

00

Regulation of Risk: A Psychological Perspective

Paul Slovic, Baruch Fischhoff and Sarah Lichtenstein

Decision Research
A Branch of Perceptronics
Eugene, Oregon 97401

DRAFT

April 1982

Regulation of Risk: A Psychological Perspective

Paul Slovic, Baruch Fischhoff and Sarah Lichtenstein

Decision Research
A Branch of Perceptronics
Eugene, Oregon 97401

Decision Research Report 82-11
April 1982

This paper was prepared for the conference on Social Science and Regulatory Policy, sponsored by the Regulation Program of the National Science Foundation and held at the Sheraton International Conference Center, Reston, Virginia, January 22-23, 1982.

Table of Contents

	<u>Page</u>
Introduction	1
Overview of this chapter	2
Confronting Human Limitations	4
Bounded Rationality	5
Judgmental Biases in Risk Perception	6
Availability	6
Overconfidence	9
Uncertain Preferences: Difficulties in Evaluating Risk	11
Labile values	12
Unstable decision frames	13
Implications for Regulation	21
Implications for policy	21
Implications for assessing values	23
Implications for informing people about risk	24
Studies of Protective Behavior	26
National Flood Insurance Program	26
Seat Belts	29
Characterizing Perceived Risk	30
The Psychometric Paradigm	31
Implications of Risk Perception Research	34
Forecasting public response	34
Comparing risks	36
Deciding "How Safe Is Safe Enough?"	37
Approaches to Acceptable Risk: A Critique	38
Toward a Safety Goal	43
Justification	43
Technical issues	45
Social issues	47
Implications of risk perception research	49
References	55

Draft April 1982

Regulation of Risk: A Psychological Perspective

Technology has enhanced society's ability to utilize the environment, eradicate dread diseases, and fashion a life of comfort and leisure. But it has become increasingly apparent that these benefits are accompanied by a variety of hazardous side effects. Hardly a day passes that does not reveal some new danger in our foods, our homes, our workplaces, our leisure activities, or our natural environments.

In a world where more and more Americans are coming to see themselves as the victims rather than the beneficiaries of technology, it is not surprising that the control of hazards has become a major concern of society and a growing responsibility of government. During the past decade, and unprecedented assemblage of powerful regulatory bureaucracies has been created and charged with answering myriad forms of the question "How safe is safe enough?" The cost of these efforts to manage risk has been estimated at \$1.4-\$4.2 billion per year (Fischhoff, Hohenemser, Kasperson & Kates, 1978). Yet despite this massive effort (some would say because of it), the public feels increasingly vulnerable to the risks from technology and believes that the worst is yet to come (see Tables 1a and 1b). Regulatory agencies have been embroiled in rancorous conflicts, caught between a fearful and unsatisfied public on one side and frustrated technologists and industrialists on the other (Table 1c). The latter see the pursuit of a "zero-risk society" as trampling individual rights, thwarting innovation, and jeopardizing the nation's economic and political stability. The way in which these conflicts are resolved will affect not just the fate of particular technologies but the fate of our society and its social organization as well.

Insert Table 1 about here

The urgent need to reduce conflict and improve the management of risk poses a challenge to scholars and scientists from many disciplines. Technical issues, such as the identification of hazards and the assessment of the probability and magnitude of their consequences, require the efforts of physical scientists, biological scientists, and engineers. Social issues, pertaining to the evaluation of these identified hazards and the decision making invoked in their management, require the knowledge and skills of lawyers, political scientists, geographers, sociologists, and economists.

Psychology, too, can contribute to the management and regulation of risk. In recent years, empirical and theoretical research on the psychology of decision making under risk has produced a body of knowledge that should be of value to those who seek to understand and improve societal decisions. This chapter aims to describe some components of this research. Rather than being comprehensive, we focus on investigations grounded in cognitive psychology, psychometrics, and decision theory, with emphasis on our own research program and the related work of a few colleagues. Our efforts are guided by the assumption that those who promote and regulate high-risk technologies need to understand the ways in which people think about risk, the outcomes of technology that people deem to be important, and the values that they attach to these outcomes. Without such understanding, well-intended policies may be ineffective, perhaps even counterproductive.

Overview of this chapter

This chapter is divided into four sections. In the first section we describe some of the research on human intellectual limitations, with

particular focus on the problems that occur when people seek to make sense out of a probabilistic environment and attempt to resolve the value conflicts arising from decisions about beneficial but hazardous activities. We point out the difficulties people have in thinking intuitively about risk and uncertainty. We argue that people's perceptions of the world are deficient and that their preferences are sometimes unstable, vague, or inconsistent. The results of this research run counter to the traditional presumptions of knowledge and rationality that underlie economic approaches to decision making under risk. In doing so, these results challenge the viability of market mechanisms for managing risk and thereby suggest that governmental regulation of risk is needed. A common reaction of industry and government officials to evidence of ignorance, misinformation, or faulty thinking has been to call for educational programs to correct these shortcomings. While we applaud such efforts and do believe that people are educable, our emphasis is on the obstacles such programs must overcome.

The second section of this chapter shows how some of the findings of the first section can be applied to two specific policy problems. Both the problems described here concern the extent to which people engage in actions that protect them from risk: purchasing flood insurance and wearing seat belts. We argue that educational programs designed in the light of cognitive research are more likely to be effective than programs not based on an understanding of the causes of people's resistance to protective behavior.

In the third section, we describe research on the perception of risk. This research explores what people mean when they say that a technology or activity is risky. We find that many attributes other than death rates

determine judgments of riskiness. Such attributes include catastrophic potential, risk to future generations, and dread. In contrast with the first section, the tone of this section is optimistic. We find that laypeople have strong, consistent and reasonable views about risk. In fact, their model of what constitutes risk appears to be much richer than that held by most technical experts.

The final section is a discussion of the problems encountered in trying to answer the question, "How safe is safe enough?" We first discuss the general nature of acceptable-risk problems and the approaches one could take to solve them. We then discuss the issue of establishing safety goals and standards, with particular emphasis on nuclear reactor safety. In our view, social issues have been neglected in past proposals for safety goals. In this section we discuss the implications of risk-perception research for the incorporation of social issues in safety-standard decisions.

Confronting Human Limitations

The traditional view of human mental processes assumes that we are intellectually gifted creatures. Shakespeare referred to man as "noble in reason, infinite in faculties...the beauty of the world, the paragon of animals." Economic theory, with its presumptions of well informed, rational (utility maximizing) decision makers, has echoed this theme. As economist Frank Knight put it, "We are so built that what seems reasonable to us is likely to be confirmed by experience or we could not live in the world at all" (1921, p. 227). Given freedom of choice, adequate information, and rational decision making, people's behavior in the marketplace should manage risk adequately. In such a world, there would be little need for regulatory interventions.

Over the years, economists and others have come to recognize that the idealized views on which the free market approach is based do not hold. The real issue is not whether to regulate, but how much to regulate and how to do it. Provision of information to consumers, standards, restrictions, and bans compete as alternative modes of intervention (Joskow & Noll, 1980). To make effective choices among these options, it is necessary to know something of the nature and extent of human intellectual limitations. Although we are far from having complete understanding, much has been learned in recent years that may have implications for regulatory policy.

Bounded Rationality

An important early critic of the rational model's descriptive adequacy was Herbert Simon, who drew upon psychological research to challenge traditional assumptions about the motivation, omniscience, and computational capacities of "economic man." As an alternative to utility maximization, Simon (1957) introduced the notion of "bounded rationality," which asserts that cognitive limitations force people to construct simplified models of the world in order to cope with it. To predict behavior "...we must understand the way in which this simplified model is constructed, and its construction will certainly be related to 'man's' psychological properties as a perceiving, thinking, and learning animal" (p. 198).

During the past twenty years, the skeleton theory of bounded rationality has been fleshed out. We have learned much about human cognitive limitations and their implications for behavior--particularly with regard to decisions made in the face of uncertainty and risk. Numerous studies show that people (including experts) have great difficulty judging probabilities, making pre-

dictions, and otherwise attempting to cope with uncertainty. Frequently these difficulties can be traced to the use of judgmental heuristics, which serve as general strategies for simplifying complex tasks. These heuristics are valid in many circumstances, but in others they lead to large and persistent biases with serious implications for decision making. Much of this work has been summarized by Kahneman, Slovic and Tversky (1982), Nisbet and Ross (1980), Slovic, Fischhoff and Lichtenstein (1977), and Tversky and Kahneman (1974). In the remainder of this section we shall discuss two general manifestations of bounded rationality that are particularly relevant to regulation of risk. These topics are judgmental biases and uncertain preferences.

Judgmental Biases in Risk Perception

If people are to respond optimally to the risks they face, they must have reasonably accurate perceptions of the magnitude of those risks. Yet the formal education of most laypeople rarely includes any serious instruction in how to assess risks. Their subsequent learning is typically restricted to unsystematic personal experience and news media reports. Perhaps it should not be surprising that people often are misinformed, rely on suboptimal risk assessment heuristics, and fail to understand the limits of their own knowledge.

Availability. On inferential strategy that has special relevance for risk perception is the availability heuristic (Tversky & Kahneman, 1973). People using this heuristic judge an event as likely or frequent if instances of it are easy to imagine or recall. Because frequently occurring events are generally easier to imagine or recall than are rare events, availability is often an appropriate cue. However, availability is also affected by factors

unrelated to frequency of occurrence. For example, a recent disaster or a vivid film could seriously bias risk judgments.

