Chapter 9

Humans and Automation: System Design and Research Issues

Thomas B. Sheridan

Wiley Series in System Engineering and Management
HFES Issues in Human Factors and Ergonomics Series

Chapter 9

SOCIAL ISSUES OF HUMANS AND AUTOMATION

Social Benefits of Automation

The Obvious Advantages of Automation

Automation has obvious advantages over humans, which are usually the reasons given for implementing it in any given context:

Information-processing speed. Modern computers compute at 1 GHz or greater. Human deductive capability appears to be much slower, but a general number is impossible to fix. For complex pattern recognition and reasoning, the difference may not be so great.

Movement speed. Modern robots respond at 1 kHz, whereas humans respond at 1 Hz, 1000 times slower.

Mechanical power. A robot can lift and carry many tons, whereas a human can handle no more than about 45 kg (100 pounds).

Accuracy and precision. A machine can work at the atomic scale, whereas a human has barely enough visual acuity and dexterity to thread a needle.

Duration of effort. Both information and mechanical automation can work steadily until breakdown from wear. A human loses attention after 30 min of monitoring and, in any case, needs to take breaks for eating and sleeping.

Robustness to environmental conditions. Automation is far more robust than the human with respect to both high and low temperature, high and low pressure, vibration, and ionizing radiation. The human can endure only within a relatively narrow range of these variables without detriment to health.

Reliability. Performance with respect to all these variables makes automation more reliable, at least within the nominal range of its capabilities.

Dollar cost. Automation is usually cheaper than a human in those tasks within the range of machine capabilities, for all the reasons mentioned. This includes consideration of the peripheral infrastructure required to maintain the machine (routine maintenance and repair, etc.) as compared with that required for human labor (food, housing, toilet facilities, medical benefits, insurance, recreation, etc.).

Enabling of Remote Control by Humans

Modern broadband communication means that automation can be put to work anywhere and controlled from anywhere. Computers can now be accessed over the Internet from any terminal to provide whatever information they have been programmed to provide. Air traffic control centers and electrical power dispatching centers need not be located near where the end operations are performed but can be located at sites most convenient from the standpoint of land cost, housing, or transportation for employees.

However, remote control of mechanical devices that perform physical tasks is usually limited to authorized operators. This is true of pilots of remotely controlled air, land, or undersea vehicles; of train dispatchers who operate track switches from centralized control facilities; and of operators in control rooms for manufacturing or process control plants who activate motors or sensors at various locations out in the plants.

Soon, however, the Internet may be used to control physical tasks. It may be difficult today to imagine why one might wish to exercise remote control for heating, cooling, or lighting the home; for cooking, cleaning, watering the lawn, security, and so forth — but surely the capability is there, and more efforts to market such capability will be seen. Perhaps the desire to remotely *monitor* the state of many variables in the home is easier to comprehend. Perhaps more interesting is monitoring other people (children, elderly persons) or pets. Remote monitoring of people raises a special set of ethical issues concerning privacy.

The ability of head-mounted visual displays and their auditory and haptic counterparts to make people feel telepresent at some location other than where they really are, or to allow people to experience presence in a computer-programmed virtual reality, has resurrected old philosophical questions of *being* and *existence* in the field of ontology (Sheridan, 1999).

Understanding Ourselves by Emulating Ourselves

Research in automation and its derivative engineering disciplines, such as artificial intelligence (AI), biomedical engineering, cybernetics, and systems engineering generally, have included many efforts to emulate human and other animal behavior — in sensing, walking, manipulating, and information processing (memory and recall, pattern recognition, decision making). The challenge of building and programming robots is in part motivated by a latent desire of humankind to emulate itself. Though emulating animal behavior is not always the best technological solution

(e.g., modern airplanes do not fly like birds), these efforts do lead to a deeper understanding of humans by humans.

The Turing test for intelligence (named after the British mathematician who invented the Turing machine, a logical analysis of the nature of computers as we know them today) has evolved into a popular annual contest. Computers with different question-answering software compete against each other to try to fool human judges, who freely ask questions (typewritten through a computer network) and judge whether the answerer is a person or a computer.

A current trend in AI is *embodied intelligence*, referring to the capability of robots to appear to have human mannerisms (e.g., with their eyes and body language in responding to human actions in a way that seems humanlike).

These efforts to emulate and understand ourselves are bound to have important long-term effects on science, ethics, religion, and culture in general.

Cathedral Building

In ancient times communities of people labored over long periods to construct cathedrals, temples, pyramids, and other sacred structures to honor their deities and prove their devotion to the greater cause. The results were magnificent feats of technology. Notwithstanding that in many cases the labor was involuntary, these communities did feel pride in the results.

