. Demonstrations, rallies, teach-ins, celebrity speak-outs, op-ed articles, and letter-writing and petition campaigns—100,000 signers in Rome alone—demanded that Abu-Jamal at least be granted a new trial. Scores of organizations—from Amnesty international and Human Rights Watch to the NAACP Legal Defense and Education Fund and the Southern Christian Leadership Conference—supported this demand and opposed Abu-Jamal's impending execution.

But not the National Association of Black Journalists. In a written statement following a vote by the group's 18-member board of directors, NABJ president Dorothy Butler Gilliam said: "As an organization of journalists, [NABJ] does not see this unfortunate circumstance as an issue of journalism upon which it feels compelled to take a stand at this time." In spite of Gilliam's attempt to make her position sound more reasonable by calling NABJ "an organization of journalists" rather than an organization of blacks, the group's stand infuriated many people, especially the black community. One critic spoke for them all when he attributed the groups decision to its domination by members "attuned to the subtle grunts and imagined nods of their employers in the corporate media." In response to the barrage of criticism that it received, the NABJ latched onto and took an extra-strong stand on a sanitized issue: restrictions that prison officials had put on Abu-Jamal's communication with the outside world after he contracted to write *Live From Death Row* (Addison-Wesley, 1995), a book sharply critical of the justice and prison system. Thus the NABJ was "outraged" not because the state planned to kill Abu-Jamal, but because it was violating his First Amendment rights, "which we find to be a legitimate issue," the group explained.

The judge in charge of Abu-Jamal's case was a tough, cop-on-the-bench type who had sentenced 32 people to death—more than twice as many as any other judge in the country. He had never before granted a stay of execution and was, in the words of the *New York Times*, "openly contemptuous" of Abu-Jamal. Yet, ten days before the planned execution, he succumbed to the growing popular pressure and granted Abu-Jamal an indefinite stay, allowing Abu-Jamal's lawyers to appeal once again to higher courts for a new trial.

Generally speaking, the greater the power, whether corporate or state or even oppositional, the more eager professionals are to subordinate themselves to it. The power's morality or immorality usually has only a secondary effect on the professional's eagerness to serve, because good subordinates don't make moral judgments about their superiors. This is the unfortunate but invaluable lesson of history. Historian Konrad Jarausch notes, for example, that "in the spring of 1933, most German professionals rushed to curry favor with the new Nazi government." The prestigious German engineering association, the prominent lawyers association, the secondary teachers association and hundreds of other groups all across Germany pledged their loyalty. The behavior of people in my own field, physics, has been far from exemplary. Before and during World War II, the world's top physicists were German, and these individuals typically accepted invitations to work in support of the Nazi war effort. Two decades later, during the Vietnam War, the world's top physicists were American, and these individuals typically jumped at the invitation to become members of the Defense Department's Jason organization and work in support of the U.S. side in Vietnam. (Jason is still active.)

At the workplace, experts can be somewhat independent in informal discussions, but almost never within their professional work itself; it is considered "unprofessional" for experts to bring independent political thinking to bear in their work. On the job, their "legitimate professional concerns" are largely confined to carrying out their assignments. Thus, while some of the nuclear weapons designers mentioned at the beginning of this chapter worried about computer troubles, others—also well trained in confining themselves to their "legitimate professional concerns"—worried that international agreements might further restrict testing and thereby make it more difficult for them to carry out their assignment of weapons design. They did not allow the notion that such restrictions on nuclear testing might actually represent social

progress to interfere with their work.

This view of what is legitimate holds hegemony over professionals in every major area of their employment. It is extreme in some cases, such as the aerospace industry, employer of thousands of scientists; the very notion of an aerospace scientist bringing a critical social perspective to his work is so unusual as to be jarring. The social function of the individual produced by the qualification system is to work uncritically within the political hierarchy, bolstering it through his example of eager participation as well as through his actual work product. When the professional leaves unchallenged the moral authority of his employer to dictate the political contour of his work, he surrenders his social existence, his control over the mark he makes on the world.