Availability bias is illustrated by several studies in which people judged the frequency of 41 causes of death (Lichtenstein et al., 1978). In one study, these people were first told the annual death toll in the United States for one cause (50,000 deaths from motor vehicles accidents) and then asked to estimate the frequency of the other 40 causes. Figure 1 compares the judged number of deaths per year with the number reported in public health statistics. If the frequency judgments equalled the statistical rates, all data points would fall on the identity line. Although more likely hazards generally evoked higher estimates, the points were scattered about a curved line that lay sometimes above and sometimes below the line of accurate judgment. In general, rare causes of death were overestimated and common causes of death were underestimated. In addition to this general bias, sizable specific biases are evident in Figure 1. For example, accidents were judged to cause as many deaths as diseases, whereas diseases actually take about 15 times as many lives. Homicides were incorrectly judged as more frequent than diabetes and stomach cancer deaths. Pregnancies, births, and abortions were judged to take about as many lives as diabetes, though diabetes actually causes about 80 times more deaths. In keeping with availability considerations, causes of death that were overestimated (relative to the curved line) tended to be dramatic and sensational (accidents, natural disasters, fires, homicides), whereas underestimated causes tended to be unspectacular events that claim one victim at a time and are common in non-fatal form (e.g., smallpox vaccinations, stroke, diabetes, emphysema).

Insert Figure 1 about here

The availability heuristic highlights the vital role of experience as a determinant of perceived risk. If one's experiences are misleading, one's perceptions are likely to be inaccurate. Unfortunately, much of the information to which people are exposed provides a distorted picture of the world of hazards. One result of this is that people tend to view themselves as personally immune to certain kinds of hazards. Research shows that the great majority of individuals believe themselves to be better than average drivers (Svenson, 1981), more likely than average to live past 80 years old (Weinstein, 1980), less likely than average to be harmed by products that they use (Rethans, 1979), and so on. Although such perceptions are obviously unrealistic, the risks may look very small from the perspective of each individual's experience. Consider automobile driving: despite driving too fast, following too closely, etc., poor drivers make trip after trip without mishap. This personal experience demonstrates to them their exceptional skill and safety. Moreover, their indirect experience via the news media shows that when accidents happen, they happen to others. Given such misleading experiences, people may feel quite justified in refusing to take protective actions such as wearing seat belts (Slovic, Fischhoff & Lichtenstein, 1978).

In some situations, failure to appreciate the limits of available data may lull people into complacency. For example, we asked people to evaluate the completeness of a fault tree showing the problems that could cause a car not to start when the ignition key was turned (Fischhoff et al., 1978).

Respondent's judgments of completeness were about the same when looking at the

full tree as when looking at a tree in which half of the causes of starting failure were deleted. In keeping with the availability heuristic, what was out of sight was also out of mind.

Overconfidence. A particularly pernicious aspect of heuristics is that people typically have too much confidence in judgments based upon them. In another follow-up to the study on causes of death, people were asked to indicate the odds that they were correct in choosing the more frequent of two lethal events (Fischhoff et al., 1977). Odds of 100 : 1 or greater were given often (25% of the time). However, about one out of every eight answers associated with such extreme confidence was wrong (fewer than 1 in 100 would have been wrong had the odds been appropriate). At odds of 10,000 : 1, people were wrong about 10% of the time. The psychological basis for this unwarranted certainty seems to be an insensitivity to the tenuousness of the assumptions upon which one's judgments are based. For example, extreme confidence in the incorrect assertion that homicides are more frequent than suicides may occur because people fail to appreciate that the greater ease of recalling instances of homicides is an imperfect basis for inference.

Overconfidence manifests itself in other ways as well. A typical task in estimating uncertain quantities such as failure rates is to set upper and lower bounds so that there is a certain fixed probability that the true value lies between them. Experiments with diverse groups of people making many different kinds of judgments have found that true values tend to lie outside of the confidence boundaries much too often. Results with 98% bounds are typical. Rather than 2% of the true values falling outside such bounds, 20-50% usually do so (Lichtenstein, Fischhoff & Phillips, 1982). Thus people

think that they can estimate uncertain quantities with much greater precision than they actually can.

Unfortunately, once experts are forced to go beyond their data and rely on judgment, they may be as prone to overconfidence as laypeople. Fischhoff et al. (1978) repeated their fault-tree study with professional automobile mechanics (averaging about 15 years of experience) and found them to be about as insensitive as laypersons to deletions from the tree. 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. None of the bounds specified by those individuals actually enclosed the true failure height.

Further evidence of expert overconfidence may be found in many technical risk assessments. For example, an official review of the Reactor Safety Study concluded that despite the study's careful attempt to calculate the probability of a core meltdown in a nuclear reactor, "we are certain that the error bands are understated. We cannot say by how much. Reasons for this include an inadequate data base, a poor statistical treatment [and] an inconsistent propagation of uncertainties throughout the calculation" (U.S. NRC, 1978, p. vi). The 1976 collapse of the Teton Dam provides another case in point. The Committee on Government Operations attributed this disaster to the unwarranted confidence of engineers who were absolutely certain that they had solved the many serious problems that arose during construction (U.S. Government, 1976).

Uncertain Preferences: Difficulties in Evaluating Risk

As difficult as it may be to perceive hazards and to deal with their inherent uncertainties, the process of evaluating the good and bad outcomes associated with hazardous activities and making decisions as to the management of these activities appears to be at least as demanding. Yet evaluation is the heart of decision making for the individual and the policy maker. Simon (1956) is well known for his assertion that boundedly rational individuals are motivated to obtain only some satisfactory, and not necessarily maximal, level of achievement, a tendency he labelled "satisficing." In recent years, criticism of the traditional utility maximization theory has taken a somewhat different turn. With regard to the consequences of hazardous activities, many investigators now question the assumptions that people have stable, precise values or utilities, which are readily measurable and are integrated with probabilistic considerations by means of the expected utility principle (see, e.g., Arrow, 1981; Fischhoff, Goitein & Shapira, 1982; Kahneman & Tversky, 1979; March, 1978; Schoemaker, 1982; Tversky & Kahneman, 1981; Fischhoff, Slovic & Lichtenstein, 1980a, b; Slovic, Fischhoff & Lichtenstein, 1982). Evidence is mounting in support of the view that our values are often not clearly apparent, even to ourselves, that methods for measuring values are intrusive and biased, that the structure of any decision problem is psychologically unstable, and that the processes whereby elusive values are integrated into decisions within such unstable structures lead to actions that differ in dramatic ways from the predictions of utility theory. Although space does not permit a full discussion of this new view, some of its elements are discussed briefly below.

Labile values. Along with contemplating the probabilities of various decision consequences, we must assess how desirable they are. What do we want to happen? What do we want to avoid? How badly? Such questions would seem to be the last redoubt of unaided intuition. Who knows better than the individual what he or she prefers? When one is considering simple, familiar events with which people have direct experience, it may be reasonable to assume that they have well-articulated preferences. But that may not be so in the case of the novel, unfamiliar consequences potentially associated with outcomes of events such as surgery, carbon-dioxide-induced climatic change, nuclear meltdowns, or genetic engineering. In these and other circumstances, our values may be incoherent, not sufficiently thought out (Fischhoff, Slovic & Lichtenstein, 1980a, b). When we think about acceptable levels of risk, for example, we may be unfamiliar with the terms used in such debates (e.g., social discount rates, miniscule probabilities, megadeaths). We may have contradictory values (e.g., a strong aversion to catastrophic losses of life but an awareness that we are no more moved by a plane crash with 500 fatalities than one with 300). We may occupy different roles in life (parents, workers, children) each of which produces clear-cut but inconsistent values. We may vacillate between incompatible but strongly held positions (e.g., freedom of speech is inviolate, but it should be denied authoritarian movements). We may not even know how to begin thinking about some issues (e.g., the appropriate tradeoffs between the outcomes of surgery for cancer vs. the very different outcomes from radiation therapy). Our views may change so much over time (say, as we near the hour of decision or of experiencing the consequences) that we are disoriented as to what we really think.

At least one regulator has appreciated the difficulties posed by uncertain and labile values. Former FDA Commissioner, Donald Kennedy, wrote:

There is genuine public ambivalence about risk in this country—not about how much risk there is, but about how much risk we want. Our citizens are uncertain about how much government intervention in the interest of their health they will tolerate, and they are also uncertain about how many other things—new inventions, creature comforts, old habits, progress—they are prepared to sacrifice to become safer. One does not resolve such doubts merely by stating the risks more precisely (Kennedy, 1981, p. 60).

Competent technical analyses may tell us what primary, secondary, and tertiary consequences to expect, but not what these consequences really entail. To some extent we are all prisoners of our past experiences, unable to imagine drastic changes in our world or health or relationships.

Unstable decision frames. In addition to the uncertainties that sometimes surround our values, perceptions of the basic structure of a decision problem are also unstable. The acts or options available, the possible outcomes or consequences of those acts, and the contingencies or conditional probabilities relating outcomes to acts make up what Tversky and Kahneman (1981) have called the "decision frame." Much as changes in vantage point induce alternative perspectives on a visual scene, the same decision problem can be subject to many alternative frames (see Figure 2). Whichever frame a decision maker adopts is determined in part by the external formulation of the problem and in part by the standards, habits, and personal characteristics of the decision maker. Tversky and Kahneman demonstrate that normatively inconsequential

changes in the framing of decision problems significantly affect preferences. These effects are noteworthy because they are sizable (often complete reversals of preference), because they violate important consistency and coherence requirements of economic theories of choice, and because they influence not only behavior, but how the consequences of behavior are experienced.

Insert Figure 2 about here

Tversky and Kahneman (1981) present numerous illustrations of framing effects, one of which involves the following pair of problems, given to separate groups of respondents.

1. Problem 1. Imagine that the U.S. is preparing for the outbreak of an unusual disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the consequences of the programs are as follows: If Program A is adopted, 200 people will be saved. If Program B is adopted, there is $1/3$ probability that 600 people will be saved, and $2/3$ probability that no people will be saved. Which of the two programs would you favor?

2. Problem 2. (Same cover story as Problem 1.) If Program C is adopted, 400 people will die. If Program D is adopted, there is $1/3$ probability that nobody will die, and $2/3$ probability that 600 people will die. Which of the two programs would you favor?