Those efforts might be likened to modern-day space probes, robots, computer chess players, and other high-tech achievements that are built to demonstrate scientific prowess rather than to solve pressing human problems or to sell to commercial markets. Just as the building of cathedrals and temples in the Middle Ages posed technological challenges and demonstrated technical achievement, so too do the modern feats of automation yield community satisfaction in the results. Just as the cathedrals were supposedly built to glorify God (and perhaps a few lesser leaders of the day claimed some attention in the process), the modern technological cathedrals may be said to glorify the scientific theories (with the scientists and engineers hoping to share the glory).

Human Irrationality

Infatuation with Technology

Some people delude themselves into thinking that more capability for control necessarily means better system performance. But anyone who has struggled with an unfamiliar VCR remote control knows otherwise. Still, manufacturers of home appliances play to this human frailty by adding relatively useless "features" to otherwise sensible designs and charging more money. One artist portrays this situation by showing a naked fellow helplessly suspended by a rope inside a towering silo, the walls of which are covered with sophisticated-looking dials, knobs, and buttons. The man's smile suggests a feeling of great pleasure and confidence and a sense of great power. Maybe the knobs, dials, and buttons are not connected to anything, but that may not matter. As a small child one may derive pleasure in make-believe.

Factors that Bias Decision Makers

In Chapter 4, I discussed biases of decision makers, especially with regard to how probability judgments tend to deviate from what is predicted by classical normative decision theory — for example, Bayesian updating (see Appendix A). Researchers have proposed many reasons, some more clearly understood than others. Human-automation system designers should be aware of them. Other factors that affect human decision making are not explicitly biases in probability estimation. Some such factors are the following.

People are unable to keep multiple hypotheses in their minds; they are not good at aggregating data. This may explain why they tend to hedge their bets toward the prior expectation or the mean of the range. In light of the brain's "computational noise," this might be considered rational.

People tend to regard situations as more risky if persons other than themselves are in control, and especially if they are coerced into such situations. People also tend to trust themselves or other people more than they do automation, even though the automation may be reliable. This would be rational behavior if they had no knowledge of the other person's or the automation's capabilities and were risk-averse.

In comparing humans with machines (in determining allocation of functions between humans and machines and assessing reliability of human-machine systems), there is a tendency to regard people as having a few simple failure modes, much as do machines. In fact, humans differ enormously from machines, in that they are inherently variable and unreliable in their detailed behavior, while simultaneously being hyperadaptable and metastable in their overall behavior because they perceive and correct their own errors. Thus they have uncountably many failure modes and are not amenable to being characterized by simple reliability numbers.

People are hedonistic in their decisions. That is, they favor benefit to themselves, their own loved ones, social class, community (nation, race, etc.), and generation over benefit to persons who are removed with respect to these factors. This is certainly what one would expect, but it is not

according to ideals of social equity, morality, and religious altruism. Nor is it in the best interest of propagation of the human species. Loss of a peasant's life to AIDS or famine in a third-world African village and loss of the life of an astronaut, a corporate chief executive, or a president in a first-world nation tend to be evaluated differently by many orders of magnitude, at least in terms of the money allocated to protect against such loss. Although this is a bias, from purely economic criteria, some might judge it to be rational.

Most people exhibit a lack of understanding of the basic concepts of decision science, such as probability and utility. Further, they often seem to lack an evidentiary basis for their own beliefs, but they are willing to accept hearsay, fashion, and isolated sensational anecdotes in forming their beliefs.

An Impediment of Democracy: Lack of Standardization

In some Western nations the principles of democracy can lead to irrationality. An example is found in the construction of nuclear power plants in the United States. Every U.S. nuclear power plant is essentially different because architect-engineers have selected, cafeteria style, from a wide array of vendors for each piece of equipment, making up a unique hodgepodge combination for each plant. Even in plants located immediately next to each other and operated by the same management, such as the two on Pennsylvania's Three Mile Island, the combination is different from one plant to the other. This contrasts with France, where nuclear plants are standardized.

Variety is fine for homes, clothing, or other areas of life in which variety and self-expression add to the culture. However, for nuclear power plants, transport aircraft, or automobiles, where safety and predictability are essential, variety is wasteful in cost and time and hinders understanding of safety and reliability. Standardization has proven to lead to greater safety and cost-effectiveness.

Quantitative Modeling as Sorcery for the Powerful

Technologists consider it their business to quantify things. They are hired to do this. Quantification is also their identity, their pleasure, and their reward, so they are naturally motivated to quantify whenever the opportunity presents itself. I include myself in this group. The clear and present danger is to overdo it.

