Articles

Perceived Risk, Real Risk: Social Science and the Art of Probabilistic Risk Assessment

WILLIAM R. FREUDENBURG

Risk assessment is commonly seen as the domain of physical and biological sciences, with social scientists focusing instead on risk management and communication. This division is unnecessary, and it may lead to errors in risk assessments. Social science input is needed for more accurate calculations of risk consequences and probabilities and for identifying potential biases created by certain risk assessment procedures, as well as in analyzing and explaining public responses to risk. Findings, moreover, suggest that the dichotomy between "real" and "perceived" risk is less "real" than is often assumed, particularly in cases involving controversial technologies.

"WHY DOESN'T ANYBODY BELIEVE US ANYMORE?" This question was recently asked by a Ph.D.-level biologist, a man who has risen over the years to a position of considerable authority in a federal resource management agency. The specific context was the public insistence that his agency stop promoting a "risky technology," even though the evidence convinced him the risk was low. The broader problem is that he is not the only scientist asking such questions lately.

As the average life-span increases, the public perceives many of the risks around them to have become more severe (*1*). As scientists, we often assume that public perceptions are simply at variance with the real risks. The scientific method, however, calls for us to test our assumptions before accepting them uncritically, and there is "no law of nature that requires us to abandon the scientific method merely because questions of human behavior are involved" (*2*, p. 27).

Instead, this is a problem that calls for input from the social and behavioral sciences, which offer at least three contributions to the burgeoning field of risk assessment. The first and most obvious is in providing tools—and increasingly, a set of relevant findings—to help clarify the differences between the scientific community and the general public in the assessment of technological risks. Such a contribution is generally seen to lie in the area of risk communication, risk perception, and risk management. A second but less-understood contribution, however, comes from social scientists' input to risk assessments themselves—to actual calculations of the probabilities and consequences of undesired outcomes. Third, social sciences offer insights into the processes by which risk assessments are carried out. In each of these three areas, the cost of ignoring social scientists' contributions can be an unnecessary bias.

In this article I discuss three potential sources of error—in calculating risk consequences, in calculating risk probabilities, and in paying insufficient attention to the person-intensive nature of performing the assessments themselves. I conclude with a discussion of the rationality of public risk perceptions.

The author is an associate professor in the Department of Rural Sociology, University of Wisconsin, Madison, WI 53706.

Risk Consequences

The potential social and economic consequences of technological failures tend to be reasonably easy to comprehend and to be increasingly recognized in the technical community. Three categories of consequences, however, deserve greater attention than they have received to date.

Impacts of serious accidents. Easiest to consider are the consequences of rare but genuinely serious accidents—including accidents less extreme than those at Bhopal or Chernobyl. An examination of the Goiânia event in Brazil (*3*) is instructive. Two men entered an abandoned clinic in search of scrap metal; they found a small capsule and pried it open, releasing approximately 100 grams of cesium-137. This led to 121 cases of skin contact with the material and four deaths, with another three to five expected in the next 5 years. The death toll was not wildly out of line with "any other industrial accident," but just the labor costs of decontamination exceeded $20 million (U.S.) by December 1987. The broader economic and social costs were far greater. Within 2 weeks of the time the event was announced, the wholesale value of agricultural products from the entire state of Goiás fell by 50%; demand for manufactured goods (including textiles, clothing, and other finished products) was also affected—even though the study was unable to find "even a published suggestion" that the agricultural products or manufactured goods could have been contaminated (*3*). Severe impacts were also felt through treatment and research costs, declining property values, and a decline in the tourist trade. More than 100,000 residents lined up at monitoring stations to be checked for radioactive contamination, and more than 8,000 residents requested and received certification that they were not contaminated (*4*).

Uncertainty costs. In the second category are the less obvious costs of dealing with uncertainty and risk, even when "nothing goes wrong." When people buy automobile, health, home, or fire insurance, they incur real costs even if the insurance is never "needed." Insurance companies keep the premiums even if the house does not burn down, the automobile is not involved in an accident, the insured person is not hospitalized, and so on. Insurance premiums provide examples of costs created by the possibility that something may go wrong, not by the actual occurrence of the event itself (*5*). Similarly, real costs are incurred when communities invest in emergency-preparedness training or the preparation of evacuation plans, when societal strains are created by inequitable distributions of technological risks, or even when individuals "invest" in the psychic costs of worrying about potential disasters, whether such disasters actually occur or not (*6*).

"Signal" incidents. The third category includes a special class of the cases in between: not cases where best hopes or worst fears are realized, but cases where problems indicate that a technology may not be fully under control. Slovic refers to such events as having "signal" value—as signifying or sending a signal to the broader public that there may be reasons for concern (*7*)—and Kasperson *et al.* note that such events have the potential to lead to the "social amplification of risk" (*8*).

The explosion of the space shuttle Challenger sent a "signal" not only to the public and the policy system, but to the scientific and technical community as well. In the context of major airline accidents, the quantitative death toll from the Challenger was small, yet the accident led to an expensive, lengthy reanalysis of the nation's space program. A clearer example is the accident at Three Mile Island (TMI), which was found by official investigations to have released very little radioactivity, although it did lead to significant mental health consequences for nearby populations (*9*). More broadly, TMI sent a signal that nuclear power plants were less safe than the public had previously been led to believe (*10*), and its consequences were little short of disastrous for the nuclear power industry.