The preference patterns tend to be quite different in the two problems. In a study of college students, 72% of the respondents chose Program A over Program B and 78% chose Program D over Program C. Another study, surveying

physicians, obtained very similar results. On closer examination, we can see that the two problems are essentially identical. The only difference between them is that the outcomes are described by the number of lives saved in Problem 1 and the number of lives lost in Problem 2.

To explain these and many other violations of utility theory, Kahneman and Tversky (1979; see also Tversky & Kahneman, 1981) have developed a descriptive model, called prospect theory. The primitives of this theory are a value function, $v(x)$, which attaches a subjective worth to each possible outcome of a gamble or prospect, and a probability weighting function, $\pi(p)$, which expresses the subjective importance attached to the probability of obtaining a particular outcome. The attractiveness of a gamble that offers a chance of p to gain x and a chance of q to lose y would be equal to $\pi(p)v(x) + \pi(q)v(y)$. The value function is defined on gains and losses relative to some psychologically meaningful (neutral) reference point. A second feature of the value function is that the function is steeper for losses than for gains, meaning that a given change in one's status hurts more as a loss than it pleases as a gain. A third feature is that it is concave above that reference point and convex below it, meaning, for example, that the subjective difference between gaining (or losing) \$10 and \$20 is greater than the difference between gaining (or losing) \$110 and \$120.

Perhaps the most notable feature of the probability weighting function is the great importance attached to outcomes that will be received with certainty. Thus, for example, the prospect of losing \$50 with probability 1.0 is more than twice as aversive as the prospect of losing the same amount with

probability .5. For intermediate probabilities, the weighting function is somewhat insensitive to changes in probability. For example, a .5 chance of winning \$50 would not be 25% more attractive than a .4 chance of winning \$50.

The way a problem is framed determines both the reference point (the zero point) of the value function and the probabilities that are evaluated. If π and v were linear functions, preferences among options would be independent of the framing of acts, outcomes, or contingencies. Because of the characteristic nonlinearities of π and v , however, different frames often lead to different decisions.

One important class of framing effects deals with a phenomenon that Tversky and Kahneman (1981) have called "pseudocertainty." Mental representations of protective actions may be easily manipulated so as to vary the apparent certainty with which they prevent harm. For example, an insurance policy that covers fire but not flood could be presented either as full protection against the specific risk of fire or as a reduction in the overall probability of property loss. Because outcomes that are merely probable are undervalued in comparison with outcomes that are obtained with certainty, Tversky and Kahneman hypothesized that the above insurance policy should appear more attractive in the first context (pseudocertainty), which offers unconditional protection against a restricted set of problems. We have tested this conjecture in the context of one particular kind of protection, vaccination. Two forms of a "vaccination questionnaire" were created. Form I (probabilistic protection) described a disease expected to afflict 20% of the population and asked people whether they would volunteer to receive a vaccine

that protects half of the people receiving it. According to Form II (pseudocertainty), there were two mutually exclusive and equiprobable strains of the disease, each likely to afflict 10% of the population; the vaccination was said to give complete protection against one strain and no protection against the other.

The participants in this study were college students. Half received Form I; the other half received Form II. After reading the description, they rated the likelihood that they would get vaccinated in such a situation. Although both forms indicated that vaccination reduced one's overall risk from 20% to 10%, we expected that vaccination would appear more attractive to those who received Form II (pseudocertainty) than to those who received Form I (probabilistic protection). The results confirmed this prediction: 57% of those who received Form II indicated they would get vaccinated compared with 40% of those who received Form I.

The pseudocertainty effect highlights the contrast between the reduction and the elimination of risk. As Tversky and Kahneman have indicated, this distinction is difficult to justify on any normative grounds. Moreover, manipulations of certainty would seem to have important implications for the design and description of other forms of protection (e.g., medical treatments, insurance, flood- and earthquake-proofing activities).

Another category of framing effects comes from the way in which one is to respond to a decision problem. Although people are sometimes free to choose their response mode, typically some external source defines the problem as involving either judgment (of individual options) or choice (selecting one

from two or more options). Many theories of decision making postulate an equivalence between judgment and choice, assuming that each option, X , has a value, $v(x)$, which determines its attractiveness in both contexts (see e.g., Luce, 1977). However, the descriptive validity of these theories is now in question. Recent research has demonstrated that the information-processing strategies used in making choices are often quite different from the strategies employed in judging single options. In particular, much of the thinking prior to choice appears to be aimed at constructing a concise, coherent set of reasons that justify the selection of one option over the others (Tversky, 1972; Slovic, 1975). Judgment of single options are based either on different justifications or on a variety of non-justificatory processes. As a result, choices and evaluative judgments of the same options often differ, sometimes dramatically.

An example of the differences that can occur between evaluation and choice comes from two experiments (Lichtenstein & Slovic, 1971, 1973), one of which was conducted on the floor of the Four Queens Casino in Las Vegas. Consider the following pair of gambles used in the Las Vegas experiment:

Bet A: 11/12 chance to win 12 chips
 1/12 chance to lose 24 chips
 Bet B: 2/12 chance to win 79 chips
 10/12 chance to lose 5 chips

where the value of each chip has been previously fixed at 25 cents. Notice that bet A has a much better chance of winning, but bet B offers a higher winning payoff. Subjects indicated, in two ways, the attractiveness of each bet in many such pairs. First they made a simple choice, A or B. Later, they

were asked to assume that they owned a ticket to play each bet, and they were to state the lowest price for which they would sell this ticket.

Presumably, these selling prices and choices are both governed by the same underlying quality, the attractiveness of each gamble. Therefore, subjects should have stated higher selling prices for the gambles that they preferred in the choice situation. In fact, subjects often chose one gamble, yet stated a higher selling price for the other. For the particular pair of gambles shown above, A and B were chosen about equally often. However, bet B received a higher selling price about 88% of the time. Of the subjects who chose bet A, 87% gave a higher selling price to bet B, thus exhibiting an inconsistent preference pattern. Grether and Plott (1979), two skeptical economists, replicated this study with numerous variations designed to show that the observed inconsistencies were artifactual. They obtained essentially the same results as Lichtenstein and Slovic.

What accounts for this inconsistent pattern of preferences for gambles? Lichtenstein and Slovic concluded that subjects used different cognitive strategies when setting prices and making choices. Subjects often justified the choice of Bet A in terms of its good odds, but they set a higher price for B because they were greatly influenced by its large winning payoff. For example, people who found a gamble basically attractive used the amount to win as a starting point. They then adjusted the amount to win downward to accommodate the less-than-perfect chance of winning and the fact that there was some amount to lose as well. Typically, this adjustment was small and, as a result, large winning payoffs caused people to set prices that were inconsistent with their choices.

Another example of the interaction between framing, response mode, and justificatory processes comes from a study in which we presented college students and members of the League of Women Voters with two different but logically related tasks. Task 1 was a variation of the civil defense problem shown in Figure 2. The cover story was the same but the options read as follows:

Option A: carries with it a .5 probability of containing the threat with a loss of 5 lives and a .5 probability of losing 95 lives. It is like taking the gamble:

.5 lose 5 lives

.5 lose 95 lives

Option B: carries with it a .5 probability of containing the threat with a loss of 40 lives and a .5 probability of losing 60 lives. It is like taking the gamble:

.5 lose 40 lives

.5 lose 60 lives

Subjects were asked to select one of these two options. The second task asked them to rate their agreement with each of three functions representing the way that society should evaluate lives in multi-fatality situations (see Figure 3). The instructions for the second task (omitted in Figure 3) provided elaborate rationales for adopting each of the functional forms over a range of between 0 and 100 lives lost in a single accident. Briefly, the linear form (curve 1) represents the view that every life lost is equally costly to society. The exponentially increasing function (curve 2) represents the view that large losses of life are disproportionately serious (e.g., loss of 20

lives is more than twice as bad as loss of 10 lives). Curve 3 represents a reduced sensitivity to large losses of life (e.g., loss of 20 lives is less than twice as bad as loss of 10 lives). After studying each curve and its rationale, subjects were asked to indicate which they agreed with.

Insert Figure 3 about here

More than half of all subjects chose option A in Task 1 and agreed most with curve 2 in the second task (Figure 3). However, option A indicates a risk-seeking attitude toward loss of life, whereas curve 2 represents risk aversion. Choice of option A would be consistent with curve 3, which was the least favored curve. These inconsistent results were not changed appreciably by changing the degree of elaboration in the rationales given for the three curves.

Subjects who were confronted with the inconsistency in their responses refused to change. They claimed to see no connection between the two tasks. Most appeared to be relying on some variant of the following justification for choosing option A: "It would be immoral to allow the loss of 40 or more lives when option A presents a good chance of coming out of the situation with only a few lives lost." This perspective was evoked by the structure of the choice problem but not by the task of evaluating the three functional relationships.

Implications for Regulation

Implications for policy. Bounded rationality, in the form of difficulties in probabilistic thinking and risk perception, ill-defined and labile

preferences, information-processing limitations, and the presence of potent and easily manipulable decision frames, provides a rather startling contrast to the presumption of rationality upon which much economic and regulatory theory is based.

Although bounded rationality undoubtedly characterizes much individual decision making, its implications for regulatory policy remain unclear. At one extreme, some regulators have used the psychological evidence on judgmental limitations to argue that "public intervention [is] necessary for the protection of working men and women in this country" (Bingham, 1979). At the other extreme are those who argue that market decisions will be more resistant than regulatory decisions to cognitive biases because in the marketplace rational agents will survive at the expense of others. Scientific evidence regarding the prevalence and impact of cognitive biases in market decisions is lacking. Contradictory notions abound. Hogarth (1981) hypothesizes that dynamic, continuous, interactive environments (such as the marketplace) will reduce the influences of biases, which have typically been observed in discrete, static settings. On the other hand, Mackay (1982) provides rather dramatic anecdotal evidence demonstrating that financial markets may actually compound and exaggerate the biases of individual investors¹; more recent evidence of this is provided by Dreman (1977, 1980). Clearly, research to clarify this issue should be given high priority. Epstein (1979) poses the challenge well in the regulatory context:

The cognitive biases so well identified need not exhibit themselves at a constant rate across environments. The question therefore arises: what type

of institutional arrangements make it more likely that the individual decision makers will be able to overcome their own personal limitations and what type does not?