Mathematical models are sometimes paraded before politicians to present a facade of objectivity. This may be done with full knowledge that the audience is incompetent to appreciate the limitations of the models or to ask meaningful questions about their validity. Sometimes this is called "lying with statistics." One can liken this practice to the kings of old employing court sorcerers, who used incense and uttered strange incantations to frighten royal visitors and ensure the king's power.

Some nations, the United States being a prime example, seem compelled to put numbers on safety of nuclear plants and other complex regulated systems of humans and machines when public safety is of concern. Sometimes the numbers have little validity. Other nations (I experienced expressions by Japanese officials of this viewpoint) seem satisfied that if the people responsible for system safety are trying their hardest and doing their best with available resources, that is good enough, and therefore their system is acceptably safe.

The obligation for the technically literate, it seems to me, is to help bring rationality and truth to all discussions of model predictability, especially as concerns safety and security. One hundred percent predictability or safety, which some politicians pompously and irresponsibly demand, is unattainable; it is possible, however, to continue to make models and systems increasingly better. Absolute reliability in predictions about relatively new or changing technology or complex human-machine interactions, for which only the most primitive analytical tools are available, make little sense. What is more believable, in my judgment, is sensitivity analysis, which assumes and admits to some rough reliability number, then analyzes whether that number increases or decreases as a function of certain actions that are taken to configure or to control the system.

For some kinds of technology, reliability modeling seems almost infeasible. One particularly worth mentioning is computer reliability. Not so many years ago, I experienced how, when planning intelligent transportation systems, automotive engineers (who fully appreciated reliability problems with brakes and concrete structures) had the idea that once a computer program was compiled, it was guaranteed to be reliable. They did not appreciate that the number of failure modes of modern computers interacting with external environments, especially with human users, is virtually infinite — and can never be exhaustively evaluated.

Advocacy versus Objectivity

Scientists and engineers like to claim objectivity. Perhaps that is the reason many interests of human users and human operators are neglected — the designer simply hasn't found an objective way to accommodate the "soft" considerations. So those considerations are ignored, and the effort is made to "optimize" on the basis of only those variables that can be quantized.

However, such a practice is suboptimal. Just because a variable cannot be quantized does not mean it is unimportant; indeed, if it is left out of consideration, the system may not be optimized.

In contrast to science, Western jurisprudence operates not on the basis of objectivity but on the basis of advocacy — advocates for various sides of a contest try to convince judges or juries. Everyone knows that real decision making in technological design and implementation involves advocacy at many levels. Unfortunately, the human user or operator has often been left without an advocate. This is changing. The human factors profession, while trying for objectivity when it is feasible, has taken to advocacy when worthwhile considerations do not lend to objectivity. So have unions, government regulators, and consumer groups. The need is for a greater understanding of when objectivity can be employed and when a fair process of advocacy should be invoked.

Peer Influence

The magnitude of peer social influence on people of all ages is well established. Although people often decry the extent of peer influence on children, we allow ourselves to be influenced by the media and by our own local cultures much more than we might admit.

A simple and oft-repeated psychophysical experiment is illustrative. Two parallel lines of slightly different length are drawn on a piece of paper. The experimental participant is allowed to get close enough so that without any external influence, he or she would almost always be able to select which is the longer of the two. However, the experimenter leaks a false report to the participant indicating that the majority of participants selected as the longer line the one that is actually shorter. Experimental results almost always show that participants so treated are heavily influenced by the leaked data and therefore make the wrong selection — even though their direct visual perception tells them otherwise.

Safety: Normative Objectivity versus Media-Motivated Acceptability

Safety, to a risk analyst (*normative* safety), is the expected value of avoiding bad consequences, where weighting of different consequences (apart from probability) is based on some combination of human suffering and economic loss. In the real world, what is safe is what is judged to be acceptably safe. Social belief and social acceptability seem to be the ultimate criteria of safety, quite apart from the product of probabilities and consequences — but what underlies social acceptability?

There is something about sensational accident situations that gives rise to far more importance weighting than is warranted, at least in comparison with mundane accident situations. A commercial airplane accident is picked up by the wire services and makes headlines everywhere. It is sensational because it is unexpected (and therefore, in an information theory sense, amounts to lots of information). An isolated single- or multiple-fatality automobile accident often will only make the back pages of a local newspaper. It is expected and accepted that there will be many automobile accidents and that some will result in loss of life, so there is not much information there. Somehow, the seriousness of an accident, like the seriousness of any public event, is judged according to the extent of its media coverage, which in turn is truly dependent on its unexpectedness or information magnitude, apart from the extent of deaths, injuries, or property damage.