Risk Probabilities

Less widely recognized is the need for social science input in calculating risk probabilities. Three sets of factors illustrate the types of systematic biases that can be created by the failure to pay adequate attention to human behaviors.

Human error and human factors. When the President's commission on the accident at TMI began its investigation, its members expected to focus on problems of "the technology," meaning the nuclear hardware, rather than the ways in which the hardware was managed and operated. Ultimately, however, the commission found "the fundamental problems [were] people-related problems and not equipment problems" (*11*, p. 8). Included were several human errors that helped to turn a malfunctioning valve into the most expensive accident in the history of domestic nuclear power.

"Human error" is a value-laden term that has often been used to describe situations that actually involve mismatches between people and machinery, including those at TMI (*12*). With this caveat in mind, however, "human error" has been implicated in accidents that range from the Chernobyl disaster to the more prosaic problems of transporting hazardous materials, where more than 50% of the risk may be due to "driver error" and other human factors (*13*). Such figures make it clear that attention to "hardware" of vehicles, sign designs, road materials, and so forth, although valuable, may address only the minority of the causes of transportation accidents or other technological risks.

Organizational error and "organizational factors." The behaviors that need to be explained, moreover, include not just those of individuals, but those of organizations. Empirical research on organizational and institutional factors often leads to conclusions that appear counterintuitive to persons trained in non–social science fields (*14*). The people problems identified at TMI were not limited to the operators in the plant at the time of the accident, but included "problems with the 'system' that manufactures, operates, and regulates nuclear power plants. There are structural problems in the various organizations, there are deficiencies in various processes, and there is a lack of communication among key individuals and groups" (*11*, p. 8; *15*). Even after TMI, some utilities have operated nuclear facilities well, but others have been less successful. The Peach Bottom nuclear power plant, 60 kilometers downstream from TMI, was extensively criticized both by the Nuclear Regulatory Commission (NRC) and the nuclear industry's own Institute of Nuclear Power Operations, which called the plant "an embarrassment to the industry and to the nation," calling attention to "the very culture of the company and how it was managed" (*16*, p. 6; *17*). Problems of operators sleeping on the job were eventually reported not by the utility, but by engineers from the reactor manufacturer, leading the NRC to close down the plant in March 1987, "the first time a nuclear power plant had been closed by the agency for a nonmechanical reason" (*16*, p. 6).

In another example that received broad attention, the immediate cause of the loss of the Challenger may have been the failure of an O-ring, but the official investigation into the accident also found problems in the broader management and organization of the National Aeronautics and Space Administration (NASA). Investigators recommended sweeping reforms of NASA's management structure, flight operations, and safety and risk analysis procedures—in addition to changes in the O-ring itself (*18*).

More broadly, it is no surprise to persons who study complex organizations, although it sometimes surprises or frustrates many of us in the technical community, that behaviors of janitors, operational personnel, and others are often different in empirical reality than in organizational policy. Such departures from official expectations can cause the actual risks to be significantly different from risk estimates that are based on the assumption that official policies will simply be followed (*19*).

External social factors. Receiving still less attention are behaviors outside of the responsible organizations altogether. The obvious examples are threats of terrorism and sabotage. Although risks from terrorism appear to be relatively low in the United States, at least to date, they may be nonnegligible, particularly for technologies that are highly controversial. Disruptions might be more likely in cases where people feel they have not been accorded an adequate voice through traditional or acceptable channels, where they believe they are responding to a higher form of morality, or where they are interested in altering public perceptions of a technology's safety. Probabilities might be further increased for acts of sabotage that carry relatively low risks of capture but have high signal value—as in ambushing nuclear waste trucks on a little-patrolled section of an interstate highway or contaminating city water supplies with genetically engineered organisms stolen from experimental farms. For the longer term, there is a need for systematic research on terrorism, as well as other external factors such as political and economic pressures, to allow decisions to be guided by more than conjecture. For the near term, however, prudence suggests that external human factors cannot be safely ignored. For certain types of technologies, the probability of disruption might ultimately prove to be anything but negligible—perhaps closer to the range of several percent per year than to a one-in-a-million chance.

"Human Error" in Estimation Techniques

Intriguingly, the group that generally receives the least attention of all is the one that may have the greatest influence on the assessments—those of us who do the calculations. A growing body of evidence, however, suggests that scientists may be subject to some of the same foibles that affect the general public, and to a few more besides (*20*). Three sets of problems, in particular, are worthy of mention here.

Calibration and overconfidence. Like other human beings, scientists may fail to foresee all factors that can introduce errors into estimates. Examples are provided by difficulties in seeing all the ways in which components of a system are interrelated (*21*), by failing to foresee interactions among individually minor problems (*22*), by temptations to overlook the aspects of technological systems that are "nontechnical" or outside a given field of expertise (*23*), by insufficient sensitivity to the fragility of assumptions or to the problems of small sample sizes (*24*), or even by conscious decisions to simplify analysis by excluding low-probability events from consideration.