These days one finds students and professionals who have some awareness of the big picture but who *cynically* adjust their behavior for the system. This is quite acceptable to the hierarchy because these individuals, even as they blast distant power figures such as the president, carefully avoid any confrontation with those who hold immediate power over them. As Max Horkheimer said in 1946, in what may be taken as one of the most succinct criticisms of many professionals on today's postmodern scene, "Well-informed cynicism is only another mode of conformity."³

However, more than professionalism or cynicism, it is lack of social vision that assures conformity, and professional training does anything but produce people who envision a more democratic social order. Professionals may complain to you about the unfair treatment that they witness firsthand at work, and they may tell you in excruciating detail about the latest cases of corruption in business and government, just as they read it in the newspaper. But most of them are unable to move from concern to action. Professionals are angry about such abuses of power, but having no vision of how power in the schools, in time workplace and in the larger society could be distributed more democratically, they naturally look for ways to make the present hierarchical power structures work. Here the choices are limited—restaff the hierarchy with "better people" or give those at the top even more power so they can "act decisively." So even the most well-meaning individuals end up reinventing some such elitist or authoritarian solution.

Group action by the rank and file is disobedient and antithetical to making hierarchical authority structures work, so many professionals who are well-informed and concerned about abuses of power will nevertheless not engage in collective acts of solidarity with the victims. They don't seek solidarity even when they themselves are the victims, and it is not unusual to see them leave their jobs rather than speak out openly and improve the situation through collective action. For the same reason, they will not identify with a specific movement or work with organizations that have independent social agendas.

Those who have no vision of greater democracy are paralyzed even further by the individualism inherent in their outlook. They retreat in fear at the mere suggestion of joining with others in struggle, for those who act as part of a group admit to being less than autonomous individuals and give up the comforting fiction that they meet their bosses as equals.

The professional, like any employee, does have conflicts with his employer, but because he is an intellectual employee, he is not free to arrive at just any understanding of the root cause of these on-the-job disputes. Specifically, under normal circumstances he cannot allow himself to view his problems with his employer as an outgrowth of a fundamental conflict of interest, for to do so would sabotage the ideological discipline that allows him to serve his employer's interest in his work and keep his job as a professional. Thus, the professional sees his clashes as originating in conflicting technical judgments over how best to pursue universal interests. He sees conflicting strategies or personalities but doesn't see himself as having a fundamental conflict of interest with his employer—or with the powerful in society in general. That is, he doesn't see his own conflict with his employer as part of a larger conflict between labor and capital. When those who wield power act against his and his fellow employees' interests, the professional does not see them as opponents acting against employee interests, but as incompetents acting against universal interests. Thus, he calls not for breaking down the hierarchy and distributing the power democratically to those who do the work, but for more "intelligence" at the top—an elitist approach, which weakens alliances with nonprofessionals. He challenges the staffing, not the structure. He fumes, "Incompetents! Stupid bureaucrats! Those idiots don't know what they're doing!" In the eyes of the professional, those with authority at worst lack intelligence or information; he dare not admit to himself that those he serves may be smart and well-informed but simply have different class interests—that is, he cannot risk admitting to himself that he has been hired to serve interests that conflict with his own.

³ Max Horkheimer, *Eclipse of Reason*, Oxford University Press, New York (1947), p. 113.

This restricted understanding renders the professional weak as a force for his own defense and impotent as a force for change in society. His protestations are impotent because, no matter how militantly stated, they are not threats to break ideological discipline. They don't threaten to affect the political content of his work, as having an independent political agenda certainly would. Even his strongest indictment of decisions made by management—"it's all political!"—suggests a mythical nonpolitical approach rather than an alternative distribution of political power. The louder he shouts his carefully restricted criticisms the more he proclaims his subordination.