So, does the multiple-fatality accident garner media importance out of proportion to the number of fatalities? Not always. Strangely, some people see it the other way around. I recall executives of a particular non-U.S. national rail system expressing the opinion that a train wreck in which N people are killed is not equivalent to N train wrecks in each of which one person is killed. They saw all train wrecks as almost equally bad — in terms of media coverage. Maybe the public empathizes as much with the size of the physical wreckage as it does with the number of fatalities!

Alienation of the Individual

The term *alienation* is used in this section to characterize a variety of negative impacts of automation on the individual human. When systems are automated and the human is removed not only spatially but also temporally, functionally, and cognitively from the ongoing physical process he or she is directing, that can be called *alienation*. This is because what the human does and thinks are likely to have different timing, different explicit form, and different logical content from what the automation does and thinks. The machine's behavior is alien to the human's.

People who have been pushed sideways into jobs that do not make use of the skills on which they pride themselves and with which they identify will feel frustrated and resentful. People who no longer understand the basis of what they are asked to do will become confused and distrustful. Especially if they perceive a powerful computer to be mediating between them and what is being produced at the other end, people become mystified about how things work, question their own influence and accountability, and ultimately abandon their sense of responsibility. All this can lead to alienation (Sheridan, 1980; Zuboff, 1988). The components of potential alienation are detailed in the following sections.

Threatened or Actual Unemployment

This is the factor that has been most often cited in opposition to automation. In the English knitting mills of the mid-19th century, angry mobs, allegedly led by "King Ludd," rioted and smashed the newly "automated" machines out of fear of losing their jobs. They came to be known as *Luddites*, and the term has come to be used against any group opposing the advance of technology, especially out of fear of unemployment.

For many years the advance of automation seemed to have little significant impact on unemployment, because people replaced by automation were simply moved to other jobs, such as machine supervision and maintenance. However, organizations have become more efficient in their use of human supervisors: Fewer people can now supervise more machines. Furthermore, automation is becoming more able to detect its own failures, and in some cases it can even repair itself. The real threat today is unemployment of the unskilled and technologically illiterate.

Erratic Mental Workload and Work Dissatisfaction

Automation affects not only the nature but also the pace of work, and it makes that pace vary between extremes. Airline pilots and nuclear plant operators refer to their work as "hours of boredom punctuated by moments of terror." Some workers may feel that although automation has eliminated many mundane tasks from their work, it has not relieved the stress induced by on-the-job decisions. For example, some pilots feel that in times of emergency or high stress, the automation is not sufficient to ensure safety and that even more human involvement is needed than in earlier times of pure manual control.

Although manufacturing workers are often glad that automation has broken the dull monotony of repetitive manual work, they may not get as much satisfaction from their new supervisory monitoring role. The satisfaction derived from putting in a hard day's work and producing something with their own wits and efforts may be gone.

Already, some airline pilots are complaining about automation, some manufacturing engineers are having second thoughts about robots, and some power plant operators who have experienced older, simpler, and very reliable plants are wondering if things haven't gone too far.

Centralization of Management Control and Loss of Worker Control

One result of automation and the introduction of electronic technology often feared by workers is the possibility that management can secretly record and monitor their work output through the equipment's sensors and communication capabilities. Initially managers assured workers that they did not engage in such activities, but today it is clear that productivity data are so easy to record that they are being recorded. Such monitoring is occurring not only in industrial plants but also in hospitals, offices, trucks, and aircraft, where high-bandwidth communication is becoming plentiful and makes such monitoring from a central location easy.

The mere possibility of being monitored in this way is often sufficient to produce worker anxiety, including fear that private data stored electronically may be accessed by persons other than those authorized.

Centralization of monitoring and control, so easy with advanced automation, is often seen by management as synonymous with efficiency. In some cases centralization may enhance productivity, but in other cases it may prove detrimental.

Desocialization

Although cockpits and control rooms now require teams of two or three, the trend is toward fewer people per team, and eventually one person will be adequate in many workplaces. This is certainly true in factories and offices. Thus cognitive interaction with computers is gradually replacing that with other people. As supervisory control systems are interconnected, the computer will mediate increasingly more of what interpersonal contact remains, as has already happened in many cases with e-mail, pagers, and associated software for management coordination and computer-supported cooperative work. Although interpersonal communication is in one sense becoming easier, in another sense it is becoming artificial (Turkle, 1995).

Deskilling

Skilled workers "promoted" to supervisory controllers (sometimes derogatorily referred to as "button pushers") may resent the transition. In part, this may be out of fear that when called on to take over and do the job manually, they may not be able to do so. A skill such as manual machining, developed over a long period, provides the worker a sense of dignity and self-respect. The same may be said for the seamstress or other manual artisan. If the skilled worker becomes a button-pushing supervisor and monitor, that original sensory-motor skill may atrophy. Many researchers have decried the tendency for skill to atrophy when automation is used (Bainbridge, 1987; Kaber & Riley, 1999; Wiener & Curry, 1980).