Perhaps partly because of these problems, even scientists may have excessive confidence in estimates, as illustrated by the well-developed science of physics and the value of a quantity as fundamental as the speed of light. An instructive compilation found 27 published surveys of the speed of light between 1875 and 1958 that included formal estimates of uncertainty. The measurements differed from the official 1984 value by magnitudes that would be expected to occur less than 0.005 of the time, by chance alone, when the original estimators' own calculations of the uncertainties in their estimates were used (*25*). The straightforward conclusion is that "the respective investigators' uncertainties . . . must [have been] significantly underestimated" (*25*, p. 793). Although the absolute magnitude of the errors declined significantly over time, there was no significant improvement in the degree to which the remaining uncertainty was estimated. The 1984 estimate of the speed of light (which has since been used to calibrate the length of a meter, rather than vice versa) falls entirely outside the range of standard error (1.48 × "probable error") for all the recommended values for the true velocity of light that were reported between 1930 and 1970 (*25*, figure 2).

Other examples can be reported for scientists ranging from engineers to physicians. In one study a group of internationally known geotechnical engineers were asked for their 50% confidence bands on the height of an embankment that would cause a clay foundation to fail; when an actual embankment was built, not one of the expert's bands was broad enough to enclose the true failure height (*26*). In another study, physicians reviewed medical histories and examined more than 1500 patients with coughs; of the group they diagnosed as having more than an 80% chance of having pneumonia, less than 20% actually did (*27*). Other studies of the ability to assess probabilities accurately, the problem of calibration, have found that calibration errors are unaffected by differences in intelligence or expertise (*28*) but may be increased by the importance of the task (*29*). More prosaically, faculty members routinely underestimate the time required for faculty meetings, and I underestimated by 50% the time that would be needed to finish this article, despite having enough experience in writing such papers to be presumed to "know better" (*30*).

Overall, one would expect that only about 2% of the estimates made with a confidence level of 98% would prove to be surprises, but nonspecialist assessors may have a "surprise index" on the order of 20 to 40% (*31*), and even technical specialists can exhibit overconfidence. Indeed, it may be more instructive to turn to the relatively rare examples of experts who have been found not to exhibit high degrees of overconfidence. I am aware of only two such groups: weather forecasters (*32*) and the group of experts who publish forecast prices for horse bets (*33*). Both sets of experts receive enough feedback to calibrate the accuracy of their estimates, and both are subject to considerable scrutiny from lay experts if they fail to recognize calibration errors themselves.

Statistical vulnerability of low-probability estimates. In one important area, moreover, those of us with training in probability theory may be subject to a potential bias that is rarely found among the general public. Nearly all of us have had the frustration of attempting to explain to the lay public, or even to students, that events with one-in-a-million probabilities are not impossible; a low estimated probability is not necessarily called into question if an unlikely event does in fact occur. Of all the events that are expected to occur only once every thousand years or so, some can be expected to occur each year, and a tiny proportion may even occur more than once per year. Yet our understanding of these principles may sometimes cause our hypotheses about extremely low probabilities to become effectively nonfalsifiable.

The familiar statistical problem of type I and type II errors—of rejecting hypotheses that are ultimately found to be true, on the one hand, or failing to reject those that are actually false, on the other—takes on new complexity in cases of incidents that are expected to occur once in a million reactor years, for example, but that actually occur twice in a single year (*34*). If empirical operating experience is limited, there is little scientific basis for deciding whether probability estimates are too low or too high. If we stick with our estimates, we avoid discarding them on the basis of what may prove to be isolated experiences, but in doing so, we make a de facto decision to trust reasoning that may be incorrect. Although many areas of risk assessment provide enough experience to correct such errors, events that are truly rare, or technologies that are still new or untried, may provide too little information to permit the needed corrections.

The problem is exacerbated by the statistical power of the hidden flaw. Low-probability estimates are especially vulnerable to the inaccuracies created when calculations fail to foresee even a small number of problems. Contrary to "common sense" expectations, the failure to recognize a problem in one portion of a probabilistic analysis is often not offset by an exaggerated conservatism in another portion of the analysis.

Consider a technology estimated to have a one-in-a-million chance of failing. For simplicity's sake, assume that risk assessors had succeeded in identifying all potential risk factors but two—one of which made the technology safer than the official estimate, and the other of which made it less safe. Imagine that the technology would still have a one-in-a-million level of risk during 80% of its operational life, but that 10% of the time, the real risk would be 1 in 1000, and 10% of the time the risk would be one in a billion. Then the true risk of the technology would be $(0.1 \times 10^{-9} + 0.8 \times 10^{-6} + 0.1 \times 10^{-3})$, that is, 10% times 10^{-9} (one in a billion), plus 80% times one in a million, plus 10% times 1 in 1,000, respectively—for an overall probability of 0.0001008001, or slightly more than 1 in 10,000. Rather than being offset by the presence of the unexpected safety factor, the unexpected problem dominates the ultimate probability. Indeed, even if the risk assessment were to have been so conservative in other respects that the "real" risks were to be no higher than one in a trillion except for the 10% of the operating experience where the 1 in 1,000 estimate would hold, the overall probability would still be higher than 1 in 10,000.

Monetary and political pressures. So far in this article I have limited the discussion to cases in which no deliberate bias or distortion has been introduced into the risk assessment process. The biases or failings that have been noted are those of good-faith scientists and practitioners who are not subject to political, economic, or other pressures that might create biases of their own. The empirical world, however, is not always so tidy; scientists are sometimes subjected to distinctly unscientific pressures, and scientific results are not immune to being used by persons who may not fully share a scientific commitment to the fair and balanced reporting of evidence (*35*).

