# How Safe is Safe Enough?

## Determinants of Perceived and Acceptable Risk

Paul Slovic and Baruch Fischhoff

Decision Research
A Branch of Perceptronics
Eugene, Oregon

Revised

January, 1979

This is a preliminary draft of a chapter in the forthcoming book, Too Hot To Handle: Social and Policy Issues in the Management of Radioactive Wastes, edited by L. Gould and C. A. Walker, Yale University Press. It is subject to change.

# How Safe is Safe Enough?

## Determinants of Perceived and Acceptable Risk

Paul Slovic and Baruch Fischhoff

Decision Research
A Branch of Perceptronics
Eugene, Oregon

Citizens of modern industrial societies are learning a harsh and discomforting lesson: the benefits from technology must be paid for not only with money, but with environmental degradation, anxiety, illness, injury, and premature loss of life. Through the news media, the American public has experienced a relentless parade of new and exotic hazards. As Rabinowitch (1972) observed:

> One day we hear about the danger of mercury, and run to throw out cans of tuna fish from our shelves; the next day the food to shun may be butter, which our grandparents considered the acme of wholesomeness; then we have to scrub the lead paint from our walls. Today, the danger lurks in the phosphates in our favorite detergent; tomorrow the finger points to insecticides, which were hailed a few years ago as saviors of millions from hunger and disease. The threats of death, insanity and--somehow even more fearsome--cancer lurk in all we eat or touch (p. 5).

The daily discovery of new threats and their widespread publicity threatens to create a national neurosis characterized by mistrust of technology and obsessive preoccupation with risk. It is in this context that the debate over nuclear waste is set.

One reaction to these perceived hazards is expressed by writer

Elizabeth Gray (1976):

> I have dealt with the situation basically by trying to drop out. I go to my doctor as seldom as possible; I take as few medications as possible. I . . . go into my supermarket, buy non-additive bread, take my vegetables and fresh fruits home and scour them to attempt to get the pesticides off.
>
> Now when you come at me with a nuclear decision, I think a person like me is . . . going to be propelled into the picket lines when you want to invade the very precarious private space that I am trying to forge . . . to protect myself from . . . [the] very dangerous effects of my society.
>
> And when you begin to invade my space with nuclear plants, people like me will say, 'No way. Someplace else . . . . ' (p. 200).

The option of dropping out gives us some control over the level of technological risk to which we are exposed. However, reduction of risk typically entails reduction of benefit. Where individual control is possible, each of us must cope with such dilemmas by personally weighing the costs against the benefits. Other dilemmas, such as the management of nuclear wastes, can only be resolved by society as a whole.

The urgent need to help society cope with risks has produced a new intellectual discipline, "risk assessment" (Kates, 1978; Lowrance, 1976; Otway & Pahner, 1976; Rowe, 1977). Risk assessment aims to determine how serious a hazard is and whether society should be exposed to it. Doing this requires an extraordinary degree of cooperation between technology sponsors, the public, its representatives and specialists from many fields. Technical issues require the efforts of physicists, biologists, chemists, and engineers. Social issues involve lawyers, political scientists, geographers, sociologists, economists and psychologists. Specialists in decision making attempt to coordinate this diverse expertise. They ask, in effect: "Given our society's values and all this knowledge, what actions should be recommended?"

However valid such assessments might be, the decisions eventually taken by society will reflect social and political pressures as well as the calm, analytic weighing of costs and benefits. Before acting, participants in those decisions must engage in an intellectual process for which the risk assessment may be an important input. They must judge for themselves the possible consequences of a technology, the likelihood that these consequences will occur, their importance and the combined implications of these various considerations.

Despite an appearance of objectivity, risk assessment is inherently subjective. Rarely will relevant statistical data (e.g., historical failure rates) be available. When it is, interpretation of such data is still subjective (e.g., Is the situation the same now as it was in the past?). More often, especially with new technologies, the risks must be estimated by applying engineering judgment to blueprints or to data on related systems and test trials. For the lay person, lacking specialized training, and access to data, the decision process will be even more subjective.

This chapter explores some of the intellectual elements in risk assessment that are critical to the nuclear debate. Its basic premises are that both the public and the experts are necessary participants in that debate, that there is a subjective element in all judgments, and that understanding the limitations of judgmental processes and proposed decision-making techniques is crucial to effective hazard management.

Coping Intellectually with Risk

Decisions about nuclear energy require high-level thinking and reasoning on the part of experts and non-experts alike. They require an appreciation of the probabilistic nature of the world and the ability to think intelligently about low-probability (but high-consequence) events. As Weinberg (1976) noted, ". . . we certainly accept on faith that our human intellect is capable of dealing with this new source of energy" (p. 21). Recently, however, the faith of many of us who study human decision processes has been shaken.

Consider probabilistic reasoning. Because of its importance to decision making, a great deal of recent research has been devoted to understanding how people deal with the probabilities of uncertain events. This research has found that intelligent people systematically violate the principles of rational decision making when judging probabilities, making predictions and otherwise attempting to cope with uncertainty. For example, people fail to recognize randomness when they encounter it; instead they perceive systematic patterns and lawful relationships in situations where none exist. In some situations, they overvalue small samples of data and unreliable data. Yet, when attempting to make forecasts or predictions, they desire too much information, often of the wrong type. Frequently, these difficulties can be traced to the use of judgmental "heuristics," mental strategies (or rules of thumb) that allow people to reduce difficult tasks to simpler judgments (Tversky & Kahneman, 1974). These heuristics are useful guides in some circumstances, but in others they lead to large, persistent biases with serious implications for decision making.[1]

Availability Bias

One heuristic particularly relevant for judgments about risks from nuclear power is "availability." This heuristic involves viewing an event as likely or frequent if it is easy to imagine or recall instances of it. Generally, instances of frequent events are more easily recalled than instances of infrequent events, and likely occurrences are easier to imagine than unlikely ones. Thus, availability is often an appropriate cue for judging frequency and probability. However, availability is also affected by numerous factors unrelated to likelihood. As a result, reliance on it may lead people to exaggerate the probabilities of events that are particularly recent, vivid, or emotionally salient.

Availability helps explain distortions in our perceptions of risk. Consider fears about grizzly bear attacks in our national parks. Although many people are concerned about the dangerousness of grizzlies, the rate of injury is only 1 per 2 million visitors and the rate of death is very much lower (Herrero, 1970). Sensational media reports contribute to the imaginability of death at the claws of an enraged grizzly, but the media ignore the multitude of uneventful visits. The motion picture, "Jaws," has likewise increased the availability (and perceived likelihood) of shark attacks. Some nuclear power proponents feel that the risks of that technology are exaggerated in the public's eye because of excessive media coverage and association with the vivid, imaginable, memorable dangers of nuclear war. As Zebroski (1976) notes, "fear sells;" the media dwell on potential catastrophes, not on the successful day-to-day operations of power plants.

Availability bias is illustrated in a recent study in which college students and members of the League of Women Voters were asked to judge the annual frequencies of death from each of 41 causes, including diseases, accidents, homicide, suicide, and natural hazards (Lichtenstein, Slovic, Fischhoff, Layman and Combs, 1978). These frequency judgments were greatly in error for many of the causes. Table 1 lists the causes whose frequencies were most seriously misjudged. Consistent with availability considerations, overestimated items tended to be dramatic and sensational. Underestimated items tended to be unspectacular events, which claim one victim at a time and are common in nonfatal form.