Intimidation of Greater Power

Automation and supervisory control encourage larger aggregations of interconnected equipment, higher speeds, greater complexity, and probably greater economic risk if something goes wrong and the supervisor doesn't

take the appropriate corrective action. The human supervisor will be forced to assume increasingly more ultimate responsibility (although in most cases the responsibility probably should reside with some combination of the manager and the system designer). Depending on one's personality, this could be very intimidating and lead to stress, inattention to detail, anxiety that one is not up to the job's requirements, or paranoia.

Technological Illiteracy

In older manual control systems, the operator could understand how things worked. In the role of supervisory controller, the operator may lack the technological understanding of how the computer and the rest of the complex technology do what they do. What is really going on with the communications and control software may be too specialized even for most technicians involved with the newer systems. The automobile mechanic's inability to understand the computer/electronic systems in modern automobiles is an example. Operators may come to resent this and to resent the elite class of technologists who do understand.

Mystification and Misplaced Trust

Human operators of computer-based systems sometimes become mystified by, and superstitious about, the power of the computer, even seeing it as a kind of magic or a "big brother" authority figure. This leads naturally to naive and misplaced trust. This was particularly well articulated in Norbert Wiener's (1964) *God and Golem, Incorporated*, in which he used a classic in horror literature, W. W. Jacobs's *The Monkey's Paw*, as a metaphor:

"In this tale, an English working family sits down to dinner in its kitchen. The son leaves to work at a factory, and the old parents listen to the tales of their guest, a sergeant major back from service in the Indian army. He tells of Indian magic and shows them a dried monkey's paw which, he tells them, is a talisman that has been endowed by an Indian holy man with the virtue of giving three wishes to each of three successive owners. This, he says, was to prove the folly of defying fate.

"He claims he does not know what were the first two wishes of the first owner, but that the last one was for death. He himself was the second owner, but his experiences were too terrible to relate. He is about to cast the paw on the coal fire when his host retrieves it, and despite all the sergeant-major can do, wishes for £200.

"Shortly thereafter there is a knock at the door. A very solemn gentleman is there from the company that has employed his son. As gently as he can, he breaks the news that the son has been killed in an accident at the factory. Without recognizing any responsibility in the matter, the company offers its sympathy and £200 as a solatium . . .

"The theme (here) is the danger of magic. This seems to lie in the fact that the operation of magic is singularly literal-minded, and that if it grants you anything at all it grants what you asked for, not what you should have asked for or what you intend. If you ask for £200 and do not express the condition that you do not wish it at the cost of the life of your son, £200 you will get whether your son lives or dies.

"The magic of automation, and in particular the magic of automatization in which the devices learn, may be expected to be similarly literal-minded" (p. 58).

To a naive user the computer can be simultaneously so wonderful as to seem faultless, and if the computer produces other than what its user expects, that can be attributed to its superior wisdom. Such discrepancies are usually harmless, but if they are allowed to continue they can, in some complex and highly interconnected systems, endanger lives. It is therefore crucial, as new computer and control technology is introduced, that it come to be accepted by users for what it is — a tool meant to serve and be controlled ultimately by human beings.

Sense of Not Contributing

Though the efficiency and mechanical productivity of a new supervisory control system may far exceed that of an earlier manually controlled system, the operator may come to feel that with automation he or she is no longer the source of value added, no longer a significant contributor. This sense of deprivation has been expressed by factory workers, farmers, craftspeople, accountants, and others who feel threatened by automation. The sense of personal productivity — what psychologist Erich Fromm (1995) called the *productive orientation* — is allegedly fundamental to humans' sense of self-worth. Without it, who are we?

Abandonment of Responsibility

As a result of the factors just described, human supervisors of automation may eventually feel they are no longer responsible for what happens but that the computers are. A worker with his or her own set of hand tools or a simple self-powered but manually controlled machine, though sometimes placing the blame for difficulties elsewhere, has a clear responsibility for use and maintenance of the tools or machine. That worker is accountable for what is produced, and everyone knows it.

When workers' actions in using a machine are mediated by a powerful computer, however, the lines of responsibility are not so clear, and the workers may not be sure which should get the credit or the blame for a

situation — the computer or themselves. As a result these workers may, in effect, abandon their responsibility for the task performed or the good produced, believing instead that it is in the "hands" of the computer. Even when computers are installed to aid information flow from one worker to another, or to act as a processor or storer of data, individuals using the system may feel that the machine is in complete control and disclaim personal accountability for any error or performance degradation.