One obvious source of pressure would be the need to control costs, particularly in a competitive environment. Heimer (*36*) examined the willingness of technically trained industrial workers to take risks, even though most people tend to be quite reluctant to take risks for the sake of monetary gain (*37*). Drawing from analyses of accidents involving North Sea oil exploration (*38*), Heimer suggests the workers take risks "*to avoid costs* rather than to make gains. . . . When they are deciding whether to take a risk or not, offshore oil workers typically are not considering whether they will get some bonus for taking the risk, but instead whether they will be fired if they refuse" (*36*, p. 503).

Incentives to minimize or understate risks can also be created by political motivations, such as the desire to avoid public embarrassment, and by high levels of commitment to organization goals. It is easy to see the potential for abuse in totalitarian countries, or to be skeptical of the claim about Chernobyl and nearby reactors in the February 1986 issue of *Soviet Life* that "the odds of a meltdown are one in 10,000 years" (*39*, p. 4A). Yet our own system may also be vulnerable to such pressures—a point that can be illustrated by the Challenger incident.

As made clear in part through independent assessments and news reports (*40, 41*), NASA's official estimate of space shuttle risks, 1 in 100,000, is spectacularly at variance from the empirical record. Historical rates of failure have been on the order of 1 in 25 to 1 in 50, depending on the calculations used. Yet NASA "pressured one consultant to produce a more optimistic estimate of booster safety and disregarded even more pessimistic predictions contained in two subsequent studies (*41*, p. 1). A 1981 report calculated the historic failure rate of solid-fuel rockets to be 1 in every 57 firings, but concluded shuttle booster risks would be in the range of 1 in 1,000 to 1 in 10,000; a second study in 1983 concluded that the chance of a booster blowup was 1 in 70, and a third, "commissioned by the Air Force in 1984 to resolve the discrepancies between the first two, suggested the booster failure rate would be about 1 in 210" (*41*, p. 1). Finally, officials of NASA's Marshall Space Flight Center published their own risk estimate in February 1985, predicting a failure rate of 1 in 100,000, a rate of safety approximately 2,000 times better than actual experience (*42*). With the Challenger accident occurring on the 25th shuttle launch, and with two boosters per launch, the failure rate ultimately proved to be almost identical to the historical track record.

Implications: Reassessing "Rationality"

Despite flaws in risk assessments, many scientists worry about the problems that could be created by allowing a greater public role in technological decisions (*43*). Scientific workshops, for example, focus on topics such as "risk assessment and the 'misinformed' public." To some extent, of course, excluding the public from decisions is a luxury that a democracy does not offer. Some observers, moreover, have argued we would not want to indulge in this luxury even if it were available; the role of the scientist is to serve society, not to run it (*44, 45*). But how can these beliefs be reconciled with the supposedly "well-known fact" that members of the general public "are poor decision makers" (*46*, p. 2)?

I suggest three ways. The first is to keep in mind that "science often fluctuates between hubris and humility" (*47*, p. 24). It is important to remember in the policy realm what we often tell students in the classroom—that science is not infallible and that scientists do err from time to time. The need to avoid hubris, or even simple overconfidence, may be particularly high when our probability estimates are particularly low.

The second suggestion is to realize that the social science literature on decision-making in the general public is less clear-cut than is sometimes assumed. Berkeley and Humphreys (*48*) found that the *Science* article on decision-making by Tversky and Kahneman (*49*) had been cited 227 times between 1975 and 1980, with roughly one-fifth of the citations in sources outside the field of psychology—100% of which used the citation to support the unqualified claim that people are poor decision-makers. Yet the original article was somewhat less sweeping in its claims, and the broader literature provides a more mixed picture. Another survey found 84 articles in which the key words decision-making, judgment, and problem-solving were used (and a comparison to an explicit normative model was provided); 47 reported poor performance in decision-making, and 37 reported good performance (*50*). Other fields of the social and behavioral sciences tend to report that people's decision-making processes are more rational than they may at first appear (*51*). In short, just as scientists' estimates may need to be treated with something less than reverence, the views of the public may need to be treated with something better than contempt.

The third suggestion is to reconsider the notion of "irrationality" even in cases when the public does fail to understand the scientific details of a technological decision. This suggestion requires that we go beyond the notion of the public as irrational, but also beyond the notion of the public as economically rational—making selfish risk-benefit comparisons on the basis of whatever information and values may apply—and consider instead the possibility of prudence.

The public as prudent. It is instructive to turn from the large number of people who make up the general public to the small number who make up boards of trustees and boards of directors—the people who are expected to exercise prudence in directing society's largest and most influential organizations. In general, trustees and directors need to keep an eye on the big picture, not the little details; their job is to establish policy, not to pursue the particulars. A university's board of trustees is expected to look after the welfare of the overall organization, not the dealings between deans and departmental chairs, and certainly not the daily performance of tasks in research laboratories.

Given that the specialist is expected to look after specific details, there is a potential for friction between the trustee and the technician, even if each is doing his or her job in a way that is basically appropriate; at a minimum, the job of one is to ask questions that are of little concern to the other. At what point, however, would a prudent trustee decide not to defer to specialists who have been hired to look after the technical details? I asked this question to a small group of high-level policy-makers; they identified two sets of factors (*52*). The first has to do with characteristics of the technical specialists or experts, and the second with characteristics of the broader situation.