No professional maintains perfect ideological discipline, and every straying leads to a run-in with management. Of course, some professionals have more themselves clashes than others. In particular, those who are the least strict about subordinating their own vision to that of the institution that employs them are the ones who find themselves in trouble most often. But these conflict-plagued employees rarely understand that their poor ideological discipline is the source of their clashes. They avoid such an understanding because it is inherently radical: It exposes their employer's ideology and is critical of it. ...

Institutions demand conformity and obedience and yet hire professionals to do work that requires creativity and questioning. Does this make employer-employee conflict inevitable? Liberals say yes. They enjoy believing that intellectuals are unbridled thinkers and therefore a threat to those in power. (This is a corollary of their elitist belief that nonintellectual workers support the status quo.) But I would argue that institutional demands for political conformity lead to conflict only when individual creative workers have independent political agendas and are not willing to subordinate them. For if professionals adopt their employers' agendas, then their creativity and questioning work toward meeting their employers' goals. The work product in that case is essentially the same as it would be had the employers done the creative work themselves.

The sad fact is, mainstream professionals don't need political freedom to do their creative work. And they don't demand that their employers allow them to exercise political freedom in their work. Only when professionals have an independent political agenda do they need and demand freedom, because only then might their creative work displease their employers.

Scientists are a good example. During Josef Stalin's reign of terror in the Soviet Union, tens of thousands of scientists and engineers were arrested, imprisoned and sometimes executed. Yet Soviet science advanced rapidly and came to lead the world in many fields, including mathematics and theoretical physics. Until the mid-1950s, some of the Soviet Union's most eminent scientists worked in prison laboratories. At the height of the repression, Soviet physicists did work that won them five Nobel prizes. One of those physicists, a Soviet citizen named Pyotr Kapitsa, had been living in England for thirteen years when, upon a routine visit to the Soviet Union to attend a conference, Soviet authorities seized him on Stalin's orders and wouldn't let him return home to England. Within a few years of this kidnapping, Stalin had Kapitsa running a Soviet laboratory and doing the most creative work of his career.

As Loren H. Graham, a science historian at Harvard University and the Massachusetts Institute of Technology, has documented, scientists do not require academic freedom to do their creative work—they just need funding. One haunting image that Graham describes is that of the young scientist Andrei Sakharov sitting at his desk at Arzamas-16 doing his famous work in theoretical physics and gazing out the window at brutal armed guards marching rows of political prisoners to their jobs at the scientific installation, which was the Soviet equivalent of Los Alamos National Laboratory in the United States. Years later Sakharov became a dissident, but that was unusual for a scientist. As Graham notes, even when the Soviet Union was on the verge of collapse, the leaders of Soviet science sided with the old order.

Those naive enough to believe that the professional's creative thinking alone leads to conflict with management probably also subscribe to the myth that the intellectual worker's "professional obligations" lead to conflict, too. No one illustrates better than the mainstream journalist that there is no tension between fulfilling a "professional obligation" or doing a "professional job" and institutional demands for conformity and obedience. The reporters who write front-page stories for the *New York Times* are considered to be among the top journalists in the profession. It is abundantly evident that the paper they work for requires that the stories be written within a framework of general support for the U.S. political and economic system (and that the stories anticipate and head off any possible faith-threatening interpretations of the facts being reported). *Times* reporters conform strictly to the paper's politics and at the same time feel that they are fulfilling their professional obligation to "get the story." There's rarely a serious complaint from either side.

Only when professionals have an independent political agenda do they argue with their bosses about what constitutes a "professional job." When *Times* editors assign one of their politically reliable, top-of-the-line journalists to cover a sensitive story, they don't worry that professional obligations will lead

their reporter to frame the story in a way that skewers the paper's fundamental tenets. Thus, for example, most mainstream media in the end reported the Watergate affair not as evidence of the political system's tendency toward corruption, but as evidence that the system works and cleanses itself.