Insert Table 1 about here

Overconfidence

A particularly pernicious aspect of heuristics is that people are typically very confident in the judgments based upon them. For example, Fischhoff, Slovic and Lichtenstein (1977) asked people to indicate the odds that they were correct in their judgments about which of two causes of death was more frequent. Odds of 100 : 1 or greater were given often (25% of the time). However, about one in eight answers associated with such extreme confidence was wrong (fewer than 1 in 100 should have been wrong if the odds had been appropriate). To take but one example, about 30% of the judges gave odds greater than 50 : 1 to the incorrect assertion that homicides are more frequent than suicides. The psychological basis for this unwarranted certainty seems to be people's insensitivity to the tenuousness of the assumptions upon which their judgments are based (in this case, the validity of the availability heuristic). The danger from such overconfidence is that we may not realize how little we know and how much additional information is needed about the various problems and risks we face.

Overconfidence manifests itself in other ways as well. For example, a typical task in estimating failure rates or other uncertain quantities is to set upper and lower bounds such that there is a 98% chance that the true value lies between them. Experiments with diverse groups of people making many different kinds of judgments have shown that, rather than 2% of true values falling outside the 98% confidence bounds, 20-50% do so (Lichtenstein, Fischhoff & Phillips, 1977). Thus people think that they can specify such quantities with much greater precision than is actually the case.

Unfortunately, experts seem as prone to overconfidence as lay people. Hynes and Vanmarcke (1976) asked seven "internationally known" geotechnical engineers to predict the height of an embankment that would cause a clay foundation to fail and to specify confidence bounds around this estimate that were wide enough to have a 50% chance of enclosing the true failure height. The bounds specified by these experts were too narrow. None of them enclosed the true failure height. The multi-million dollar Reactor Safety Study (U.S. Nuclear Regulatory Commission, 1975), in assessing the probability of a core melt in a nuclear reactor, used the very procedure for setting confidence bounds that has been found in experiments to produce the highest degree of overconfidence. The Committee on Government Operations (U.S. Government, 1976) has attributed the 1976 collapse of the Teton Dam to the unwarranted confidence of engineers who were absolutely certain they had solved the many serious problems that arose during construction. Indeed, in routine practice, failure probabilities are not even calculated for new dams even though about 1 in 300 fails when its reservoir is first filled.

Desire for Certainty

Every technology is a gamble of sorts and, like other gambles, its attractiveness depends on the probability and size of its possible gains and losses. Both scientific experiments (e.g., Lichtenstein & Slovic, 1973) and casual observations show that people have difficulty thinking about and resolving the risk/benefit conflicts even in simple gambles (e.g., which would you rather play: a gamble with favorable chances of winning a modest amount or a gamble with less favorable odds but a larger possible payoff?).

One way to reduce the anxiety generated by uncertainty is to deny the uncertainty. The denial resulting from this anxiety-reducing search for certainty constitutes another source of overconfidence, in addition to those described earlier. Denial is illustrated by many people exposed to natural hazards who view their world as either perfectly safe or predictable enough to preclude worry. Thus some flood victims interviewed by Kates (1962) flatly denied that floods could ever recur in their areas. Some thought (incorrectly) that new dams and reservoirs in the area could contain all potential floods, while others attributed previous floods to freak circumstances unlikely to recur. Denial, of course, has its limits. Many people feel that they cannot ignore the risks of nuclear power. For these people, the search for certainty is best satisfied by outlawing the risky technology.

Scientists and policy makers who point out the gambles involved in societal decisions are often resented for the anxiety they provoke. Borch (1968) noted how annoyed corporate managers get with consultants who give them the probabilities of possible future events instead of telling them exactly what will happen. Just before hearing a blue ribbon panel of scientists report being 95% certain that cyclamates do not cause cancer, Food and Drug Administration Commissioner Alexander Schmidt said, "I'm looking for a clean bill of health, not a wishy-washy, iffy answer on cyclamates" (Eugene Register Guard, January 14, 1976). Senator Muskie has called for "one-armed" scientists, who do not respond "on the one hand, the evidence is so, but on the other hand . . ." when asked about the health effects of pollutants (David, 1975).

The search for certainty is legitimate if it is done consciously, if residual uncertainties are acknowledged rather than ignored, and if people realize the costs. If extreme certainty is sought, those costs are likely to be high. Eliminating the uncertainty may mean eliminating the technology and foregoing its benefits. Often, some risk is inevitable. Efforts to eliminate it may only alter its form. We must choose, for example, between the vicissitudes of nature on an unprotected flood plain and the less probable, but potentially more catastrophic, hazards associated with dams and levees.

Perseverence of Beliefs

The difficulties of facing life as a gamble contribute to the polarization of opinion about nuclear power. Some people view it as extraordinarily safe, while others view it as a catastrophe in the making. It would be comforting to believe that these divergent beliefs would converge towards one "appropriate" view as new evidence was presented. Unfortunately, this is not likely to be the case. As noted earlier in our discussion of availability, risk perception is derived in part from fundamental ways of thinking that lead people to rely on fallible indicators such as memorability. Furthermore, a great deal of research indicates that people's beliefs change slowly, and are extraordinarily persistent in the face of contradictory evidence (Ross, 1977). Once formed, initial impressions tend to structure and distort the way in which subsequent evidence is interpreted. New evidence appears reliable and informative if it is consistent with one's initial belief; contradictory evidence is dismissed as unreliable, erroneous, or unrepresentative. Ross (1977) concluded his review of this phenomenon as follows:

> Erroneous impressions, theories, or data processing strategies, therefore, may not be changed through mere exposure to samples of new evidence. It is not contended, of course, that new evidence can never produce change--only that new evidence will produce less change than would be demanded by any logical or rational information-processing model. (p. 210).

The "Wisdom" of Hindsight

Technologists, policy makers and regulators must constantly consider how their decisions will be judged in hindsight by history, by the public, and by courts of law. Their decisions will be judged harshly if it appears that they failed to anticipate important, foreseeable difficulties. Psychological research indicates that, in hindsight (i.e., knowing how things actually turned out), people consistently exaggerate what could have been anticipated in foresight (Fischhoff, 1975a, b). They not only tend to view events that happened as having been inevitable, but also believe (often incorrectly) that those events appeared "relatively inevitable" before they happened and that "others should have known they were going to occur." They also misremember their own predictions, exaggerating in hindsight what they themselves knew in foresight (Fischhoff, 1977a; Fischhoff & Beyth, 1975).

An extreme response to this form of bias is to claim that one's critics are always guilty of capitalizing unfairly on hindsight knowledge. This would not be entirely justified, for although the bias is pervasive, there is still some relation between what was and what seems to have been anticipatable. Just evaluation requires a method for improving hindsight. One bit of advice we can give to decision makers is to leave a clear record of what they knew and the uncertainties surrounding their actions. Critics of these decision makers might improve the acuity of their hindsight by attempting to state how the event might have turned out otherwise (Slovic & Fischhoff, 1977) or by seeking the opinions of persons not already contaminated by outcome knowledge.

Forecasting Public Response Towards Nuclear Power

Given this research into how people generally view risks and uncertainties, can we predict how they will respond to nuclear power? Probably not very well. Depending on how nuclear risks are presented by the media, in public debates and in private discussions and on whether or not there are major accidents, near misses or energy shortages, nuclear power may come to be viewed as increasingly safe or increasingly dangerous.

Implications of the availability heuristic. It is easy to see how accidents, near misses, or even minor problems, coupled with the attention the news media give such events, would increase the perceived risk from nuclear power. But a more subtle and disturbing implication of the availability heuristic is that any discussion of low-probability hazards, regardless of its content, will increase the memorability and imaginability of those hazards and, hence, increase their perceived risks. This poses a major barrier to open, objective discussions of nuclear safety. Consider an engineer demonstrating the safety of waste disposal in a salt bed by pointing out the improbability of the various ways radioactivity could be released (see Figure 1). Rather than reassuring the audience, the presentation might lead them to feel that "I didn't realize there were that many things that could go wrong".