Blissful Enslavement

To many writers the worst form of alienation, the worst tragedy, occurs when a worker is happy to accept a role in which he or she is made to feel powerful but, in actuality, he or she is enslaved. Engelberger (1981), the founder of the world's first major robot company, reminded readers that "it will always be far easier to make a robot of a man rather than to make a robot like a man." Both Huxley's *Brave New World* and Orwell's *1984* are famous for the theme of blissful enslavement.

With regard to the affects of the computer on a person's self-perception, I am fond of citing historian Mazlish (1967), who referred to the computer as the "fourth discontinuity" in this self-perception. At first, said Mazlish, humans saw themselves as the center of all things, but Copernicus jarred this perception by showing that the human race was an isolated dweller of a tiny planet of a minor star (which turned out to be at the edge of only one galaxy). Darwin came along and rudely asserted that we were descended from the apes and they, from lower creatures still. Later Freud suggested that we humans are not even consciously in charge of our own faculties, that our egos and ids drive us. Now the computer may be dealing us the ultimate affront — surpassing us intellectually — beating us in our best suit, and seemingly inviting us to give in.

Surely automation promises greater product quality, better energy and economic efficiency, and improved worker safety — all benefits that motivate development of the new technology. Nevertheless, these potential negative impacts must be examined and reduced in order to achieve a satisfactory acceptance of automation by society.

Alienation of the Community

Automation can be alienating not only to individuals but also to communities of people as groups. Using that somewhat arbitrary dichotomy, this section considers additional forms of automation alienation.

Distributive Justice: Who Decides and Who Benefits

Automation in any particular application has its proponents and its opponents. The proponents are usually the technologists who design, build, market, install, and manage it. They are the ones who usually benefit directly from its use. The operators can be opponents, as discussed in the previous section, or proponents, seeing themselves as its caretakers. Others, who have no direct participation in it, may be threatened by various aspects of it, because of economic competition, because they are jealous or intimidated, or for other reasons. In a democracy it is an individual's or a corporate entity's right to automate, so long as legal constraints on safety and environmental impacts are met. However, as with any human endeavor that affects many people, there is bound to be disagreement over who gains and who loses, who decides and who has the decision forced upon them. These are issues of *distributive justice* (Rawls, 1999), and they are not easy.

Figure 9.1 illustrates the problem of distributive justice that occurs when rich nations or individuals have the power to automate and poor nations or individuals do not.

Questions of distributive justice occur at two levels. At the global level, implied by the preceding paragraph, all affected parties are included. However, there is a lower level having to do with the allocation of power among system designer, system manager, and human operator of the automation. There is a tension between the prescriptions and proscriptions explicitly specified by the designer (or implicit in the system design constraints), the policies and constraints established by system management, and the freedom of the operator (free will) to make and execute decisions in the course of system operation. How can this tension be resolved?

Technophiles and Technophobes

Perhaps it is good that some people are *technophiles* and others are *technophobes*. The technophiles want to exercise technology wherever possible, seeing it as a challenge. However, they can succumb to the temptations of promoting the technology for self-gain at the expense of others. The technophobes see technology as a threat to tradition and to well-established and comfortable ways of doing business. They might have society return to what they believe would be a safer, healthier world without technology, modern travel and commerce, and large-scale interconnected communications, power, food, and water supply systems that are vulnerable to terrorists. However, one cannot escape the fact that individual longevity and health are far better than they were in earlier times, and in any case the clock cannot be turned back.

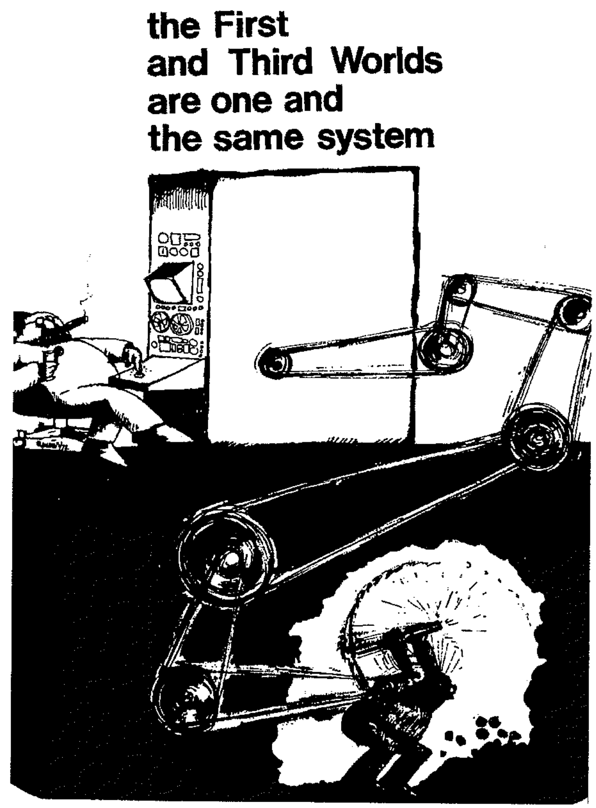

Figure 9.1 A political cartoon by Claudius (Brazil). Used with permission.