It is a prime order of business to organize our social institutions to facilitate decisions that are free from the pervasive influence of cognitive bias.

Many important questions remain unanswered at this time. However, we know enough, we believe, to argue that approaches to risk management should start with the presumption of bounded rationality and investigate the extent to which various market and regulatory mechanisms can overcome the limitations of individual minds.

Implications for assessing values. Decision problems with high stakes tend to be unique and unfamiliar. They draw us into situations in which we have not adequately thought through the implications of values and beliefs acquired in simpler, more familiar settings. Yet, at the same time, many have argued that values must be made explicit and incorporated into regulatory decision making (see, e.g., Kennedy, 1981). But how should this be done? Some call for direct elicitation of values through surveys, hearings, and the like, whereas others prefer to infer values from the preferences "revealed" in ongoing decisions. Both approaches assume that people know their own values and that elicitation methods are unbiased channels that translate subjective feelings into analytically usable expressions.

These assumptions may not always be valid. The strong effects of framing and information-processing considerations, acting upon inchoate preferences, can make elicitation procedures major forces in shaping the expression of

values (Fischhoff, Slovic & Lichtenstein, 1980). In such cases, the method becomes the message. Subtle aspects of how problems are posed, questions are phrased, and responses are elicited can have substantial impact on judgments that supposedly express the precise nature of people's preferences. Nevertheless, we are not wholly pessimistic on this score. Research to be described later in this chapter indicates that psychometric survey methods can identify the general categories of issues and outcomes that people deem important in the evaluation and management of risk. The macrostructure of preference may be better defined, more stable, and more readily measurable than the microstructure.

Implications for informing people about risk. Many theorists have argued that the most important function of a regulatory agency is to ensure that workers, patients, and citizens are properly informed about the risks they face (see, e.g., Joskow & Noll, 1980). Such information can presumably help people make better decisions in the marketplace and thus obviate the need for severe regulatory interventions. In recent years:

- The Food and Drug Administration mandated patient information inserts for an increased number of prescription drugs.

- The Department of Housing and Urban Development began to require the sellers of homes built before 1950 to inform buyers about the presence of lead-based paints.

- A proposed federal products liability law placed increased weight on adequately informing consumers and workers about risks they are likely to encounter.

- The White House directed the then-Secretary of Health, Education, and Welfare to develop a public information program on the health effects of radiation exposure.

Clearly, provision of information about risk is crucial to making better personal decisions and to facilitating the market and regulatory processes whereby societal standards are developed and enforced. However, despite good intentions, creating effective informational programs may be quite difficult (Slovic, Fischhoff & Lichtenstein, 1980b; 1981). Doing an adequate job means finding cogent ways of presenting complex, technical material that is clouded by uncertainty and may be distorted by the listener's preconceptions (and perhaps misconceptions) about the hazard and its consequences. Moreover, as we have seen, people are often at the mercy of the way problems are formulated. Those responsible for determining the content and format of information programs thus have considerable ability to manipulate perceptions. Indeed, since these effects are not widely known, people may inadvertently be manipulating their own perceptions by casual decisions they make about how to organize their knowledge.

The stakes in risk problems are high--product viability, jobs, energy costs, willingness of patients to accept treatments, public safety and health, etc. Potential conflicts of interest abound. When subtle aspects of how (or what) information is presented make a significant difference in people's responses, one needs to determine which formulation should be used. Making that decision takes one out of psychology and into the domains of law, ethics, and politics.

We have been emphasizing here the difficulties people have in comprehending and estimating risks. Some observers, cognizant of these difficulties, have concluded that the problems are insurmountable. We disagree. Although the broad outlines of the psychological research just described seem to support a pessimistic view, the details of that research give some cause for optimism. Upon closer examination, it appears that people understand some things quite well, although their path to knowledge may be quite different from that of the technical experts. In situations where misunderstanding is rampant, people's errors can often be traced to inadequate information and biased experiences, which education may be able to counter.

Studies of Protective Behavior

In this section we shall describe studies of two kinds of protective behavior, insurance and the use of seat belts. This research was designed to provide basic knowledge that would also have relevance for regulatory decisions.

National Flood Insurance Program

There has been much governmental concern over the fact that, whereas few residents of flood and earthquake areas voluntarily insure themselves against the consequences of such disasters, many turn to the federal government for aid after suffering losses (Kunreuther et al., 1978). Policy makers have argued that both the government and the property owners at risk would be better off financially under a federal insurance program. Such a program would shift the burden of disasters from the general taxpayer to individuals living in hazardous areas and would thus promote wiser decisions regarding the use of flood plains.

Without a firm understanding of how people perceive and react to risks, however, there is no way of knowing what sort of disaster insurance program would be most effective. For example, it seems reasonable to expect that lowering the cost of such insurance would stimulate people to buy it, yet there is evidence that people do not voluntarily purchase flood insurance even when the rates are highly subsidized.

Research on this topic by Kunreuther et al. (1978) and Slovic et al. (1977) was designed to determine the critical factors influencing the voluntary purchase of insurance against the consequences of low-probability events such as floods or earthquakes. Both laboratory experiments and field surveys were used. Analysis of the survey data revealed widespread ignorance and misinformation regarding the availability and terms of insurance and the probabilities of damage from a future disaster. The laboratory experiments showed that people preferred to insure against relatively high-probability, low-loss hazards and tended to reject insurance in situations where the probability of loss was low and the potential losses were high. These results suggest that people's natural predispositions run counter to economic theory (e.g., Friedman & Savage, 1948), which assumes that risk-averse individuals should desire a mechanism to protect them from rare catastrophic losses.

When asked about their insurance decisions, subjects in both the laboratory and survey studies indicated a disinclination to worry about low-probability hazards. Such a strategy is understandable; limitations on people's time, energy, and attentional capacities create a finite reservoir of concern. Unless people ignored many low-probability threats they would become

so burdened that productive life would become impossible. Another insight gleaned from the experiments and the survey is that people think of insurance as an investment. Making claims and receiving payments (by insuring against more probable losses) seems to be viewed as a return on the premium, hence a good investment. The popularity of low-deductible insurance plans (Fuchs, 1977; Pashigian, Schkade & Menefee, 1966) provides confirmation from outside the laboratory that people prefer to insure against probable events with small consequences.

One surprising result was that homeowners' lack of interest in disaster insurance did not seem to be due to expectations that the federal government would bail them out in an emergency. The majority of individuals interviewed said they anticipated no aid at all from the government in the event of a disaster. Most appeared not to have considered at all how they would recover from flood or earthquake damage.

This research led us to conclude that the primary cause of failure in the disaster insurance market is lack of consumer interest. If insurance is to be marketed on a voluntary basis, consumers' attitudes and information-processing limitations must be taken into account. Policy makers and insurance providers must find ways to communicate the risks and arouse concern for the hazards. One method found to work in the laboratory experiments is to increase the perceived probability of disaster by lengthening the individual's time horizon. For example, considering the risk of experiencing a 100-year flood at least once during a 25-year period, instead of considering the risk in one year, raises the probability from .01 to .22 and may thus cast flood insurance

in a more favorable light. Another step would have insurance agents play an active role in educating homeowners about the proper use of insurance as a protective mechanism and providing information about the availability of insurance, rate schedules, deductible values, etc. Of course, such actions may not be effective. It may also be necessary to institute some form of mandatory coverage. Recognizing the difficulty of inducing voluntary coverage, the National Flood Insurance Program now requires insurance as a condition for obtaining federal money to build in flood-prone areas.

Seat Belts

Another form of protection that people do not often use is the automobile seat belt. Promotional efforts to get motorists to wear seat belts have failed dismally (Robertson, 1976). In the wake of expensive advertising campaigns and buzzer systems, fewer than 20% of all motorists "buckle up for safety." Policy makers have criticized the public for failing to appreciate the risks of driving and the benefits of seat belts. However, results from risk perception research provide an alternative perspective that seems at once more respectful of drivers' reasoning and more likely to increase seat belt use. As noted above, people often disregard very small probabilities. By like token, motorists' reluctance to wear seat belts might be due to the extremely small probability of incurring a fatal accident on a single automobile trip. Because a fatal accident occurs only about once in every 3.5 million person-trips and a disabling injury only once in about every 100,000 person-trips, refusing to buckle up one's seat belt may seem quite reasonable. It may look less reasonable, however, if one frames the problem within a

multiple-trip perspective. This is, of course, the perspective of traffic safety planners, who see the thousands of lives that might be saved annually if everybody buckled up on every trip. For the individual driver, during 50 years of driving (about 40,000 trips), the probability of being killed is .01 and the probability of experiencing at least one disabling injury is .33. In laboratory experiments, we found that people induced to consider this lifetime perspective responded more favorably toward the use of seat belts (and air bags) than did people asked to consider a trip-by-trip perspective (Slovic, Fischhoff & Lichtenstein, 1978). More recent studies suggest that television and radio messages based on this lifetime-cumulative-risk theme will effectively increase actual seat belt use (Schwalm & Slovic, 1982).

Characterizing Perceived Risk

If it is to aid hazard management, a theory of perceived risk must explain people's extreme aversion to some hazards, their indifference to others, and the discrepancies between these reactions and experts' recommendations. Why, for example, do some communities react vigorously against the location of a liquid natural gas terminal in their vicinity, despite the assurances of experts that it is safe? Why, on the other hand, do many communities situated on earthquake faults or below great dams show little concern for experts' warnings? Over the past few years researchers have been attempting to answer such questions as these by examining the opinions that people express when they are asked, in a variety of ways, to characterize and evaluate hazardous activities and technologies. The goals of this descriptive research are (a) to develop a taxonomy of risk characteristics that can be used to understand

and predict societal responses to hazards and (b) to develop methods for assessing public opinions about risk in a way that could be useful for policy decisions.