---

Insert Figure 1 about here

---

Availability magnifies fears of nuclear power by blurring the distinction between what is remotely possible and what is probable. As one nuclear proponent lamented, "Whan laymen discuss what might happen, they sometimes don't even bother to include the 'might'" (B. Cohen, 1974, p. 36). Another analyst has elaborated a similar theme in the misinterpretation of "worse case" scenarios:

> It often has made little difference how bizarre or improbable the assumption in such an analysis was, since one had only to show that some undesirable effect could occur at a probability level greater than zero. Opponents of a proposed operation could destroy it simply by exercising their imaginations to dream up a set of conditions which, although they might admittedly be extremely improbable, could lead to some undesirable results. With such attitudes prevalent, planning a given nuclear operation becomes somewhat perilous since it requires predicting the extent to which the adversaries can employ their imagination (J. Cohen, 1972, p. 55).

Qualitative factors. Opponents of nuclear power appear to be responding not just to the probability of a mishap, but also to a number of frightening qualitative perceptions of the harm it might produce. Some insight into these qualitative factors may be found in a study by Fischhoff, Slovic, Lichtenstein, Read and Combs (1978), in which 76 members of the League of Women Voters rated the risks associated with 30 activities and technologies on nine qualitative scales: voluntariness, familiarity, controllability, potential for catastrophe (multiple fatalities), immediacy of consequences, the degree to which the risks are known to the public and to scientists, the extent to which the risks are common (as opposed to dread), and lethality (the likelihood that a mishap would prove fatal). Participants also rated the total risk and benefits accruing to society from each activity or technology, as well as how acceptable those risks were.

One of the most interesting findings of that study was the evaluation of nuclear power. For one, the benefits of nuclear power were not appreciated, being judged lower than those of home appliances, bicycles, and general aviation. Perhaps this is because nuclear power is seen merely as a supplement to other, essentially adequate, sources of energy. Second, its risks were seen as extremely high. Only automobile accidents, which take about 50,000 lives each year, were viewed as comparably risky. Third, its current level of risk was judged as unacceptably high. Participants in this study wanted nuclear power to be far safer than they now perceive it to be. The frightening character of nuclear power emerged clearly in the rating scales. Figures 2a and 2b show its unique risk profile. Nuclear power was rated at or near the extreme on all of the characteristics associated with high risk: involuntariness, uncontrollability, dread, lethality, etc. These figures also contrast nuclear power with two ostensibly similar technologies, X rays and non-nuclear electric power. Although both X rays and nuclear power involve radiation, nuclear power was judged much more catastrophic and dread. The comparison in Figure 2b shows that, where risk is concerned, nuclear power is not seen as just another form of energy.

Insert Figures 2a and 2b about here

Nuclear power was, in fact, rated higher on "dread" than any of the 29 other items studied by Fischhoff et al. This may stem from the association of nuclear power with nuclear weapons and from fear of radiation's invisible and permanent bodily contamination that causes genetic damage and cancer (Lifton, 1976; Pahner, 1975). Subsequent studies using students, business people and professional risk experts as participants have revealed remarkable agreement in judgments of the qualitative characteristics of the risks encountered with these 30 technologies.

Pathways toward acceptance. With all this working against it, how could nuclear power ever gain acceptance? Public response to X rays and nerve gas provides some clues. Widespread acceptance of X rays suggests that a radiation technology can be tolerated once its use becomes familiar, its benefits clear, and its practitioners trusted. However, although nuclear power might someday attain the low "dread" level of X rays, its perceived potential for catastrophic accidents seems less likely to change. Whether a continued high score on that one characteristic would render it permanently unacceptable is unclear.

Nerve gas may provide an enlightening case study. Few human creations could be more dread or more potentially catastrophic than this deadly substance. When, in December of 1969, the army decided to transfer nerve gas from Okinawa to the Umatilla Army Depot in Hermiston, Oregon, citizens of Oregon were outraged--except those in Hermiston. Whereas public opinion around the state was more than 90% opposed, residents of hermiston were 95% in favor of the transfer, despite the warning that the fuses on the gas bombs deteriorate with age, but that the gas does not (Eugene Register Guard, December 18, 1969, and January 11, 1970). Several factors seem to have been crucial to Hermiston's acceptance of nerve gas. For one, munitions and toxic chemicals had been stored safely there since 1941 (so the record was good and presence of the hazard was familiar). Second, there were clear economic benefits to the community from continued storage at the depot of hazardous substances, in addition to the satisfaction of doing something patriotic for the country. Finally, the responsible agency, the U.S. Army, was respected and trusted.

These examples illustrate the slow path through which nuclear power might gain acceptance. It requires an incontrovertible long-term safety record, a responsible agency that is respected and trusted, and a clear appreciation of benefit. However, since people are generally willing to accept increased risks in exchange for increased benefits (Starr, 1969; Fischhoff et al., 1978), a quicker path to acceptance might be forged by a severe energy shortage. Brownouts, rationing, or worse would undoubtedly enhance the perceived benefits from nuclear power and increase society's tolerance of its risks. A recent example of this process is the oil crisis of 1973-4 which broke the resistance to offshore drilling, the Alaska pipeline and shale oil development, all of which had previously been delayed because of their environmental risks.

Resistance to change. One likely possibility is that people will maintain their present positions, pro or con. This would be consistent with the research on perseverance of beliefs showing that rather than modifying existing beliefs, new evidence tends to be distorted in a way that confirms them. In the context of nuclear power, Zebroski (1976) observed that opponents often interpret intense effort to reduce nuclear risks as evidence that the risks are great, not as a sign that the technologists are responsive to the public's concern. Likewise, these same individuals may view minor mishaps as near catastrophes and dismiss the opinions of experts disputing such claims on the grounds that they are biased by the experts' vested interests in the industry. From a statistical standpoint, convincing a skeptic that nuclear power is safe would be difficult under the best of conditions. Any mishap can be seen as proof of high risk, whereas demonstrating high reliability requires a massive amount of evidence (Green & Bourne, 1972). Nelkin's (1974) case history of a nuclear plant siting controversy provides a good example of the inability of technical arguments to change opinions. Each side in that debate interpreted technical ambiguities in ways that reinforced its own position.

## Aiding the Decision-Making Process

The intellectual limitations described above portend continued and severe conflict over the safety and desirability of nuclear energy. The prospect of poor decisions with extreme costs to society seems all too likely. Sinsheimer (1971) observed that the human brain has evolved to cope with real, immediate and concrete problems and thus lacks the proper framework with which to encompass other sorts of phenomena. People have only recently faced decisions such as those involving nuclear energy. Following Sinsheimer's reasoning, it might be argued that we have not had the opportunity to evolve the intellect needed to deal with uncertainties of this nature. We are essentially trial-and-error learners, in an age where errors are increasingly costly.

But there may be some positive steps we can take to minimize the consequences of these limitations. Besides being less confident in our intellect, we can attempt to develop procedures for combatting the biases to which risk assessments are susceptible. The simplest "procedure" is to be wary of bias and hope that alertness will suffice. Alternatively, we may employ several methods for making the same judgment in the hope that their respective biases will be detected or perhaps balance one another. A third possibility is to restructure judgment tasks, perhaps by decomposing them into simpler judgments that can be made with less bias. The use of fault trees is a widely used decomposition method for estimating risks that we shall discuss later in this section.