There are four problematic characteristics of specialists, the first of which occurs when an expert might have a vested interest in outcomes; for example, the suspicions of trustees tend to be raised if one department claims to need more money than all of the other departments combined. Concerns are also raised if the expert's recommendations or activities have implications for other parts of the organization, or if the expert fails to recognize constraints imposed by the larger picture. A third reason for concern is created if a given expert has been wide of the mark or has caused problems and difficulties in the past. The fourth arises when another expert warns the policy-maker that something seems to be seriously amiss. In all four cases, the gut-level credibility of the expert in question is likely to play a key role in guiding the policy-maker's response.

The second set of factors includes three types of situations that call for particular scrutiny. First are situations that incorporate a large element of the unknown: Activities that are familiar or draw on a large body of experience are not nearly as worrisome as those that get into areas in which the organization has less of an experience base, or in which the ratio of knowledge to guesswork is lower. Second are situations when the potential consequences of a mishap, specifically including an "unexpected" mishap, might be especially severe. It is unwise to put all of the organization's eggs into any one basket, for example, even if it is a good-looking basket. Third are situations when potential problems, if experienced, would be difficult or impossible to correct.

A more exhaustive search could develop a longer list, but the seven interrelated points in Table 1 should be sufficient. To return to technological controversies, everything on this list could be applied

Table 1. Typical "warning signs" that would cause a prudent board of trustees to question the recommendations of their technical experts.

Characteristics of specialists
Specialists have direct interest in outcomes
Specialists' past recommendations were wrong
Specialists' activities and recommendations have broader implications
Other experts indicate there may be reason for worry
Characteristics of situations
Those that contain a large element of the unknown
Those in which potential consequences of mistakes could be especially severe
Those in which errors have potential to be irreversible

to examples such as nuclear power and to the group that in many ways is the ultimate board of trustees—the public. The public has little knowledge of technical details of the industry or the nuclear fuel cycle, just as trustees or directors tend not to have detailed familiarity with given technical processes. What the public does have, however, is a set of all seven of the kinds of warning signals that would draw the attention of a "real" board of directors (*53*).

Although the key proponents of nuclear power tend to be scientifically trained, they stand to benefit from its implementation, if only through research and employment opportunities. Early efforts to promote the industry may have been too enthusiastic in describing the benefits or downplaying its problems. Rather than producing electricity "too cheap to meter" (*54*), nuclear power plants completed in recent years have suffered pervasive problems with cost overruns (*55*), and attention to management and waste control at some facilities has been less vigorous than it might have been (*16, 17*). A widespread commitment to nuclear power has been criticized as having serious implications, ranging from disposing of nuclear wastes to the threat of increased "police state" characteristics (*56*), and difficulties in identifying sites for nuclear waste disposal have lent some credence to the argument (*57*). Some respected scientists argue that nuclear technologies are unsafe or insufficiently understood, and TMI and Chernobyl have given the criticisms an increased credibility.

A similar picture emerges when we consider the characteristics of the broader situation. Nuclear power is still widely seen as involving a large element of the unknown. Events of the recent past have sent signals not only that the technology is less well understood than people might hope, but that the consequences of accidents could be extremely severe. The notion of "irreversible" implications takes on additional meaning when the public hears that the official planning horizon for a high-level nuclear waste repository is 10,000 years—a period roughly double the age of written records.

Many of these points, of course, are associated with counterarguments. Cost overruns resulted from a variety of factors, some of which were largely outside the control of the industry. The early promotional literature was intended to get people to think about the technology in terms of peacetime potential rather than mushroom-shaped clouds, and downplaying of the problems or uncertainties might have seemed necessary at the time. Many of the earlier problems have been recognized and remedied. We all underestimate problems that might crop up in implementing an idea, whether in the form of writing a new paper or installing a new kitchen sink.

All these objections are legitimate, but to a certain extent, they all miss the point. The public is often castigated for "irrationality" in its reactions to nuclear or other controversial technologies, and most people do know as little about the technical details as we might expect directors to know about a given technology in an organization that is served by literally thousands of different technologies. Even so, the public may be playing something like the role that "ought" to be played—particularly in managing controversial technologies that require decisions about values as well as about technical details (*58*).

Conclusion

It is tempting to assume that risk management can be improved by settling scientific facts before worrying about any social implications (*59*), or to assume that scientists identify "real" risks, with additional public concerns being due to misinformation or irrationality. Such assumptions may cause few problems when the stakes are low, consensus is high, experience is vast, and decisions do not impose burdens on one group for the benefit of another. The assumptions are clearly problematic, however, for controversies that involve high stakes, low consensus, new technologies, and unequal distributions of burdens and benefits. These kinds of technological controversies are often precisely those for which the perceived-versus-real argument is pushed with the greatest passion.

In studies that examine the ways in which citizens and scientists assess risks, investigators have found that citizens often reach ill-advised conclusions—and that scientists do as well. Although citizen judgments often incorporate misinformation, they can also reflect a deeper kind of prudence than is commonly realized.

Scientists' errors appear to be most problematic in two areas—those involving human and social factors, rather than physical and biological ones, and those requiring guesswork or judgment in the face of limited or nonexistent evidence. Monetary or political pressures can create additional problems and distortions.