The Psychometric Paradigm

In recent years, psychometric scaling methods and multivariate analysis techniques have been used to produce quantitative representations of risk attitudes and perceptions (Brown & Green, 1980; Fischhoff et al., 1978; Green, 1980a, b; Green & Brown, 1980; Renn, 1981; Slovic, Fischhoff & Lichtenstein, 1979, 1980a, b). Researchers employing this psychometric paradigm have typically asked people to judge the current and desired riskiness (or safety) of diverse sets of hazardous activities, substances, and technologies, and to indicate their desires for risk reduction and regulation of these hazards. These global judgments are then related to judgments about other properties, including:

- the hazard's status on characteristics that have been hypothesized to account for risk perceptions and attitudes (e.g., voluntariness, dread, knowledge, controllability),
- the benefits that each hazard provides to society,
- the number of deaths caused by the hazard in an average year,
- the number of deaths caused by the hazard in a disastrous year, and
- the seriousness of each death from a particular hazard relative to a death due to other causes.

A number of systematic, replicable, and potentially important results have emerged from these studies. An overview of the conclusions of our own work in

this area is as follows:

(a) Perceived risk is quantifiable and predictable. Psychometric techniques seem well suited for identifying similarities and differences among groups with regard to risk perceptions and attitudes (see, for example, Table 2).

Insert Table 2 about here

(b) "Risk" means different things to different people. When experts judge risk, their responses correlate highly with technical estimates of annual fatalities. Laypeople can assess annual fatalities if they are asked to (and produce estimates not unlike the technical estimates). However, their judgments of risk are sensitive to other factors as well (e.g., catastrophic potential, threat to future generations) and, as a result, are not closely related to their own (or experts') estimates of annual fatalities (see Figures 4 and 5).

Insert Figures 4 and 5 about here

(c) Even though groups disagree about the overall riskiness of specific hazards, they show remarkable agreement when rating those hazards on characteristics of risk such as knowledge, controllability, dread, catastrophic potential, etc.

(d) Many of these risk characteristics are highly correlated with each other, across a wide domain of hazards. For example, voluntary hazards tend also to be controllable and well known, hazards that threaten future

generations tend also to be seen as having catastrophic potential, etc. Analysis of these interrelationships shows that the broader domain of characteristics can be condensed to three higher-order characteristics or factors. These factors reflect the degree to which a risk is understood, the degree to which it evokes a feeling of dread, and the number of people exposed to the risk (see Figure 6). This factor structure has been found to be similar across groups of laypersons and experts judging large and diverse sets of hazards. Making the set of hazards more specific (e.g., partitioning nuclear power into radioactive waste transport, uranium mining, nuclear reactor accidents, etc.) appears to have little effect on the factor structure or its relationship to risk perceptions (Slovic, Fischhoff & Lichtenstein, in press).²

 Insert Figure 6 about here

(e) Many of the various characteristics, particularly those associated with the factor "Dread Risk," correlate highly with laypersons' perceptions of risk. The higher an activity's score on the dread factor, the higher its perceived risk, the more people want its risks reduced, and the more they want to see strict regulation employed to achieve the desired reductions in risk (see Figure 7). The factor labeled "Unknown Risk" tends not to correlate highly with risk perception. Factor 3, Exposure, is moderately related to laypersons' perceptions of risk. In contrast, experts' perceptions of risk are not related to any of the various risk characteristics or factors derived from these characteristics.

Insert Figure 7 about here

(f) In agreement with hypotheses originally put forth by Starr (1969), people's tolerance for risk appears related to their perception of benefit. All other things being equal, the greater the perceived benefit, the greater the tolerance for risk. Moreover, that tolerance depends upon the voluntariness of the activity. In addition, we have found that risk acceptability is also influenced by other characteristics such as familiarity, control, catastrophic potential, and uncertainty about the level of risk.

Implications of Risk Perception Research

In the results of the psychometric studies, we have the beginnings of a perceptual/psychological classification system for hazards. Ultimately we need not only a better psychological taxonomy but one that reflects physical, biological, and social/managerial elements of hazards, as well as psychological ones. Such a taxonomy would be a potent device for understanding and guiding social regulation of risk. We are far from this goal, though a start towards an expanded taxonomy has been made (Clark University and Decision Research, 1982). For the present, more modest insights and applications must suffice.

Forecasting public response. Despite the lack of a complete understanding of public attitudes and perceptions, we have attempted to use the results from risk perception studies to explain and forecast reactions to specific technologies. Nuclear power has been the principle object of such analysis because of its obvious social importance. As Alvin Weinberg (1976) observed, "...the public perception and acceptance of nuclear energy...has emerged as

the most critical question concerning the future of nuclear energy." The reasonableness of these perceptions has been the topic of an extensive public debate, filled with charges and countercharges. For example, one industry source has argued that public reaction to Three Mile Island has cost "...as much as \$500 billion...and is one measure of the price being paid as a consequence of fear arising out of an accident that according to the most thorough estimates may not have physiologically hurt even one member of the public" (EPRI Journal, 1980, p. 30).

Risk perception research offers some promise of clarifying the concerns of opponents of nuclear power (Slovic, Lichtenstein & Fischhoff, 1979). In particular, psychometric studies show that these people judge its benefits as quite low and its risks as unacceptably great. On the benefit side, most opponents do not see nuclear power as a vital link in meeting basic energy needs; rather, they view it as a supplement to other sources of energy which are themselves adequate. On the risk side, nuclear power occupies a unique position in the factor space, reflecting people's views that its risks are unknown, dread, uncontrollable, inequitable, catastrophic, and likely to affect future generations (see Figure 6). Opponents recognize that few people have died to date as a result of nuclear power. However, they do have great concern over the potential for catastrophic accidents. Further analyses by Slovic, Fischhoff and Lichtenstein (1981b) have suggested that opposition to nuclear power can be understood in terms of basic psychological principles of perception and cognition and is not likely to be changed by information campaigns that focus on safety; however, information about benefits may have

some impact. Opposition might well ease if the industry maintains a superb safety record or energy shortages occur. But because nuclear risks are perceived to be unknown and potentially catastrophic, even small accidents will have immense social costs, a fact that, as discussed later, has direct implications for the setting of safety standards (Slovic, Fischhoff & Lichtenstein, 1980c).

This type of research may also forecast the response to technologies that have yet to catch the public's eye. For example, our studies indicate that recombinant DNA technology shares several of the characteristics that make nuclear power so hard to manage (Slovic, Fischhoff & Lichtenstein, in press). If it somehow seizes public attention, this new technology could face some of the same problems and opposition now confronting the nuclear industry.

Comparing risks. One frequently advocated approach to deepening people's perspectives is to present quantified risk estimates for a variety of hazards, expressed in some unidimensional index of death or disability, such as risk per hour of exposure (Sowby, 1965), annual probability of death (Wilson, 1979), or reduction in life expectancy (Cohen & Lee, 1979; Reissland & Harries, 1979). Even though such comparisons have no logically necessary implications as guides to decision making (Fischhoff, Lichtenstein, Slovic, Derby & Keeney, 1981), one might still hope that they would help improve people's intuitions about the magnitude of risks. Risk perception research suggests, however, that these comparisons will not be very satisfactory. People's perceptions and attitudes are determined not only by the sort of unidimensional statistics used in such tables but also by a variety of

quantitative and qualitative characteristics--including a hazard's degree of controllability, the dread it evokes, its catastrophic potential, and the equity of its distribution of risks and benefits. To many people, statements such as "the annual risk from living near a nuclear power plant is equivalent to the risk of riding an extra three miles in an automobile" give inadequate consideration to the important differences in the nature of the risks from these two technologies. In short, "riskiness" means more to people than "expected number of fatalities." Attempts to characterize, compare, and regulate risks must be sensitive to the broader conception of risk that underlies people's concerns.

Deciding "How Safe is Safe Enough?"

The bottom line in risk regulation is usually an answer to some variant of the question, "How safe is safe enough?" The question takes such forms as: "Do we need additional containment shells around our nuclear power plants?" "Is the carcinogenicity of saccharin sufficiently low to allow its use?" "Should schools with asbestos ceilings be closed?" Lack of adequate answers to such questions has bedeviled hazard management.

Of late, many hazard management decisions are simply not being made--in part because of vague legislative mandates and cumbersome legal proceedings, in part because there are no clear criteria for making such decisions. The decisions that are made are often inconsistent. Our legal statutes are less tolerant of carcinogens in the food we eat than of those in the water we drink or in the air we breathe. In the United Kingdom, 2,500 times as much money per life saved is spent on safety measures in the pharmaceutical industry as in agriculture (Sinclair, Marstrand & Newick, 1972). U.S. society is

apparently willing to spend about \$140,000 in highway construction to save one life and \$5 million to save a person from death due to radiation exposure (Howard, Matheson & Owen, 1978).

Frustration over this state of affairs has led to a search for clear, implementable rules that will tell us whether or not a given technology is sufficiently safe. Among the many seeking to solve the mystery of acceptable risk have been Bazelon (1979), Council for Science and Society (1977), Crouch and Wilson (1982), Comar (1979), Griesemeyer and Okrent (1980), Howard, Matheson and Owen (1978), Lave (1981), Lowrance (1976), Okrent and Whipple (1981), Rowe (1977), Salem, Solomon and Yesley (1980), Schwing and Albers (1980), Starr (1969), and ourselves (Fischhoff, Lichtenstein, Slovic, Derby & Keeney, 1981). Despite heroic effort, no magic formula has been discovered. Nevertheless, some progress has been made, not the least of which includes the development of a heightened respect for the complexities of the task.