The same technological bent that has created so many new hazards has also created methods designed to help make decisions about these hazards. Cost-benefit analysis and decision analysis are leading members of this genre. Like the technologies they are meant to evalutate, these analytic techniques have both potential benefits and inherent limitations. They can improve the decision-making process and its sensitivity to public desires, but only if the public understands the techniques and their limitations, monitors the way that analyses are done, and makes certain that their conclusions are heeded (or ignored, as suitable). To this end, the present section also describes some of the more important approaches and techniques for making decisions, the problems encountered in applying them, and the ways in which specific techniques may be led astray and produce erroneous results.

Estimating Risks

A frequently used aid for assessing and communicating the risks of a complex system is a fault tree. Construction of a fault tree begins by listing all important pathways to failure, then listing all possible pathways to these pathways and so on, as shown in Figure 1. A fault tree representing the "risks" of an automobile failing to start appears in Figure 3. The first level lists major systems problems such as battery failure; the next level traces these global failures to sources like loose terminals or weak charge; the lowest level provides even more detail. When the desired degree of detail is obtained, the experts assign probabilities to each of the component pathways (relying on judgment or available data) and then combine these to provide an overall failure rate.[2] The importance of fault-tree analysis is demonstrated by its role as the primary methodological tool in the Reactor Safety Study, which assessed the probability of a catastrophic loss-of-coolant accident in a nuclear power reactor (U.S. Nuclear Regulatory Commission, 1975).

---

Insert Figure 3 about here

---

Errors of omission. Fault-tree analysis has been attacked by critics who question whether it is valid enough to be used as a basis for decisions of great consequence (e.g., Bryan, 1974; Fischhoff, 1977a; Primack, 1975). One major danger in designing a fault tree is leaving things out and, thereby, underestimating the true risk. The car-won't-start tree would be seriously deficient if it failed to include problems with the seat-belt system (for 1974 models) or vandalism. The cartoon in Figure 4, drawn after the discovery that hydrofluorocarbons from aerosol products may damage the earth's ozone shield, dramatizes the dangers of omitting relevant pathways to disaster. It is unlikely that any fault tree created before 1974 to evaluate the major threats from technology included hair sprays and deodorants.

---

Insert Figure 4 about here

---

Several kinds of pathways seem particularly prone to omission. One type involves human error or sabotage. Can we ever be certain that we have enumerated all of the important and imaginative ways in which we, the people (as opposed to they, the machines), can create trouble? Consider the Brown's Ferry fire in one of the world's largest nuclear power plants, which was caused by a technician checking for an air leak with a candle, in direct violation of standard operating procedures (Comey, 1975). The fire got out of control, in part, because plant personnel were slow to sound alarms and begin the reactor shut-down. Disaster was averted finally when plant personnel managed to jury-rig pumps normally used to drive control rods into the reactor to, instead, pump water to cool the reactor core. Identifying such possibilities for human error (or ingenuity) obviously poses a difficult challenge for risk analysts (U.S. Nuclear Regulatory Commission, 1978).

A second source of omissions is failure to consider unanticipated changes in the world in which the technology functions (Coates, 1976; Hall, 1975). Risk assessments always assume some constancies in the external environment. At times, these assumptions may be both unrecognized and questionable. For example, nuclear power plant design assumes implicitly the continued availability of properly trained personnel. Even though a tree's designers might not realize that they are making this assumption, it is possible to imagine a future world in which such individuals are in short supply.

Omissions may also result from failing to see how the system functions as a whole. For example, the rupture of a liquid natural gas storage tank in Cleveland in 1944 resulted in 128 deaths, largely because no one had realized the need for a dike to contain spillage (Katz & West, 1975). The DC-10 failed repeatedly in its initial flights because none of its designers realized that decompression of the cargo compartment would destroy vital parts of the plane's control system (Hohenemser, 1975). Green and Bourne (1972, p. 547) caution us not to forget that backup systems may not function when needed because they are undergoing routine maintenance and testing or because they have been damaged by the testing process.

Another omission in this category is provided by a National Academy of Sciences study of the effects of thermonuclear war. The Academy panel decided that the anticipated reduction of the earth's ozone shield would not imperil the survivors' food supply because many crops could survive the increased ultraviolet radiation. The study failed to point out, however, that increased radiation would make it virtually impossible to work in the fields to raise those crops. "How was this overlooked? Because . . . it fell between the chinks of the expert panels. The botanists who considered the effects of ultraviolet radiation on plants didn't think to worry about the workers." (Boffey, 1975, p. 250).

Common-mode failures. Another sort of error in the use of fault trees and one that the Reactor Safety Study took great pains to avoid is miscalculating what are called "common-mode failures." To insure greater safety, many technological systems are built with a great deal of redundancy. Should one crucial part fail, there are others designed either to do the same job or to limit the resulting damage. Since the probability of each individual part failing is very small, the probability of all failing, thereby creating a major disaster, would seem to be extremely small. This reasoning is valid only if the various components are independent, that is, if what causes one part to fail will not automatically cause the others to fail. "Common-mode failure" occurs when the independence assumption does not hold. As an example, the discovery that a set of pipes in several nuclear power plants were all made from the same batch of defective steel (Eugene Register Guard, October, 13, 1974), suggests that the simultaneous failure of several such pipes is not inconceivable. At Brown's Ferry, the same fire that caused the core to overheat also damaged the electrical system needed to shut the plant down. Constructing a tree that takes proper account of all such contingencies may be very difficult (U.S. Nuclear Regulatory Commission, 1978).

Presentation biases. Fault trees are tools for the communication as well as for the estimation of risks. The expert who has completed an analysis of risks, with whatever success, must present the results to other experts or to the public. Doing so involves making a number of discretionary decisions. For example, with the car starting fault tree (Figure 3), the presenter must decide how much detail to provide for each branch, which minor pathways to lump into the "all other problems" category and just how to categorize various items (e.g., should the items grouped under "fuel system defective" be split into two pathways, fuel line defective and carburetion system defective?).

We have recently studied how such "discretionary" aspects of fault-tree presentation affect people's perceptions of the risks they embody (Fischhoff, Slovic & Lichtenstein, 1978). Our results indicate that the decision to put some difficulties in an "all other problems" category can have a major biasing effect. People are quite insensitive to how much has been left out of a fault tree. Deleting branches responsible for about half of all automobile starting failures only produced a 7% increase in people's estimates of what was missing. Professional automobile mechanics were about as insensitive as non-experts. Apparently, what was out of sight was also out of mind. The fault-tree presenter who, deliberately or inadvertently, fails to mention a branch (thereby implicitly or explicitly assigning it to "all other problems") may remove it completely from consideration. We also found that the perceived importance of a set of problems can be substantially increased by presenting it as two (smaller) problem categories rather than as one category (e.g., splitting "fuel system" into "fuel-line problems" and "carburetion problems").

The fact that subtle differences in how risks are presented can have large effects on how they are perceived suggests that people attempting to communicate information about risks have considerable ability to manipulate others' perceptions without making any overt misrepresentations. Indeed, since these effects are not widely known, people may inadvertently be manipulating their own perceptions by decisions they make about how to organize their knowledge.

As with the other research we have described, studies of presentation biases have two lessons. One is to be wary, realizing that judgment is fallible and readily influenced by irrelevant factors. The second is that these effects might be counteracted by adopting a variety of perspectives. Ask yourself: What is left out? How would this problem look if the categories were rearranged? Is the presenter interested in manipulating my perceptions and, if so, what strategies might make that possible?