There are no easy solutions to the problems of political, monetary, or personal biases, although such problems cannot be safely ignored. Nor are there easy solutions to the nonavailability of necessary data—although experts' resultant guesswork can be significantly worse than is often realized. Data gaps on human and social influences, however, are often correctable, but only if the needed research is performed.

In short, the social sciences should be asked to provide not just an improved understanding of public perceptions, but also significantly improved quantitative estimates of the probabilities as well as the consequences of important risks. This will require, however, that all of us explicitly begin to integrate human behaviors into our thinking about "technological" systems—and that we begin devoting approximately the same level of resources to understanding the human components of technological systems as to the hardware. Although often overlooked, human and social factors play vital roles in technological systems; real-world risks, far from being free of such inconvenient "people factors," are indeed often dominated by them.

REFERENCES AND NOTES

1. L. Harris *et al.*, *Risk in a Complex Society* (Marsh & McLennon, New York, 1980); V. T. Covello, J. Menkes, J. Mumpower, *Risk Evaluation and Management* (Plenum, New York, 1986); B. Fischhoff, P. Slovic, S. Lichtenstein, in *Analysis of Actual Versus Perceived Risks*, V. T. Covello, G. Flamm, J. Rodericks, R. Tardiff, Eds. (Plenum, New York, 1983), pp. 235–240.
2. W. R. Freudenburg and E. A. Rosa, Eds., *Public Reactions to Nuclear Power: Are There Critical Masses?* (Westview/AAAS, Boulder, CO, 1984).
3. J. S. Petterson, in *Waste Management '88*, R. G. Post and M. E. Wacks, Eds. (Univ. of Arizona Press, Tucson, in press).
4. This "was not simply a case of 'ignorant peasants' flopping around in confusion—or of pervasive cultural or information-base differences. For example, doctors and dentists trained in the U.S. routinely refused to treat patients without certificates . . . politically and economically well-placed individuals sought preferential treatment [certificates of noncontamination] to protect against stigma" (*3*).
5. If the possibility of a negative event is accompanied by uncertainty or "ambiguity" about the actual probability of the event, this creates further costs. A fuller discussion of insurance actuaries' pricing of risk premiums under situations of ambiguity is given by R. M. Hogarth and H. Kunreuther (*J. Risk Uncertainty*, in press.); H. J. Einhorn and R. M. Hogarth [*J. Bus.* **59**, S225, 1986).
6. J. F. Short, *Am. Soc. Rev.* **49**, 711 (1984); A. Baum, in *Cataclysms, Crises and Catastrophes: Psychology in Action*, G. R. VandenBos and B. K. Bryant, Eds.