Approaches to Acceptable Risk: A Critique

Our own efforts in this area during the past three years have been instigated and supported by the Nuclear Regulatory Commission (NRC). Nowhere are the problems of determining acceptable levels of risk more important or more complex than in the regulation of nuclear power. It has always been known that nuclear reactors could be made safer—at increased cost. But since quantifying safety was a problem, the question of how much safety at what price was rarely addressed directly. Beginning with the Reactor Safety Study (U.S. NRC, 1975), the technology of measuring risk has advanced rapidly in recent years. Now that quantitative estimates of safety are thought to be

accessible for many aspects of the nuclear fuel cycle, the need to determine how safe reactors should be has taken on greater significance.

At the urging of Congress, the nuclear industry, and the Advisory Committee on Reactor Safeguards, the NRC has been working intensively to develop an explicit, possibly quantitative, safety goal or philosophy. Presumably this goal would clarify the Commission's current vague mandate to "avoid undue risk to public health and safety" and, by doing so, would serve to guide specific regulatory decisions.

We were asked, at an early stage in this endeavor, to take a comprehensive, critical look at the philosophical, sociopolitical, institutional, and methodological issues crucial to determining how best to answer the question of "How safe is safe enough?". Limitations of time, resources, and our own capabilities forced us to narrow the scope of this endeavor somewhat, although we did structure the task in a general way, not restricted to nuclear power or any other specific technology. Our efforts, reported in Fischhoff et al. (1981), are summarized below.

Guided by our own psychological perspective and aided by decision theorists Stephen Derby and Ralph Keeney, our examination of approaches to acceptable risk attempted to:

- (a) Clarify the nature of acceptable-risk problems and examine some frequently proposed, but not entirely adequate, solutions.
- (b) Characterize the essential features of acceptable-risk problems that make their resolution so difficult. These features included uncertainty about how to define acceptable risk problems, difficulties in obtaining crucial

facts, difficulties in assessing social values, unpredictable human responses to hazards, and problems of assessing the adequacy of decision-making processes.

(c) Create a taxonomy of decision-making methods, identified by how they attempt to address the essential features of acceptable-risk problems. The major approaches we discussed were professional judgment: allowing technical experts to devise solutions; bootstrapping: searching for historical precedents to guide future decisions; and formal analysis: theory-based procedures for modeling problems and calculating the best decision, such as risk/benefit, cost/benefit, and decision analysis.

(d) Specify the objectives that an approach should satisfy in order to guide social policy. These included comprehensiveness, logical soundness, practicality, openness to evaluation, political acceptability, institutional compatibility, and conduciveness to learning.

(e) Evaluate the success of the approaches in meeting these objectives.

(f) Derive recommendations for policy makers and citizens interested in improving the quality of acceptable-risk decisions.

The following conclusions emerged from our analysis:

1. Acceptable-risk problems are decision problems, that is, they require a choice among alternatives. That choice is dependent on the set of options, consequences, values, and facts invoked in the decision process. Therefore, there can be no single, all-purpose number that expresses the acceptable risk for a society. At best, one can hope to find the most acceptable alternative in a specific problem, one that will represent the values of a specific constituency.

2. None of the approaches we considered is either comprehensive or infallible. Each gives special attention to some features of acceptable-risk problems but ignores others. As a result, not only does each approach fail to give a definitive answer, but it is biased toward particular interests and particular solutions. Hence choosing an approach is a political act that carries a distinct message about who should rule and what should matter. The search for an objective method is doomed to failure and may obscure the value-laden assumptions that will inevitably be made.

3. Acceptable-risk debates are greatly clarified when the participants are committed to separating issues of fact from issues of value. Nonetheless, a clearcut separation is often impossible. Beliefs about the facts of the matter shape our values; those values in turn shape the facts we search for and how we interpret what we find.

4. The determining factor in many acceptable-risk decisions is how the problem is defined (i.e., which options and consequences are considered, what kinds of uncertainty are acknowledged, and how key terms are operationalized). Until definitional disputes are settled, it may be impossible to reach agreement regarding what course of action to take.

5. Values, like beliefs, are acquired through experience and contemplation. Acceptable-risk problems raise many complex, novel, and vague issues of value for which individuals may not have well-articulated preferences. In such situations, the values one expresses may be greatly influenced by transient factors, including subtle aspects of how value questions are posed by interviewers, politicians, or the marketplace. The conflicts in each individual's values are above and beyond the conflicts between different individuals' values. Both types of conflict require careful attention.

6. Although a distinction is often made between perceived and objective risks, for most new and intricate hazards even so-called objective risks have a large judgmental component. At best, they represent the perceptions of the most knowledgeable technical experts. However, even such experts may have an incomplete understanding. Indeed, their professional training may have limited them to certain traditional ways of looking at problems. In such cases, nonexperts may have important supplementary information or viewpoints on hazards and their consequences.

We concluded that no one solution to acceptable-risk problems is now available, nor is it likely that a single solution will ever be found. Nonetheless, the following recommendations, proposed to regulators, citizens, legislators, and professionals and designed to enhance society's ability to make decisions, were put forth:

1. Explicitly recognize the complexities of acceptable-risk problems. The value judgments and uncertainties encountered in specific decision problems should be acknowledged. More generally, we should realize that there are no easy solutions and not expect society's decision makers to come up with them.

2. Acknowledge the limits of currently available approaches and expertise. As no approach is infallible, we should at least avoid the more common mistakes. Our aim should be a diverse and flexible approach to decision making that emphasizes comprehensiveness.

3. Improve the use of available approaches. Develop guidelines for their conduct and review. Make them sensitive to all aspects of the problem and to the desires of as many stakeholders as possible. Analysis should proceed iteratively in order to sustain its insights and absorb its criticisms.

4. Make the decision-making process compatible with existing democratic institutions. The public and its representatives should be involved constructively in the process both to make it more effective and to increase the public's understanding of hazard issues.

5. Strengthen nongovernmental social mechanisms that regulate hazards. Decisions reached in the marketplace and political arena provide important guidelines for most approaches. Those mechanisms can be improved by various measures including reform of the product-liability system and increased communication of risk information to workers and consumers.

6. Clarify government involvement. Legislation should offer clear, feasible, and predictable mandates for regulatory agencies. The management of different hazards should be coordinated so as to build a legacy of dependable precedents and encourage consistent decisions.

Toward a Safety Goal

Justification. Our analysis of decision-making approaches was used by the NRC in the planning stages of its program to develop a safety goal (U.S. NRC, 1981). Upon completion of this analysis, we were asked to participate in the development of the goal itself. Before doing so, we felt it necessary to critique the effort in light of our earlier conclusion that, since acceptable risk is the outcome of specific decisions, there can be no single, all purpose number (standard or goal) that does the job. Whereas decision-making procedures attempt to order options according to some criterion of attractiveness, goals and the standards that are derived from goals essentially categorize options in terms of pass/no pass. Beyond the obvious

efficiency of setting a generally applicable decision rule, are there any other justifications for goals and standards? Fischhoff (1982) wrestled with this question and concluded that there were, indeed, circumstances in which standards were warranted. Any one of the following conditions might justify the development of a pass/no pass safety standard:

1. When managerial resources are severely limited.
2. When one need not choose a single best option.
3. When a single (standardizable) feature captures the most important aspect of a category.
4. When the standard accurately postdicts past decisions and predicts future ones.
5. When one wants to make a statement to reflect the goals of policy makers (who assume the symbolic standard will be reasonably compromised by those who apply it).
6. When one hopes to shape the set of future options.
7. When the decision process leading to the standard is of higher quality than could be maintained in numerous specific decisions.

In addition to providing a rationale for goals and standards, Fischhoff (1982) explored the many subtle and complex problems involved in transforming a goal from a political statement to a useful tool, one that can be unambiguously applied by regulators and understood by the regulated. Here one faces issues such as (a) defining the category governed by the standard (e.g., Is a cosmetic a drug?); (b) determining the point and time of regulation (e.g., Plant by plant or company by company? At which stage of production and use?); (c) tailoring standards to mesh with engineering and design.

capabilities; (d) deciding whether to regulate technical matters (nuts and bolts) or performance ("as long as you meet this goal, we don't care how you do it"). Once one has decided where to place the standard, the critical design question becomes how to measure the risks that arise in order to determine whether they are in compliance with the standard.

Technical issues. Having satisfied ourselves that general goals and standards had a place in the regulator's armamentarium (the NRC already knew this), we proceeded to consider the detailed process of establishing a safety goal. As we write this, a proposed goal is still in the discussion phase (U.S. Nuclear Regulatory Commission, 1982). Since this enterprise is still in progress, our discussion here pertains only to various preliminary proposals and will not necessarily apply to the goal eventually put forth by the NRC.

Our objective was to provide comments, criticisms, and suggestions, from our perspective as behavioral scientists and decision theorists, about what was seen as primarily a technical problem, dealing with the design, construction, and licensing of reactors and the ability of probabilistic techniques to assess and verify reactor risks. Before discussing our efforts on the social side, some of the technical approaches need to be described.

There has been no shortage of safety goal proposals over the years, beginning with Adams and Stone (1967) and Farmer (1967). Notable recent efforts include those of the Advisory Committee on Reactor Safeguards (U.S. NRC, 1980), the Atomic Industrial Forum (1981), Griesmeyer and Okrent (1981), Joksimovic and O'Donnell (1981), Kinchin (1978, 1979), and Zebroski (1980). Miller and Hall (1981) provide an annotated bibliography of efforts in this

field. Solomon, Nelson, and Salem (1981) surveyed these and other proposals and counted 103 safety criteria applicable to reactor accidents, which they categorized as follows:

1. Criteria for the safety of reactor systems, e.g., an upper limit for the acceptable probability of a core-melt accident.
2. Criteria for the allowable risks to individuals in the vicinity of the plant site.
3. Criteria for the allowable risks to individuals beyond the plant site. These include, for example, risks to populations resulting from accidental radiation releases.³
4. Qualitative criteria and criteria based on dollar damage.