Reaching a Decision

Muddling through. After risks have been assessed, some decision must be made. By far the most common approach towards setting risk policies, nuclear or otherwise, is "muddling through," making somewhat arbitrary initial decisions and then letting them be molded into generally accepted standards by the pressure of political and economic forces. While this process may employ analytic arguments, it is essentially non-analytic. It relies upon the internal structure of participating organizations, their interaction with one another, and the varied feedback provided by their environment to produce satisfactory decisions. The building blocks of this approach are mechanisms that are familiar and accepted, even if not entirely understood. This approach does not attempt to produce and defend acceptable risk criteria on the basis of specific decision analytic techniques, although it does allow such analyses to be presented as evidence.

Comparative analyses. One major form of input into the process of muddling through comes from various forms of comparative analyses. Comparative procedures attempt to determine the acceptable level of risk for a given hazard (e.g., nuclear power) by reference to the level of safety tolerated from other hazards, either natural or technological. For example, the allowable radiation from a particular segment of the nuclear fuel cycle might be set equal to natural background radiation or equal to a fraction of that tolerated from medical exposures. Workers in the nuclear industry might be expected to tolerate the same level of risk borne by workers in other energy-producing industries.

Comparative analysis has several attractive features. It avoids the difficult and controversial task of converting diverse risks into a common monetary unit (like dollars per life lost or per case of sterilization or per day of suffering). It presents issues in a mode that is probably quite compatible with natural thought processes. Among other things, this mode may avoid any direct numerical reference to very small probabilities, for which people have little or no intuitive feel.

A more elaborate form of comparative analysis, incorporating benefits as well as risks, is the revealed preference approach advocated by Starr (1969). This approach is based on the assumption that, by trial and error, society has arrived at nearly optimal balance between the risks and benefits associated with any activity. If this is the case, then historical data can be used to reveal acceptable risk/benefit tradeoffs. Acceptable risk for a new technology is assumed to be the level of safety associated with ongoing activities having similar benefit to society.

From this approach, Starr (1969) derived what may be regarded as "laws of acceptable risk." These included: (1) greater risks are acceptable for more beneficial activities; and (2) the public accepts much greater risks from voluntary activities (e.g., skiing) than from involuntary activities (e.g., food preservatives) providing similar levels of benefit. Thus, according to Starr's analysis, acceptable risk is determined by two factors, benefit and voluntariness.

Although the method of revealed preference is based upon an intuitively compelling logic, it has several drawbacks. It assumes that past behavior is a valid predictor of present preferences, perhaps a dubious assumption in a world where values may change quite rapidly. It is politically conservative in that it enshrines current economic and social arrangements. It makes strong (and not always supported) assumptions about the rationality of people's decision making in the marketplace and about the freedom of choice that marketplace provides. It may underweigh risks to which the market responds sluggishly such as those with a long lead time (e.g., most carcinogens). Finally, it is no simple matter to develop the measures of risks and benefits needed for its implementation (Otway & Cohen, 1975).

Cost-benefit analysis. Cost-benefit analysis attempts to quantify, in terms of dollars, the expected gains and losses from a proposed action. If the calculated benefits from a project are greater than its costs, the action is recommended. When risks to life and health represent an important component of the costs, the term risk-benefit analysis is used.

The expected cost of a project is determined by enumerating all aversive consequences that might arise from its implementation, assessing the probability that each will occur, and estimating the cost or loss to society should each occur. Next, the expected loss from each possible consequence is calculated by multiplying the amount of the loss by the probability that it will be incurred. The expected loss of the entire project is computed by summing the expected losses associated with the various possible consequences. An analogous procedure produces an estimate of the expected benefits.

Performing a full-dressed analysis assumes, among other things, that all significant consequences can be enumerated in advance; that meaningful probability, cost and benefit judgments can all be reduced to dollar equivalents; that people really know how they value consequences today and how they will value them in the future; and that people want, or should want, to maximize the difference between expected benefits and losses (Fischhoff, 1977b).

Decision analysis. Cost-benefit analyses typically produce a yes-no decision about one particular project (i.e., its costs either do or do not outweigh its benefits). Often, however, we are faced with an array of alternative actions from which we must choose the best. In addition, cost-benefit analysis only considers consequences that can be expressed in dollars. For many amenities and disamenities (e.g., aesthetic improvement or degradation), however, monetarization is difficult.

Recently, a technique called decision analysis has been developed to handle situations with multiple alternatives and varied consequences. Decision analysis is hailed by some as "the" general method of choice for coping with risky decisions (Howard, 1968, 1975; Keeney & Raiffa, 1976; Raiffa, 1968). It combines sophisticated modeling of decision problems (i.e., the critical options, events and consequences) with a theory specifying how to deal rationally with uncertainty and the inevitable subjectivity of decision makers' preferences and values. Decision analysis has been applied to problems involving hurricane modification (Howard, Matheson & North, 1972), the selection of experiments for a Mars space mission (Matheson & Roths, 1967), the decision to undergo coronary artery surgery (Pauker, 1976), and the desirability of nuclear power plants (Matheson et al., 1968; Barrager, Judd & North, 1976).

In some ways, decision analysis is a generalized form of cost-benefit analysis. After identifying feasible courses of action, possible outcomes are specified and the probabilities and values of those outcomes are determined. The alternative of choice is the one with the highest expected value, the greatest preponderance of expected benefits over expected costs. As with cost-benefit analysis, fault trees could be used to assess probabilities and dollars could be used as the common unit of value. In fact, one could do a decision analysis of just one alternative (or, rather, of the two alternatives "act" and "don't act"), reducing the decision analysis to a slightly modified cost-benefit analysis. More often, however, the decision analyst will choose a common unit of value (usually called "utility"), with no necessary relation to dollars. This makes it easier to include consequences such as increased anxiety and aesthetic degradation which have no convenient dollar equivalent.

Decision analysis is characterized not only by its conceptual framework, but also by some of the techniques it uses. Much emphasis is placed on what might be called "dynamic" modeling of the decision problem, looking at how possible actions interact with world events. For example, a decision model might include a sequence like "build plant according to Plan B--<u>federal emission standards change</u>--retool plant--<u>litigation holds up operating permit</u>--plant starts up 6 months late." Probabilities are assigned to the world events (underlined above) and a value is assigned to the end state, a plant in operation after this course of events.

Decision analysts acknowledge that many (or even all) of the probabilities and values they use are not well defined. They handle this problem with "sensitivity analysis," a technique which leads them to ask questions such as: "How much of a difference would it make in the final decision if this probability were off by a factor of 10?" If, when each probability and the value assigned each consequence is varied widely enough to encompass whatever degree of uncertainty exists, the same action is still recommended, one's confidence in the decision is greatly increased. If changes in these values produce different decisions, one should be cautious and collect more data on those values whose fluctuations most influence the decision.

Decision analysis of waste disposal options. Although a full-scale decision analysis of nuclear waste options is beyond the scope of this paper, a brief outline of how such an analysis might proceed may be instructive.

Let us assume that nuclear power is generating high-level radioactive wastes and a decision has been made to dispose of these wastes permanently. The decision of interest concerns the optimal means of permanent disposal.

Many options are possible, including storage in geologic formations, the sea bed, or polar ice, or even extraterrestrial disposal. Each option has several sub-options (e.g., different geologic formations). An analysis of risks and benefits needs to be done for each option and sub-option.

We shall focus here on one aspect of one option: some of the social costs associated with geologic disposal in a cavity within a salt bed. As an additional restriction, even though these social costs are incurred both while the wastes are deposited and during the long time period after the repository is sealed, we shall consider only the latter period.