Many find themselves between the extremes, maintaining a neutral stance. Sometimes they are conscious of the evidence tugging them in two directions. Sometimes their neutrality is based on ignorance or apathy. Although the media may provoke a war of words between the two extremes, the public is often confused by the arguments and is made to feel left out. Nevertheless, the encounter between technophiles and technophobes should be useful for elucidating the truth and directing compromise toward sanity in public policy.

Machine Productivity instead of Human Productivity

Today national governments and multinational corporations are obsessed with productivity. Balance of payments and national well-being are at stake. Computers integrated with sensors, actuators, and robots are allegedly there to make the factories, power and chemical plants, aircraft, ships, automobiles, banks, hospitals, schools, and homes more productive. Most people consider economic productivity to be the overriding social goal.

What kind of productivity is being increased? Is it the kind of productivity that at the end of the day — and indeed at the end of our lives — gives us the satisfaction of having best served other human beings? Is machine productivity synonymous with human productivity? Or are the two different — perhaps orthogonal, or even in conflict with one another?

Daily Living by Remote Control: Reduction of Social Contact

Future constraints on energy potentially will limit free travel. However, with microelectronics advancing communication in so many ways, people have a new capability to use telephones, fax machines, desktop computer terminals, and portable communication devices to order fast food or supplies delivered to our homes, to shop from electronic catalogs, to buy and sell stocks, to pay bills or make bank transfers, and to do e-business without leaving home. At the moment these personal remote control activities are not regarded the same as remote control of industrial automation that does useful mechanical work but, rather, as an extension of the telephone for communication purposes.

Future scenarios may include remote control of robots — to browse up and down the aisles of the supermarket or store to inspect and handle merchandise through remote video links, to cut the grass, or to take out the trash. It is too early to predict the benefits and hazards of such radical changes in living style. However, no significant breakthroughs in technology are required for such a scenario to occur.

The tendency of people to perform many more daily functions by remote control requires fundamental rethinking of the meaning of home, commerce, transportation, education, and human fulfillment. It probably spells a significant reduction of casual social contact with others in our communities. Our social contacts may have to become planned and controlled to a much greater extent. Does this translate to psychological or political dangers lurking ahead?

Electronic Telegovernance by the Powerful over the Powerless

The use of modern automated communication technology can provide not only feed-forward of information from those in power down to the masses of citizens but also feedback from the masses to the powerful

(DeSola Pool, 1973; Sheridan, 1975). Many experiments have been conducted with people watching television programs and "talking back" through the telephone or the Internet to vote on questions posed. Conceivably, given the bandwidth already available, democratic governance could extend to all citizens and be essentially instantaneous, possibly even eliminating layers of elected representatives.

This vision may have advantages for some purposes — for example, to educate and test public opinion on certain issues. However, such possibilities of electronic governance also conjure up visions of a society that could rapidly put too much power in the hands of whoever controlled the head end of the communication system. Indeed, such arguments have already been made against the national media networks, and similar arguments underlay the motivations for breaking up the AT&T and (alleged) Microsoft monopolies.

Automation of Ecoexploitation

The year 1970 saw the beginning of great international debate among scholars about "limits to growth" hosted by the Club of Rome and other international study organizations (Oltmans, 1974). The issue of exploitation of natural resources was central, and the role of automation was a great concern. Some assert that the debate fizzled and came to naught, whereas others claim there were lasting revelations and lessons learned. Most planners of future automation systems believe it is especially important that the so-called technological imperative — that automatic systems be built because they can be built — be restrained, especially with respect to use of energy and other precious natural resources. Many nations are lucky enough to have many such resources. Others are not so lucky.

Proper use of automation will produce the same or greater benefits for people while reducing demands on energy and other (especially nonrenewable) resources. However, automation makes it easy to consume energy and other resources for questionable ends. It would be criminal to employ automation that, for the sake of short-term productivity, gluttonously exhausts nonrenewable resources and leaves none for less-developed nations and for future generations. For example, certain wealthy nations now recover cobalt and other mineral nodules from the ocean bottom using high-tech robotic suction devices. In some cases this happens just offshore of poorer nations that are powerless to stop it.