A detailed discussion of these criteria is beyond the scope of this paper. Suffice it to say that (a) they differ from each other, sometimes by several orders of magnitude; (b) they tend to be derived on the basis of comparisons with other accident risks and with the risks from other sources of electricity, (c) they are concerned with a rather narrow view of the costs of a reactor accident, focusing on immediate and latent fatalities, physical damage to the reactor and adjoining property, and costs of cleanup and replacement electricity, and (d) they sometimes incorporate a weighting factor that attributes extra significance and cost to accidents that cause multiple fatalities. This last feature is known as "risk aversion."

Typical criteria are those recommended by the Atomic Industrial Forum (1981). They proposed a goal of $p < 10^{-5}$ /year mortality risk to a maximally exposed individual in close proximity to a reactor site as representing a non-significant (0.1%) increase in annual mortality risk (assumed to be <

10^{-2} /year). This goal was justified by noting that it was equal to or less than the annual mortality rates from other causes of accidents such as motor vehicles, violence, fires, etc. The risk of electrocution is about the same magnitude as the proposed goal. The total risk from all non-nuclear accidents is about 60 times greater than the AIF goal.

For risk to the population at large, the AIF proposes a goal of less than one statistically estimated fatality per year per 1,000 megawatts of nuclear power capacity. Again, they point out that this represents a very small percentage (.01%) increase in existing levels of risk and compares favorably with the risk of electrocution (about 2 fatalities per 1,000 megawatt capacity). The AIF also proposed a goal of 10^{-4} per reactor year for the probability of a large scale fuel melt, wherein the reactor core is uncovered for a sufficient time to cause the bulk of the fuel to become molten.

A major technical issue is whether or not methods of probabilistic risk assessment are accurate enough to be used as the basis for establishing safety goals and verifying that they are being met. Many experts doubt that the low probabilities stated in the goals can be verified. Others assert that probabilistic risk assessment, employed competently and with margins of error built into the calculations, can be a valuable aid to regulatory decision making.

Social issues. The preceding section presents an incomplete and grossly oversimplified view of the technical issues and policy motivations underlying the development of safety goals. Nevertheless, the description should suffice to provide background for a consideration of the social issues involved.

The main objective of our efforts has been to highlight the importance of

certain social issues that tend to be neglected by the standard engineering approach to establishing safety goals. The assessment of accident probabilities, accident consequences, and the broader social impacts of accidents are all vital to the goal-setting process. Although much attention has been given to the problems in modeling the first two of these components, the uncertainties in assessing social values and incorporating them into safety goals are, we believe, even more serious and much less acknowledged.

Nevertheless, social values must be considered. In promulgating a safety goal, the Commission, implicitly or explicitly, takes a stand on many important social issues. Among the many questions that need to be addressed are the following:

- a. What social values should be considered in determining safety goals? How should they be integrated with technical considerations in the process of policy formation?
- b. Are current risk levels from other hazards or competing energy technologies meaningful benchmarks against which to set standards for nuclear power?
- c. Is risk aversion (special emphasis on avoiding large accidents) desirable?
- d. Should accident costs be defined narrowly (i.e., restricted to early and latent health effects and property damage) or broadly (e.g., including higher-order or ripple effects such as possible shutdown from the nuclear industry and the costs that could result from that)?
- e. What consideration should be given to public fears that experts see as unjustified?

Ideally, it would be desirable to address these and other value issues from the standpoint of all parties involved (e.g., the various publics, regulators, industry, government, etc.), considering not only what these people do want (as expressed verbally or as inferred from their present and past behavior), but what they should want from a normative or ethical standpoint. From such analyses, qualitative and quantitative goals would be derived. Unfortunately, the state of the art does not permit such precise analysis of social values. At best, then, the Commission can:

- give these social issues a prominent place on its agenda of issues relevant to determining adequate protection of the public's health and safety,
- attempt to derive qualitative and quantitative safety goals based on reasonable assumptions with regard to these issues,
- make these assumptions explicit,
- encourage scientists, politicians and the public to study these issues and to propose ways in which the safety goals could be made more responsive to social impacts, and
- provide a management plan for future evolution of the safety goals that incorporates mechanisms for their revision in light of improved understanding of social and technical issues.

Implications of risk perception research. Although we do not have complete answers to the value questions raised above, research on risk perception has something to say in response to each of them. Regarding the question of which social values should be considered, our studies have shown that people's concerns about nuclear power are strongly influenced by their perceptions of catastrophic potential and the inequitable distribution of risks and benefits (including the transfer of risks to future generations).

Safety goals need to take these concerns into account.

Regarding the question of the appropriateness of using benchmarks based on other hazards, psychological research suggests that comparisons with other risks of life or risks from competing energy sources should not be a primary factor in determining safety goals. Nuclear power risks are unique in their level of perceived uncertainty, potential for catastrophe, threat to future generations, potential for triggering social disruption, etc. The logic of comparing qualitatively different risks is not developed well enough so that comparisons with other hazardous activities or technologies can serve as definitive guidelines for safety goals. Even if it were possible to make qualitatively different risks commensurable, comparisons would be subject to Hume's dictum that "no ought can ever follow from an is." Prescriptive guidelines must reflect not just what a society does want (assuming that we can discern this from what it currently accepts from other technologies—a questionable assumption), but what it should want.

Our research has led us to have even stronger feelings about the question of incorporating risk aversion and the question of breadth of costs to be considered. Although psychometric studies and other surveys have pinpointed perceived catastrophic potential as a major public concern, further investigation indicates that the alpha model, the model most often proposed for incorporating risk aversion into safety goals, is incorrect. According to this model, the seriousness or social impact of a loss of N lives in a single accident should be modeled by the function N^α , where α is greater than 1.0. By attributing greater social disruption to large accidents, this model strongly influences the stringency of standards. It implies that small

accidents may be tolerable but that extra money and effort should be expended to prevent or mitigate large accidents.

Research indicates that the alpha model is oversimplified and invalid; the societal costs of an accident cannot be modeled by any simple function of N . Slovic, Fischhoff, and Lichtenstein (1980a) presented data suggesting that accidents are signals containing information about the probability of their reoccurrence in similar or more destructive forms. As a result, the perceived seriousness of an accident is often determined more by the message it conveys than by its actual toll of death and destruction. An accident that takes many lives may produce relatively little social disturbance and thus be judged as not terribly serious if it occurs as part of a familiar and well understood system (e.g., a train wreck). In contrast, a small accident in an unfamiliar system whose risks are judged to be unknown and potentially catastrophic may have immense consequences if it portends further, possibly greater mishaps; such an accident may thus be perceived as very serious.

The view that accidents are signals implies that any of the usual ways of modeling risk aversion, including the alpha model, will fail to capture properly the real concerns people have about multiple-death accidents. In addition, it suggests a broadening of accident costs.

The perception of nuclear power risks as poorly understood and potentially catastrophic implies that nuclear accidents will be seen as extremely informative and ominous signals. Thus, another core-damaging accident similar to Three Mile Island would likely raise fears that the technology is out of control, even if few lives were lost and the physical damage was contained. The major costs of such an accident would not be those from immediate loss of

life, latent cancers, and direct economic expenses (e.g., property damage, repairs, cleanup), important as these may be. Instead, the dominant costs would arise from secondary impacts such as public reaction leading to shut-down of the entire industry and the resulting higher order consequences (dependence on more costly and more dangerous energy sources, economic collapse), which could total tens of hundreds of billions of dollars. These broader sociopolitical costs must be considered when determining the acceptable probability of a core-damaging accident. In other words, the design of safety criteria might be phrased in terms of the question: "What probability of an accident costing tens or hundreds of billion dollars is tolerable?" Such high costs may be as likely to result from an accident with no deaths and little property damage as from a rarer form of accident that produces many deaths and much damage. This broader impact is a direct result of public perceptions.

The final question, concerning the consideration that should be given to public fears, seems an appropriate one with which to end this section and conclude this paper, since it is fundamental, not only to nuclear safety goals, but to risk regulation in general. It is a difficult question and one that empirical data are not sufficient to answer.

There are many reasons for laypeople and experts to disagree. These include misunderstanding, miscommunication, and misinformation (Fischhoff, Slovic & Lichtenstein, 1981). Discerning the causes underlying a particular disagreement requires a combination of careful thought, to clarify just what is being talked about and whether agreement is possible given the disputants' differing frames of reference, and careful research, to clarify just what it

is that the various parties know and believe. Once the situation has been clarified, the underlying problem can be diagnosed as calling for a scientific, educational, semantic, or political solution.

The most difficult situations will be those in which the participants cannot agree on what the problem is (and have no recourse to an institution that will resolve the question by arbitration or by fiat) and those in which education is called for, yet fails (after some reasonable, diligent effort). Policy makers then face the hard choice between going against their own better judgment by using the public's assessment of risk (in which they do not believe) or going against the public's felings by imposing policies that will be disliked. Such policies may seem overly cautious (e.g., motorcycle helmet laws--to some people) or insufficiently cautious (e.g., nuclear power--to some people). When fears are ignored, the result can be stress or psychosomatic effects, which can be as real in their impact as they are illusory in their source. When strong public opinions are ignored, the result can be hostility, mistrust, and alienation. Since a society does more than manage risks, the policy maker must consider whether the social benefits to be gained by optimizing the allocation of resources in a particular decision is greater than the social costs of overriding a concerned public. A pessimistic view on "going with the public" might argue that "it only encourages the forces of irrationality (indirectly giving credence to astrology, superstition, and the like)." An optimistic view might be that risk questions are going to be with us for a long time. For a society to deal with them wisely, it must learn about their subtleties, including how appearances can be deceiving. One way of learning is by trial and error. Often, the experts will be able to say "We

told you so. It would have been better to listen to us." In other cases, the experts may be surprised. Learning is possible as long as some basic respect remains between teacher and pupil. That respect may be one of society's greatest assets.