Figure 5 depicts a model for calculating the social costs of disposal in a salt bed after closure. First, all hazards that could trigger a radioactive release (and resultant public loss) are catalogued. The magnitude of loss will depend on the spatial and temporal distribution of the release, (i.e., how widely radioactivity is spread, how long it lasts). Hence, a dispersion model integrating geologic and demographic information is needed to predict human exposure and property contamination. A human implications model is also needed to specify the expected number of deaths and illnesses and the genetic and property damage resulting from this exposure. Finally, all these various damages must be converted to a measure of expected loss that can be integrated with the risks and benefits derived from the other segments of the analysis.

Insert Figure 5 about here

Figure 6 outlines the procedures for calculating social loss from salt-bed mishaps in greater detail. First, probabilities must be assessed for each possible category of radioactive release. For some categories, there are relevant historical data from which probabilities can be derived. This is the case for meteorite impacts. Blake (1968) has calculated that the probability of a meteorite making a crater with a width of 200-300 meters, in a land storage surface of 100 $km^2$, over a period of 100,000 years is about $10^{-5}$. For other categories, such as sabotage, we have little recourse but to use probability assessments based on expert judgment, perhaps aided by fault trees or other formal models. The dispersion and human implication models would then be used to calculate expected death and illness for each of the seven possible branches of radiation release.

Insert Figure 6 about here

Finally, a dollar or utility value would have to be assigned to each consequence. For example, one might assign a social cost of $3 million for each death. In that case, if we expected 20 immediate and 20 delayed fatalities from a meteorite impact, each valued at $3 million, the expected loss from this segment of the analysis would be $3,000,000 x 40 x $10^{-5}$ which equals $1200.[3] In similar fashion, the expected loss from illness, genetic effects, etc., would be computed for meteorite impact and all the other consequences of the remaining branches. Combining the expected losses from the seven branches would produce an overall expected social cost of salt-bed disposal. This value could then be combined with similarly derived figures for economic losses and social and economic benefits from salt-bed disposal. The result, representing the overall expected value of salt-bed disposal, could then be compared with the expected values of the other waste disposal options.

Although this analysis is crude and inadequate, it does serve to highlight the following basic questions that will be raised in connection with any thorough decision analysis.

Are the estimates accurate enough for decision making? Earlier in this paper, we discussed the many biases that can affect probability assessments. Certainly all possible steps should be taken to minimize such biases in a decision analysis. However, for some problems, no effective debiasing procedure is known; for others, the presence of bias may not even be recognized. In addition, the analysis must also accommodate disagreement among experts, regarding both probabilities and values assigned to consequences. Sensitivity analysis deals with these problems by using ranges of opinion instead of single estimates. However, unless some alternative stands out as best even when widely disparate estimates are used, decision analysis will not tell us what to do. One might argue that, in such cases, we simply can't know what course of action is best.

Some critics would argue that no estimates are good enough when the events of importance are extremely rare. For example, according to Holdren (1976):

> . . . the expert community is divided about the conceivable realism of probability estimates in the range of one in ten thousand to one in one billion per reactor year. I am among those who believe it to be impossible in principle to support numbers as small as these with convincing theoretical arguments

(that is, in the absence of operating experience in the range of 10,000 reactor-years or more), even ignoring the crucial possibility of malevolence. The reason I hold this view is straightforward: nuclear power systems are so complex that the probability the safety analysis contains serious errors (for example, that it omits failure modes more important than those included) is so big as to render meaningless the tiny computed probability of accident (p. 21).

Can the value of a life be quantified? Despite the aversiveness of thinking about life in economic terms, the fact is inescapable that by our actions, we put a finite value on our lives. Decisions to install safety features, to buy life insurance, and to accept a more hazardous job for high salary all carry implicit values for the worth of life.

Economists have long debated the question of how to quantify the value of a life or the value of a specified change in survival probability. Bergstrom (1974) argued that the best way to answer these questions is by observing the actual behavior of people trading risks for economic benefits. In this tradition, Thaler and Rosen (1975) studied salary as a function of occupational risk and concluded that a premium of about $200 (in 1967 dollars) per year was required to induce people in risky occupations to accept an annual increment in probability of accidental death of .001. From this, it can be argued that society should be willing to pay about $200,000 to avoid a statistical death. An extensive replication of the Thaler and Rosen study by Rappoport (1977) obtained a value of about $2,000,000. Howard (1977) inferred a value similar to Rappoport's when he asked people directly how much they would have to be paid to incur an additional .001 probability of death.

Should all modes of death be valued equally? Some proponents of nuclear power have complained about the public's apparent willingness to spend many times more money to prevent a fatality in the nuclear industry than to prevent other types of death. Is this so, and if so, is it unreasonable?

Unfortunately, we do not have good answers to these questions. However, some data from the previously mentioned study by Fischhoff et al. (1978) are suggestive. In that study, members of the League of Women Voters indicated that an activity's acceptable risk level depended upon a varity of qualitative characteristics (e.g., voluntariness, dread, immediacy of consequences, familiarity, and controllability). It might be desirable to weight the aversiveness of deaths according to these characteristics. One could argue, for example, that deaths from risks imposed involuntarily should be counted more in the calculations of loss than deaths from voluntary risks. The size of the weighting factor would be a matter for society to decide.

Potential for catastrophe (the loss of large numbers of lives at once) is another characteristic that might deserve higher negative weight, because of the horrific nature of catastrophes and because they may pose a greater threat to the survival of a community or society than do scattered individual deaths. Wilson (1975) has argued that the cost of N lives lost at once should be weighted by $N^2$ to take this into account.

Another consideration in nuclear waste decisions is the treatment of delayed deaths. Should the death of an individual who succumbs to leukemia 30 years after exposure to radiation be assigned the same social cost as a death within weeks after exposure? Traditional economic theory argues that it is better to delay costs and, therefore, delayed deaths should be discounted (i.e., considered as less serious). Any such discounting favors options that save bad consequences for later generations. Assuming a 5% annual discount rate, 1 immediate death would be equivalent to 1730 deaths 200 years from now. Use of a discount rate has been vigorously protested (Lovins, 1976) as encouraging shoddy standards (e.g., bridges that collapse in 20 years, since a loss then wouldn't be worth much). Critics argue that the discount rate should be zero if we wish to minimize future generations' regret about our generation's choices (Schulze, 1974).

At present, there is no generally accepted method for weighing intergenerational benefits and costs. It has been suggested (National Academy of Science, 1975) that, until a method is developed, benefits and costs should be computed over a wide range of discount rates to see if the effects are significant.

Are there higher-order consequences that need consideration? The model illustrated in Figure 6 may be inadequate because it considers only the direct costs such as harm to people and damage to the environment. What happens as a consequence of these first-order consequences, may be even more important. For example, a reactor accident may also lead to the shutdown of the whole nuclear industry for some period of time. This second-order consequence would affect virtually every facet of our society (through even higher-order consequences such as power shortages, lost jobs, and unheated homes). Because society is so vulnerable to such events, even a small accident could result in massive social disruption. Although there is no inherent reason why higher-order effects could not be incorporated in an analysis, doing so is difficult and may introduce an even higher level of uncertainty. Nonetheless, it is unlikely that traditional (intuitive and political) decision procedures, unaided by decision analysis, would do any better job of considering these subtle factors.

## Future Directions

For the immediate future, the primary tools brought to bear on risk management decisions will be various formal and semi-formal analytic techniques. The results of these analyses will go through the political wringer and some decisions will emerge. The research discussed in this chapter suggests a number of actions that interested citizens might take in order to involve themselves effectively in the decision-making process.

One action involves preparing themselves for the process. Preparation is ordinarily seen as a matter of boning up on the facts of an issue. The research on risk perception indicates a need to educate one's intuitions in order to understand these facts. That means learning to be wary of systematic biases, both those that come naturally and those that might be exploited by risk presenters. It means realizing where one's intuitions are not to be trusted and decision aids are needed. It means appreciating the limits of one's knowledge. H. G. Wells once observed, "Statistical reasoning will one day be as important for good citizenship as the ability to read and write." Perhaps that day has come.