(American Psychological Association, Washington, DC, 1987), pp. 5–53.
7. P. Slovic, *Science* **236**, 280 (1987).
8. R. E. Kasperson, O. Renn, P. Slovic, *Risk Anal.* **8**, 177 (1988).
9. B. P. Dohrenwend *et al.*, in *The Three-Mile Island Nuclear Accident: Lessons and Implications*, T. H. Moss and D. L. Sills, Eds. (New York Academy of Science, New York, 1981), pp. 159–74; D. L. Sills, C. P. Wolf, V. B. Shelanski, Eds., *Accident at Three-Mile Island: The Human Dimensions* (Westview, Boulder, CO, 1982).
10. W. R. Freudenburg and R. K. Baxter, *Policy Stud. Rev.* **5**, 96 (1985); *Soc. Sci. Q.* **65**, 1129 (1984); E. A. Rosa and W. R. Freudenburg, in (*2*), pp. 3–37.
11. President's Commission on the Accident at Three Mile Island, *The Need for Change: The Legacy of TMI* (Government Printing Office, Washington, DC, 1979), p. 8.
12. C. B. Flynn, in (*2*), pp. 205–232; J. R. Egan, *Technol. Rev.* **85**, 23 (1982).
13. Office of Technology Assessment, *Transportation of Hazardous Materials* (OTA-SET-304, Government Printing Office, Washington, DC, 1986); B. E. Sabey and H. Taylor, in *Societal Risk Assessment: How Safe Is Safe Enough?* R. C. Schwing and W. A. Albers, Eds. (Plenum, New York, 1980), pp. 43–65.
14. S. G. Hadden, *Risk Analysis, Institutions, and Public Policy* (Associated Faculty Press, Port Washington, NY, 1984); A. W. Lerner, *Admin. Soc.* **18**, 334 (1986). One paradoxical example is that the longer a technology operates without accidents, the harder it may be to sustain alertness and to continue operating safely. The problem may be worsened if redesigns cause operators to "depend on the equipment" instead of feeling sometimes-stressful levels of personal responsibility; G. I. Rochlin, T. R. LaPorte, K. H. Roberts, *Nav. War Coll. Rev.* **40**, 77 (1987). See also (*36*).
15. Nuclear Regulatory Commission Special Inquiry Group, *Three-Mile Island: A Report to the Commissioners and to the Public* (Government Printing Office, Washington, DC, 1980); Nuclear Regulatory Commission, *TMI-2 Lessons Learned Task Force Final Report* (NRC, Washington, DC, 1979).
16. M. L. Wald, "The Peach Bottom syndrome," *New York Times*, 27 March 1988, section 3, p. 1.
17. B. Dumaine, *Fortune* **114**, 40 (27 October 1986).
18. W. Rogers *et al.*, *Report to the President by Presidential Commission on the Space Shuttle Challenger Accident* (NASA, Washington, DC, 9 June 1987).
19. Larger organizations may be more prone to experience such departures from official policies, along with problems of top-level managers not being fully aware of problems, if only because of the number of communications links involved. In a hypothetical organization with an average correlation of 0.7 between what a given employee knows about a problem and what that person's supervisor knows, just two layers of bureaucracy would yield a correlation of less than 0.5 ($0.7 \times 0.7 = 0.49$), and seven layers would reduce the correlation to 0.082 (0.7^7). Overcoming this "organizational attenuation of information" would appear to require a concerted effort to highlight rather than to hide bad news. However, this is just the opposite of what often happens, particularly if an organization sees itself as facing a risk of failing to meet an important deadline, losing a key contract, or seeming not to be in control of a problem.
20. See the summary provided by B. B. Johnson and V. T. Covello [B. B. Johnson and V. T. Covello, Eds., *The Social and Cultural Construction of Risk: Essays on Risk Selection and Perception* (Reidel, Dordrecht, 1987)].
21. C. Perrow, *Normal Accidents: Living with High-Risk Technologies* (Basic, New York, 1984).
22. P. Slovic, B. Fischhoff, S. Lichtenstein, in (*2*), pp. 115–135.
23. J. P. Holdren, *Bull. At. Sci.* **32**, 20 (1976); J. P. Holdren, G. Morris, I. Mintzer, *Annu. Rev. Energy* **5**, 241 (1980).
24. B. Fischhoff, S. Lichtenstein, P. Slovic, S. L. Derby, R. L. Keeney, *Acceptable Risk* (Cambridge Univ. Press, New York, 1981), pp. 33–46.
25. M. Henrion and B. Fischhoff, *Am. J. Phys.* **54**, 791 (1986). In general, measurements of the physical constants are from experiments done in a single laboratory, providing a less satisfactory error estimate than from experiments involving different laboratories. The article notes highly similar problems of overconfidence in estimated errors for values such as h (Planck's constant), e (electron charge), m_e (electron mass), and N (Avogadro's number).
26. M. Hynes and E. VanMarcke, in *Mechanics in Engineering*, R. N. Darbey and N. C. Lind, Eds. (Univ. of Waterloo Press, Waterloo, Ontario, 1976), pp. 367–384.
27. J. J. J. Christensen-Szalanski and J. B. Bushyhead, *J. Exper. Psych.* **7**, 928 (1982). For related studies, see also A. A. De Smet, D. G. Fryback, J. R. Thornbury, *Am. J. Radiol.* **43**, 139 (1979); S. Lichtenstein, B. Fischhoff, L. D. Phillips, in *Judgment Under Uncertainty: Heuristics and Biases*, D. Kahneman, P. Slovic, A. Tversky, Eds. (Cambridge Univ. Press, New York, 1982), pp. 306–334.
28. S. Lichtenstein and B. Fischhoff, *Org. Behav. Human Perform.* **20**, 159 (1977).
29. J. E. Sieber, *J. Pers. Soc. Psychol.* **30**, 688 (1974).
30. Part of the problem is that these are often not statistical confidence intervals but statements of the expert's degree of belief in an estimate. "One's degree of belief in an outcome is based only on information selected as relevant from that which is available, while there is no way of knowing if even all the available knowledge is sufficient, let alone complete" [S. Rayner, in *Risk Analysis, Institutions, and Public Policy*, S. G. Hadden, Ed. (Associated Faculty Press, Port Washington, NY, 1984), p. 152]. Otway puts it more boldly: "Probability is no more than a degree of belief in an event or proposition whose truth has not been ascertained" [H. Otway, *Risk Anal.* **5**, 271 (1985)].
31. For example, B. Fischhoff, P. Slovic, S. Lichtenstein, *J. Exper. Psychol.* **3**, 552 (1977); see also (*27*).
32. A. H. Murphy and R. L. Winkler, *Nat. Weather Dig.* **2**, 2 (1977).
33. J. Dowie, *Economica* **43**, 139 (1976).
34. For example, E. Marshall, *Science* **220**, 280 (1983); see also (*27*).
35. A Schnaiberg, *The Environment: From Surplus to Scarcity* (Oxford Univ. Press, New York, 1980).
36. C. A. Heimer, *Annu. Rev. Sociol.* **14**, 491 (1988).