As we know, not all students become clones of the prototypical professional described above. But those who are headed in that direction are easy to spot, because their subordinate attitude is conspicuous early on at the training institution. These students scramble to figure out the rules of the game in their university graduate department or professional school, and then they literally compete to adjust themselves appropriately. Being not merely adjustable, but self-adjusting, they are good students in the eyes of the faculty. For the same reason, they will be good professionals in the eyes of their employers. These students do not simply refrain from acts of insubordination, such as challenging the training institution's agenda or criticizing the ways that agenda reflects the needs of the larger system. Rather, they enthusiastically embrace the system of professional qualification and defend the qualifying examination. The personal strategy of these skilled submissives is to play the game: to use the qualifying examination to demonstrate on the system's terms that they are "good" (that is, well-adapted), to be certified with a credential and to get a job with a new set of rules to submit to. In short, this means integrating themselves into the system, being dwarfed by it but surviving, if not its independent forces for change in society, then at least as well-fed biological entities serving the status quo.

These students also subordinate the dreams they once had of experiencing the totality of their subject in all its technical and social dimensions. In what can be seen as it sad attempt to imitate this forgone experience, some students treat the small problem parts assigned to them as if they were interesting enough in and of themselves to play the role of a surrogate totality. Today these assignments are the catechism-like test-preparation problems, tomorrow the narrow thesis problem and thereafter the corporate problem segment.

Many students do resist making the appropriate adjustments and heading down the designated road: Unwilling to reorient their outlook and goals, they find themselves in conflict with one or another action or policy of the training institution. These students usually struggle individually and indirectly, misunderstanding their problems in the training program as simply personal and not the inevitable result of the system-serving nature of the training institution's goals. Though they often leave the training program, they should not be looked upon as "losers," for they have not necessarily been broken and may go on to struggle elsewhere.

The system of professional training is set up to turn students into good self-adjusters or else get rid of them. Through the mechanisms of pressure and scrutiny that I have described in this book, it usually succeeds in doing one or the other. However, students can and sometimes do frustrate the system by both confronting it *and* remaining, but this is accomplished only through politically conscious, organized action, as I discuss in the next two chapters. At the core of the conflict is an unstated but highly contentious issue: Who will the student become? Professional training programs work routinely, methodically and often consciously to turn students into very different persons, and so individuals who want to control who they are must fight to do so.

Some Pledges for Scientists and Engineers

1. Student Pugwash Pledge

Student Pugwash USA, www.spusa.org/pledge/

I promise to work for a better world, where science and technology are used in socially responsible ways. I will not use my education for any purpose intended to harm human beings or the environment. Throughout my career, I will consider the ethical implications of my work before I take action. While the demands placed upon me may be great, I sign this declaration because I recognize that individual responsibility is the first step on the path to peace.

2. Scientists' and Engineer's Pledge to Renounce Weapons of Mass Destruction

Los Alamos Study Group (www.lasg.org), Natural Resources Defense Council (www.nrdc.org), Tri Valley CAREs (www.trivalleycares.org), Western States Legal Foundation (www.wslfweb.org).

I pledge never to participate in

- *the design, development, testing, production, maintenance, targeting, or use of nuclear, biological, or chemical weapons or their means of delivery.*
- *research or engineering that I have reason to believe will be used by others to do so.*

3. Graduation Pledge of Social and Environmental Responsibility

Graduation Pledge Alliance, www.graduationpledge.org

I pledge to explore and take into account the social and environmental consequences of any job I consider and will try to improve these aspects of any organizations for which I work.

4. Charles Schwartz's oath for UCB physics students

The purpose of science should be the general enhancement of life and not the causing of harm to man. I affirm that I will uphold this principle, in teaching and in practice of my science, to the best of my ability and judgement. (1970)

Historical note: Prof. Schwartz was formally reprimanded by UCB Chancellor Robert W. Heynes, for requiring students in a physics class he taught, Physics 222, to take this oath. See the interview at www.aip.org/history/ohilist/5913.html for a fascinating oral history (interview with Dr. Charles Schwartz by Patrick Catt at the University of California, Berkeley, on July 19, 1995, Niels Bohr Library & Archives, American Institute of Physics, College Park, MD USA).