The second set of tasks involves making formal analyses as useful as possible to the decision process. That means working to keep them honest. Just like eye-witness testimony, the evidence emerging from such analyses can be biased, incomplete and half-truthful. Although such analyses may be the exception, a watchful eye is prudent. The usefulness of an analysis may also depend on its modesty. The recognition of uncertainty is a central facet of decision analysis. Nonetheless, analysts enmeshed in their work may not realize the limits of their own critical powers or the limits of the scientific knowledge on which they rely. Inadequate caution is encouraged by the pressures of public situations (with their demand for definitive statements) and the desire to affect decisions.

Just as performing an analysis requires skills from many disciplines, it is difficult for any one individual to criticize all facets. Concerned citizens might formulate their own independent review teams with experts from different fields. This strategy has apparently been successfully implemented in Massachusetts (Smardon & Woodland, 1976-7).

Another task is to make sure that analysts make their assumptions explicit. There is no reason to expect most citizens to realize the ethical implications underlying the use of a positive social discount rate, nor to recognize which alternative courses of action have been summarily ignored in the analysis. The public should insist on being told what assumptions about life and data it is accepting along with the "facts" of an analysis.

We believe that the basic function of formal analyses should be to help policy makers gain insight into complex issues. In our view, the important element is not the bottom line, but the process of reaching it. Policy makers need to understand the assumptions and calculations that led to the result. The analysis should shift debate from the decision itself to the critical impacts and assumptions, highlighting the most sensitive issues for scrutiny.

Lovins (1976) warns about " . . . the delicacy of the balance between drawing on expertise and smothering democracy with it" (p. 114). Accordingly, the public must work to keep societal decision processes open and responsive. Analysts and regulators are paid out of public funds. They should make their analyses comprehensible, solicit public input, and reflect public desires in their conclusions. A strong case can be made that the future of the democratic process is being shaped in these risk-management decisions. Excluding the public here may create a new technocratic elite, in effect declaring the public to be technically incompetent and disenfranchising it from a broad range of important decisions.

References

Barrager, S. M., Judd, B. R. & North, D. W. Decision analysis of energy alternatives: A comprehensive framework for decision making. Palo Alto, CA: Stanford Research Institute, 1976.

Bergstrom, T. C. Preference and choice in matters of life and death. Appendix 1 in Report No. ENG 7478, School of Engineering and Applied Science, UCLA, Nov. 1974.

Blake, V. E. A prediction of the hazards from the random impact of meteorites on the earth's surface. Albuquerque, N. M.: Sandia Labs Aerospace Nuclear Safety Report SC-RR-68-838, 1968.

Boffey, P. M. Nuclear war: Federation disputes Academy on how bad effects would be. Science, 1975, 190, 248-250.

Borch, K. The economics of uncertainty. Princeton, NJ: Princeton University Press, 1968.

Bryan, W. B. Testimony before the Subcommittee on State Energy Policy, Committee on Planning, Land Use, and Energy, California State Assembly, February 1, 1974.

Coates, J. F. The role of formal models in technology assessment. Technological Forecasting and Social Change, 1976, 8.

Cohen, B. L. Perspectives on the nuclear debate. Bulletin of the Atomic Scientists, 1974, 30(9), 35-39.

Cohen, J. J. A case for benefit-risk analysis. In H. J. Otway (Ed.), Risk vs. benefit: Solution or dream, Report LA-4860-MS, Los Alamos Scientific Laboratory, February, 1972 (available from the National Technical Information Service).

Comey, D. D. How we almost lost Alabama. Chicago Tribune, August 31, 1975, p. 2/1.

David, E. E. One-armed scientists? Science, 1975, 189, 891.

Fischhoff, B. Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. Journal of Experimental Psychology: Human Perception and Performance, 1975a, 1, 288-299.

Fischhoff, B. Hindsight: Thinking backward? Psychology Today, 1975b, 8.

Fischhoff, B. Perceived informativeness of facts. Journal of Experimental Psychology: Human Perception and Performance, 1977a, 3, 349-358.

Fischhoff, B. Cost-benefit analysis and the art of motorcycle maintenance. Policy Sciences, 1977b, 8, 177-202.

Fischhoff, B. & Beyth, R. "I knew it would happen":--Remembered probabilities of once-future things. Organizational Behavior and Human Performance, 1975, 13, 1-16.

Fischhoff, B., Slovic, P. & Lichtenstein, S. Knowing with certainty: The appropriateness of extreme confidence. Journal of Experimental Psychology: Human Perception and Performance, 1977, 3, 552-564.

Fischhoff, B., Slovic, P. & Lichtenstein, S. Fault trees: Sensitivity of estimated failure probabilities to problem representation. Journal of Experimental Psychology: Human Perception and Performance, 1978, 4, 342-355.

Fischhoff, B., Slovic, P., Lichtenstein, S., Read, S. & Combs, B. How safe is safe enough? A psychometric study of attitudes towards technological risks and benefits. Policy Sciences, 1978, 8, 127-152.

Gray, E. F. Remarks on the problem of nuclear waste management. Proceedings of the 1976 Conference on Public Policy Issues in Nuclear Waste Management. Washington, D. C., National Science Foundation, 1976, 195-202.

Green, A. E. & Bourne, A. J. Reliability technology. New York: Wiley Interscience, 1972.

Hall, W. K. Why risk analysis isn't working. Long Range Planning, December 1975, 23-29.

Herrero, S. Human injury inflicted by grizzly bears. Science, 1970, 170, 593-597.

Hohenemser, K. H. The failsafe risk. Environment, 1975, 17, 6-10.

Holdren, J. P. The nuclear controversy and the limitations of decision making by experts. Bulletin of the Atomic Scientists, 1976, 32(3), 20-22.

Howard, R. A. The foundations of decision analysis. IEEE Transactions on Systems, Science and Cybernetics, 1968, 4, 211-219.

Howard, R. A. Social decision analysis. Proceedings of the IEEE, 1975, 63, 359-371.

Howard, R. A. Life or death decision making. Unpublished manuscript, Stanford: Department of Engineering, Stanford University, 1977.

Howard, R. A., Matheson, J. E. & North, D. W. The decision to seed hurricanes. Science, 1972, 176, 1191-1202.

Hynes, M. & Vanmarcke, E. Reliability of embankment performance predictions. Proceedings of the ASCE Engineering Mechanics Division Specialty Conference, Waterloo, Ontario, Canada: University of Waterloo Press, 1976.

Kates, R. W. Hazard and choice perception in flood plain management. Research Paper 78, Chicago: Department of Geography, University of Chicago, 1962.

Kates, R. W. Risk assessment of environmental hazard. New York: Wiley, 1978.

Katz, D. L. & West, H. H. The overall problem--risk/benefit for LNG shipping and storage. In D. Okrent (Ed.), Risk-benefit methodology and application: Some papers presented at the Engineering Foundation Workshop, Asilomar, Ca. Los Angeles: Report ENG-7598, School of Engineering and Applied Science, UCLA, December 1975, 1-40.

Keeney, R. L. & Raiffa, H. Decisions with multiple objectives: Preferences and value tradeoffs. New York: Wiley, 1976.

Lichtenstein, S., Fischhoff, B. & Phillips, L. D. Calibration of probabilities: The state of the art. In H. Jungermann & G. deZeeuw (Eds.), Decision Making and Change in Human Affairs. Amsterdam: D. Reidel, 1977.