37. D. Kahneman and A. Tversky, *Sci. Am.* **246**, 160 (January 1982).
38. W. G. Carson, *The Other Price of Britain's Oil: Safety and Control in the North Sea* (Rutgers Univ. Press, New Brunswick, NJ, 1982).
39. Translation, in "Soviets praised industry's safety," *USA Today*, 29 April 1986, p. 4A. The article continued with the claim that "the plants have safe and reliable controls that are protected from any breakdowns."
40. E. Marshall, *Science* **232**, 1596 (1986).
41. M. Thompson, "NASA defied 5 years of warnings," *Wisconsin State Journal*, 1 June 1986, p. 1.
42. Marshall Space Flight Center, *Space Shuttle Data for Planetary Mission RTG Safety Analysis* (NASA, Marshall Space Flight Center, AL, 15 February 1985).
43. B. Cohen, *Risk Anal.* **5**, 1 (1985).
44. For example, O. Renn, in *Regulating Industrial Risks: Science, Hazards, and Public Protection*, H. Otway and M. Peltu, Eds. (Butterworths, London, 1985), pp. 111–127; A. M. Freeman III, in *People, Penguins, and Plastic Trees: Basic Issues in Environmental Ethics*, D. Van De Veer and C. Pierce, Eds. (Wadsworth, New York, 1986), pp. 218–227.
45. Particularly significant problems are introduced by the fact that many decisions "about technology" actually require value-based along with fact-based decisions, especially given that the values of scientists are often quite different from those of the general public. For illustrations, see S. M. Nealey and J. A. Hebert, in *Too Hot to Handle: Social and Policy Issues in the Management of Radioactive Wastes*, C. E. Walker, L. C. Gould, E. J. Woodhouse, Eds. (Yale Univ. Press, New Haven, 1983), pp. 94–111; R. E. Dunlap and M. E. Olsen, *Policy Stud. J.* **13**, 413 (1984); also (*20*) and (*24*).
46. L. L. Lopes, paper presented at *Symposium in Mass Communication* (Madison, WI, 19 November 1987).
47. R. Kates, *Nat. Res. Counc. News Rep.* **35**, 24 (7 July 1985).
48. D. Berkeley and P. Humphreys, *Acta Psychol.* **50**, 201 (1982).
49. A. Tversky and D. Kahneman, *Science* **185**, 1124 (1974).
50. J. J. J. Christensen-Szalanski and L. R. Beach, *Am. Psychol.* **39**, 75 (1984).
51. H. J. Einhorn and R. M. Hogarth [*Annu. Rev. Psychol.* **32**, 53 (1981)] note that the superiority of formal decision-making models over people's everyday rules of thumb may be less impressive in the "messiness" of the real world than in the artificial neatness of psychological laboratories. There is a clear possibility that "evolution is nature's way of doing cost/benefit analysis" (*ibid.*, p. 54), and that everyday rules of thumb may have survived the tests of natural selection, although this is a possibility rather than a proven fact. The authors note the possibility that reports of "flawed" decision-making may be subject to half-empty–half-full interpretations on the part of experimenters, particularly given that research on "lower animals" often reports behaviors that are impressively "consistent with optimizing principles. . . . The danger of such pictures is that they are often painted to be interesting rather than complete" (*ibid.*, p. 55). See also (*45*).
52. Discussion drawn from W. R. Freudenburg [in *Waste Management '87*, R. G. Post, Ed. (Univ. of Arizona Press, Tucson, AZ, 1987), pp. 109–115.
53. In this context, it may be worth reexamining the frustration that many feel when members of the public ignore the technical considerations that scientists find important and focus instead on what seem to be impossibly broad or global questions. The "trustees" metaphor would suggest that these may be precisely the types of "overall policy" questions that should be answered by someone other than the scientist, except when the scientist is acting as a citizen-trustee.
54. L. L. Straus, remarks to National Association of Science Writers, New York City, 16 September 1954. Reprinted in *Background Info* (Atomic Industrial Forum, Washington, DC, 1987) and in D. Ford, *The Cult of the Atom: The Secret Papers of the Atomic Energy Commission* (Simon and Schuster, New York, 1982), p. 50.
55. Department of Energy, *1983 Survey of Nuclear Power Plant Construction Costs* [DOE/EIA-0439(83), Energy Information Administration, Washington, DC, 1983].
56. The classic statement is by A. B. Lovins, *Foreign Aff.* **55**, 65 (October 1976); see also the compilation by S. H. Murdock *et al.* [*Nuclear Waste: Socioeconomic Dimensions of Long-Term Storage* (Westview, Boulder, CO, 1983)]; C. E. Walker, L. C. Gould, E. J. Woodhouse, Eds., *Too Hot to Handle: Social and Policy Issues in the Management of Radioactive Wastes* (Yale University Press, New Haven, 1983).
57. After extended debate, Congress attempted to resolve the issue with the Nuclear Waste Policy Act in 1982, but lawmakers amended the act within 5 years, dropping a carefully developed process in favor of sticking "a pin in the map" at a potential site in Nevada [E. Marshall, *Science* **239**, 15 (1988)]. Like earlier "solutions," this one may or may not last. D. S. Zinberg, in (*2*), pp. 233–253.
58. I refer to "the" public in the interest of simplicity, but there are of course many publics, not all of which are equally likely to be castigated for irrationality. Thus it may be instructive that the segment of the public made up by the insurance industry, with a long history of being held accountable for pragmatic risk assessments, has become reluctant to provide liability insurance for nuclear installations and environmental hazards, in other words, for some of the same types of installations that inspire aversive reactions among the publics at large. Public opinion polls do not suggest that the public is willing to shut down properly operating reactors at this time, but there is mounting evidence that new nuclear facilities, ranging from nuclear waste repositories to recently completed reactors, are increasingly seen as unacceptable. B. C. Farhar-Pilgrim and W. R. Freudenburg, in (*2*), pp. 183–203; see also (*10*).
59. National Academy of Science, *Risk Assessment in the Federal Government: Managing the Process* (National Academy Press, Washington, DC, 1983).
60. I thank D. Anderson, K. T. Erikson, A. M. Freeman III, R. E. Kasperson, H. Kunreuther, E. Nichols, R. A. Rappaport, P. Slovic, R. Stevenson, and several additional colleagues in the informal Risk Interest Group at the University of Wisconsin–Madison, for comments on this article.