Society may be said to be facing a new stage of Hardin's (1968) oft-quoted "tragedy of the commons." English sheepherders unintentionally increased the number of sheep beyond what the common grazing lands could support by each adding a number that was (to the sheepherder) insignificant. Now the technologically advanced nations are automating and expanding their use of automatic industrial processes that foul the air and water and use up the oil and other natural resources. Would global warming progress as fast if there were no automation? The metaphor of the sorcerer's apprentice in Walt Disney's *Fantasia* also comes to mind: an unstable self-multiplication of trouble.

Telerobotic Soldiers, Spies, and Saboteurs

Consider a future manually controlled teleoperator that has good mobility on land, in the sea, or in the air or in space — all are possible, all are being developed. Assume wireless telecommunication, reasonable sensors, and some manipulative dexterity (full telepresence is not necessary). Assume a good battery or other power supply, and the possibility that the device can be controlled from afar to "plug itself in" from time to time to reenergize itself. This teleoperator may be driven by its human operator to do mischief in anyone's backyard. Its human operator need have no empathy for its welfare, for it is only a mechanical slave. Further, at any time its human operator can cease to communicate with it and with ease can abandon responsibility and accountability for its behavior.

Now consider a future telerobot, in this case with a computer to help it see, hear, and touch, but also having some motor reflex skills to conserve energy and adapt to its environment and having enough knowledge and instruction to understand how to implement its programmed goals in spite of disturbances or obstacles it encounters. This telerobot may be sent to do even more mischief because it can go surreptitiously or lie dormant for long periods without the need for communication. It can be used to spy, sabotage, set explosives, and perform a host of other duties. Because it is controlled in supervisory fashion, any one human supervisor can multiply his or her own capability manyfold. In this case it is even easier for human operators to abandon responsibility. They can claim to be too busy with their brood of telerobots, and in any case each telerobot can be said to have a mind of its own.

In past arguments, fights, and wars, individuals who felt they must resort to violence toward other individuals put themselves at some bodily risk in doing so — at least, more risk than those not so inclined. Traditionally such behavior has been lauded as courage or bravery and at least posed a high cost on the initiator. However, technology has been changing all that, as weapon telerobots have evolved from clubs to arrows to bullets to bombs dropped from airplanes to smart cruise missiles. The human operators no longer put themselves at risk so much as they jeopardize the safety of others. Telerobot technology is merely the next step.

Some have heralded the dawn of smart weapons, battlefield robotics, and "telegladiators" as a new day when international disputes will be fought on a large technological playing field (space, undersea, or surely somewhere well away from real people). Technological prowess will win the day, and no real people (at least on one's own side) will get hurt. Certain military planners seem to be moving toward that fantasy even now.

Already there are telerobotic sonar spying devices in the oceans and similar electromagnetic spying devices in orbit — all over the protests of the helpless poorer nations. Now space is becoming militarized — perhaps eventually in the form of an ultimate human-machine system programmed in top-down supervisory fashion with the purpose of defending its nation but demanding, ultimately, that people abandon responsibility to the computer. I cannot help but recall again the theme of Wiener's *God and Golem, Incorporated* (1964), mentioned previously in this chapter.

Terrorism and the Vulnerability of a Variety of Automated Systems

Since September 11, 2001, the Western world has realized that many highly automated systems throughout the fabric of society are vulnerable to terrorists and that large systems involving both many people and much capital investment can be shut down through actions that may be relatively simple for hostile agents to execute. Transportation systems, communication systems, postal systems, and supply systems for food, water, electricity, and gas are all highly automated and therefore at risk.

We had already experienced somewhat milder terrorist acts in cyberspace. A number of denial of service attacks have already been effected against commercial Internet sites — the perpetrator secretively planting automatic software agents in the computers of thousands of unsuspecting individuals and synchronizing them to launch a simultaneous attack on their victims. It is extremely difficult to trace the source of such attacks.

Major banks and financial service organizations have been forced to spend many millions of dollars to build firewalls to prevent unauthorized access to their communication networks, which are responsible for secure transfer of billions of dollars per hour. Now the threats to large human-machine systems are far greater. Military and civilian security organizations now freely admit that their own "red teams" (used in war games) have succeeded in invading control systems for electrical power distribution grids, hydroelectric dams, transportation control centers, and even defense facilities, all of which had been assumed to be secure.

Unfortunately, society is entering a new era when we will have to spend large sums of money and human capital to ensure the freedoms and openness that we have experienced in recent years. We may also have to sacrifice some of our freedoms in order to do so.