Lichtenstein, S. & Slovic, P. Response-induced reversals of preference in gambling: An extended replication in Las Vegas. Journal of Experimental Psychology, 1973, 101, 16-20.

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M. & Combs, B. Judged frequency of lethal events. Journal of Experimental Psychology: Human Learning and Memory, 1978, 4, 551-578.

Lifton, R. J. Nuclear energy and the wisdom of the body. Bulletin of the Atomic Scientists, 1976, 32(7), 16-20.

Lovins, A. B. Comments in S. M. Barrager, B. R. Judd & D. W. North, Decision analysis of energy alternatives. Palo Alto, Ca.: Stanford Research Institute, 1976.

Lowrance, W. W. Of acceptable risk. Los Altos, Ca.: Wm. Kaufmann, 1976.

Matheson, J. E. & Roths, W. J. Decision analysis of space projects: Voyager Mars. Paper presented at the National Symposium "Saturn, Apollo and Beyond," American Astronautical Society, June, 1967.

Matheson, J. E., Cazalet, E. G., North, D. W., Rousseau, W. F. & Howard, R. A. Decision analysis of nuclear plants in electrical system expansion. Final Report, SRI Project 6469, Menlo Park, Ca.: Stanford Research Institute, December 1968.

McGrath, P. E. Radioactive waste management: Potentials and hazards from a risk point of view. Report EURFNR-1204 (KFK 1992), Karlsruhe, Germany: U.S.-EURATOM Fast Reactor Exchange Program, June 1974.

National Academy of Sciences. Decision making for regulating chemicals in the environment. Appendix H, 1975, 163-196.

Nelkin, D. The role of experts in a nuclear siting controversy. Bulletin of the Atomic Scientists, 1974, 30(9), 29-36.

Otway, H. J. & Cohen, J. J. Revealed preferences: Comments on the Starr benefit-risk relationships. Research Memorandum 75-5, Laxenburg, Austria: International Institute for Applied Systems Analysis, March 1975.

Otway, H. J. & Pahner, P. D. Risk assessment. Futures, 1976, 8, 122-134.

Pahner, P. D. The psychological displacement of anxiety: An application to nuclear energy. In D. Okrent (Ed.), Risk-benefit methodology and application: Some papers presented at the Engineering Foundation Workshop, Asilomar, Ca., Los Angeles: Report ENG-7598, School of Engineering and Applied Science, UCLA, December, 1975, 557-580.

Pauker, S. G. Coronary artery surgery: The use of decision analysis. Annals of Internal Medicine, 1976, 85, 8-18.

Primack, J. Nuclear reactor safety: An introduction to the issues, Bulletin of the Atomic Scientists, 1975, 31(9), 15-17.

Rabinowitch, E. Living dangerously in the age of science. Bulletin of the Atomic Scientists, 1972, 28(1), 5-8.

Raiffa, H. Decision Analysis. Reading, Ma.: Addison-Wesley, 1968.

Rappoport, E. Unpublished Ph.D. dissertation, Department of Economics, UCLA, 1977.

Ross, L. The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), Advances in Experimental Social Psychology, New York: Academic Press, 1977, 173-220.

Rowe, W. D. An anatomy of risk. New York: Wiley, 1977.

Schultze, W. Social welfare functions for the future. American Economist, 1974, 18, 70-81.

Sinsheimer, R. F. The brain of Pooh: An essay on the limits of mind. American Scientist, 1971, 59, 20-28.

Slovic, P. & Fischhoff, B. On the psychology of experimental surprises. Journal of Experimental Psychology: Human Perception and Performance, 1977, 3, 544-551.

Slovic, P., Fischhoff, B. & Lichtenstein, S. Behavioral decision theory, Annual Review of Psychology, 1977, 28, 1-39.

Slovic, P., Kunreuther, H. & White, G. F. Decision processes, rationality, and adjustment to natural hazards. In G. F. White (Ed.), Natural hazards, local, national and global. New York: Oxford University Press, 1974.

Smardon, R. C. & Woodland, R. B. Some preliminary results of an environmental impact report review process. Journal of Environmental Systems, 1976-7, 6, 209-228.

Starr, C. Social benefit versus technological risk. Science, 1969, 165, 1232-1238.

Thaler, R. & Rosen, S. The value of saving a life: Evidence from the labor market. In N. E. Terleckyj (Ed.), Household production and consumption, National Bureau of Economic Research, 1975.

Tversky, A. & Kahneman, D. Judgment under uncertainty: Heuristics and biases. Science, 1974, 185, 1124-1131.

U. S. Government. Teton Dam Disaster. Committee on Government Operations, Washington, D. C., 1976.

U. S. Nuclear Regulatory Commission. Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants. WASH 1400 (NUREG-75/014), Washington, D. C.: The Commission, October 1975.

U. S. Nuclear Regulatory Commission. Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission, NUREG/CR-0400, Washington, D. C.: The Commission, September 1978.

Weinberg, A. M. The maturity and future of nuclear energy. American Scientist, 1976, 64, 16-21.

Wilson, R. Examples in risk-benefit analysis. Conference on Advanced Energy Systems, Chemtek, 1975, 5, 604.

Zebroski, E. L. Attainment of balance in risk-benefit perceptions. In D. Okrent (Ed.), Risk-benefit methodology and application: Some papers presented at the Engineering Foundation Workshop, Asilomar, Ca.,Report ENG-7598, School of Engineering and Applied Science, UCLA, December 1975, 633-644.

Acknowledgments

We are indebted to Sarah Lichtenstein, Leroy Gould and Scott Slovic for their helpful comments on earlier drafts of this chapter. This work was supported by the National Science Foundation under Grants ENV77-15332 and OSS-16564 to Perceptronics, Inc. Any opinions, findings and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Footnotes

1. More extensive discussions of heuristics and biases in probabilistic thinking are available in articles by Slovic, Kunreuther and White (1974), Slovic, Fischhoff and Lichtenstein (1977), and Tversky and Kahneman (1974).

2. Fault trees start with a particular undesired final event (a failure of the system) and work backward to identify the component failures needed to initiate that event. A related method uses event trees. These start from a particular initiating event (e.g., an earthquake in a [illegible] waste storage area) and project all possible outcomes of that event.

3. We could have worked through this calculation with utilities, but the (dis)utility assigned to death would have meant little intuitively without knowing what values were assigned to other consequences. That is, one utility unit per death could be a lot or a little depending on the utility assigned to each illness, etc.

Table 1

Bias in Judged Frequency of Death

| Most Overestimated | Most Underestimated |
| --- | --- |
| All accidents | Smallpox vaccination |
| Motor vehicle accidents | Diabetes |
| Pregnancy, childbirth, and abortion | Stomach cancer |
| Tornadoes | Lightning |
| Flood | Stroke |
| Botulism | Tuberculosis |
| All cancer | Asthma |
| Fire and flames | Emphysema |
| Venomous bite or sting | |
| Homicide | |

## Figure Captions

1. Fault tree indicating the possible ways that radioactivity could be released from wastes deposited in bedded salt (after closure of the repository). (From McGrath, 1974)

2a. Comparison between nuclear power and X rays on nine risk characteristics. (From Fischhoff et al., 1978)

2b. Comparison between nuclear power and non-nuclear electric power on nine risk characteristics. (From Fischhoff et al., 1978)

3. A fault tree indicating the ways in which an automobile might fail to start. (From Fischhoff, Slovic & Lichtenstein, 1978)

4. Copyright 1974 The Chicago Sun Times. Reproduced by permission of Wil-Jo Associates, Inc. and Bill Mauldin.

5. A model for calculating the social costs from waste disposal in bedded salt.

6. Human and property consequences of radiation release from salt bed disposal site.