References

- Adams, C. A. & Stone, C. N. Safety and siting of nuclear power stations in the United Kingdom. IAEA Symposium on the containment and siting of nuclear power plants. Vienna, Austria, April 1967.
- Arrow, K. J. Risk perception in psychology and economics. Presidential address to the Western Economic Association, 1981; Technical Report No. 351; Center for Research on Organizational Efficiency, Stanford University, October 1981.
- Atomic Industrial Forum. A proposed approach to the establishment and use of quantitative safety goals in the nuclear regulatory process. May 1981.
- Bazelon, D. L. Risk and responsibility. Science, 1979, 205, 277-280.
- Bingham, E. Address presented at the hearings on Risk/Benefit and the legislative process. Committee on Science and Technology (U.S. House of Representatives) and Committee on Commerce, Science, and Transportation (U.S. Senate). July 24, 1979. Available from the U.S. Government Printing Office.
- Brown, R. A. & Green, C. H. Precepts of safety assessments. Journal of the Operational Research Society, 1980, 11, 563-571.
- Clark University and Decision Research. The nature of hazard. Decision Research Report 82-8. Eugene, Oregon, 1982.
- Cohen, B. & Lee, I. A catalog of risks. Health Physics, 1979, 36, 707-722.
- Comar, C. L. Risk: A pragmatic de minimus approach. Science, 1979, 203, 319.
- Combs, B. & Slovic, P. Newspaper coverage of causes of death. Journalism Quarterly, 1979, 56(4), 837-843.
- Council for Science and Society. The acceptability of risks. London: Barry Rose, 1977.
- Crouch, E. A. C. & Wilson, R. Risk analysis. Cambridge, Mass.: Ballinger, 1981.
- Dreman, D. Psychology and the stock market, New York: AMACOM, 1977.
- Dreman, D. Contrarian investment strategy. New York: Random House, 1979.
- EPRI Journal. Assessment: The impact and influence of TMI. Palo Alto, Calif.: Electric Power Research Institute, 1980, 5(5), 24-33.

- Epstein, R. A. Cognitive biases: The institutional framework.
Unpublished manuscript, 1979.
- Farmer, F. R. Siting criteria--A new approach. Containment and siting of nuclear plants. Vienna, Austria: International Atomic Energy Agency, 1967, Pp. 303-318.
- Fischhoff, B. Standard setting standards. Decision Research Report 82-9, Eugene, Oregon, 1982.
- Fischhoff, B., Goitein, B. & Shapira, Z. The experienced utility of expected utility approaches. In N. Feather (Ed.), Expectancy, incentive and action. Hillsdale, N. J.: Erlbaum, 1982.
- Fischhoff, B., Hohenemser, C., Kasperson, R. E. & Kates, R. W. Handling hazards: Can hazard management be improved? Environment, 1978, 20(7), 16-20; 32-37.
- Fischhoff, B., Lichtenstein, S., Slovic, P., Derby, S. & Keeney, R. L. Acceptable risk. New York: Cambridge University Press, 1981.
- 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, 330-344.
- 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.
- Fischhoff, B., Slovic, P. & Lichtenstein, S. Knowing what you want: Measuring labile values. In T. Wallsten (Ed.), Cognitive processes in choice and decision behavior. Hillsdale, N. J.: Erlbaum, 1980a.
- Fischhoff, B., Slovic, P. & Lichtenstein, S. Labile values: A challenge for risk assessment. In J. Conrad (Ed.), Society, technology, and risk assessment. London: Academic Press, 1980b.
- Fischhoff, B., Slovic, P. & Lichtenstein, S. "The public" vs. "The experts": Perceived vs. actual disagreements about the risks of nuclear power. Risk Analysis, in press.
- Friedman, M. & Savage, L. J. The utility analysis of choices involving risk. Journal of Political Economy, 1948, 56, 279-304.

- Fuchs, V. R. From Bismarck to Woodcock: The irrational pursuit of national health insurance. The Journal of Law and Public Policy, 1976, 19, 347-359.
- Griesemeyer, J. M. & Okrent, D. Risk management and decision rules for light water reactors. Safety goals for nuclear power plants. NUREG-0739. Washington, D. C.: U.S. Nuclear Regulatory Commission, Advisory Committee on Reactor Safeguards, 1980.
- Green, C. H. Risk: Attitudes and beliefs. In D. V. Canter (Ed.), Behaviour in fires. Chichester: Wiley, 1980a.
- Green, C. H. Revealed preference theory: Assumptions and presumptions. In J. Conrad (Ed.), Society, technology and risk assessment. London: Academic Press, 1980b.
- Green, C. H. & Brown, R. A. Through a glass darkly: Perceiving perceived risks to health and safety. Research paper, School of Architecture, Duncan of Jordanstone College of Art/University of Dundee, Scotland, 1980.
- Grether, D. M. & Plott, C. R. Economic theory of choice and the preference reversal phenomenon. American Economic Review, 1979, 69, 623-638.
- Hogarth, R. M. Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. Psychological Bulletin, 1981, 90, 197-217.
- Howard, R. A., Matheson, J. E. & Owen, D. The value of life and nuclear design. Proceedings of the topical meeting on probabilistic analysis of nuclear reactor safety, 1978, 2, IV.2-1 - IV.2-9.
- Hynes, M. & Vanmarcke, E. Reliability of embankment performance prediction. In Proceedings of the ASCE Engineering Mechanics Division Specialty Conference. Waterloo, Ontario, Canada: University of Waterloo Press, 1976.
- Joksimovic, V. & O'Donnell, L. F. Quantitative safety goals for the regulatory process. In Current nuclear power plant safety issues. Vienna, Austria: International Atomic Energy Agency, 1981.
- Joskow, P. & Noll, R. Theory and practice in public regulation: A current overview. In G. Fromm (Ed.), Studies in public regulation, Cambridge, Mass.: MIT, 1980.
- Kahneman, D. & Tversky, A. Prospect theory: An analysis of decisions under risk. Econometrica, 1979, 47, 262-291.
- Kahneman, D., Slovic, P. & Tversky, A. (Eds.), Judgment under uncertainty: Heuristics and biases. New York: Cambridge University Press, 1982.

- Kennedy, D. The politics of preventive health. Technology Review, 1981, 84(2), 58-60.
- Kinchin, G. H. Assessment of hazards in engineering work. Proceedings of the Institute of Civil Engineers, 1978, 64, Pt. 1, 431-438.
- Kinchin, G. H. Design criteria, concepts and features important to safety and licensing. Proceedings of the international meeting on fast reactor safety technology, Seattle, Wash., 1979.
- Knight, F. H. Risk, uncertainty, and profit. Boston: Houghton-Mifflin, 1921.
- Kunreuther, H., Ginsberg, R., Miller, L., Sagi, P., Slovic, P., Borkan, B. & Katz, N. Disaster insurance protection: Public policy lessons. New York: Wiley, 1978.
- Lave, L. The strategy of social regulation. Washington, D. C.: The Brookings Institution, 1981.
- Lichtenstein, S. & Slovic, P. Reversals of preference between bids and choices in gambling decisions. Journal of Experimental Psychology, 1971, 89, 46-55.
- 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., Fischhoff, B. & Phillips, L. D. Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), Judgment under uncertainty: Heuristics and biases. New York: Cambridge University Press, 1982.
- 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.
- Lowrance, W. W. Of acceptable risk. Los Altos, Calif.: Kaufmann, 1976.
- Luce, R. D. The choice axiom after twenty years. Journal of Mathematical Psychology, 1977, 15, 215-233.
- Mckay, C. Extraordinary popular delusions and the madness of crowds. London: Office of the National Illustrated Library, 1852. Reprinted by the Noonday Press, New York, 1970.
- March, J. G. Bounded rationality, ambiguity, and the engineering of choice. The Bell Journal of Economics, 1978, 9, 587-608.

- Miller, B. & Hall, R. E. Index of risk exposure and risk acceptance criteria. NUREG/CR-1930, BNL-NUREG-51339. Washington, D. C.: U.S. Nuclear Regulatory Commission, January 1981.
- Nisbett, R. & Ross, L. Human inference: Strategies and shortcomings of social judgment. Englewood Cliffs, N. J.: Prentice Hall, 1980.
- Okrent, D. & Whipple, C. An approach to societal risk acceptance criteria and risk management. Report UCLA-ENG-7746. Los Angeles, Calif.: School of Engineering and Applied Science, University of California, June 1977.
- Pashigian, B. P., Schkade, L. & Menefee, G. H. The selection of an optimal deductible for a given insurance policy. Journal of Business, 1966, 39, 35-44.
- Reissland, J. & Harries, V. A scale for measuring risks. New Scientist, 1979, 83, 809-811.
- Renn, O. Man, technology, and risk: A study on intuitive risk assessment and attitudes towards nuclear power. Report Jul-Spez 115, Julich, Federal Republic of Germany: Nuclear Research Center, June 1981.
- Rethans, A. An investigation of consumer perceptions of product hazards. Unpublished doctoral dissertation, University of Oregon, 1979.
- Robertson, L. S. The great seat belt campaign flop. Journal of Communication, 1976, 26, 41-45.
- Rowe, W. D. An anatomy of risk. New York: Wiley, 1977.
- Salem, S. L., Solomon, K. A. & Yesley, M. S. Issues and problems in inferring a level of acceptable risk. R-2561-DOE. Santa Monica, Calif.: RAND Corp., 1980.
- Schoemaker, P. J. H. The expected utility model: Its variants, purposes, evidence and limitations. Journal of Economic Literature, in press.
- Schwalm, N. D. & Slovic, P. Development and test of a motivational approach and materials for increasing use of restraints. Final Technical Report PFTR-1100-82-3. Woodland Hills, Calif.: Perceptronics, Inc., March 1982.
- Schwing, R. C. & Albers, W. A. Jr. (Eds.), Societal risk assessment: How safe is safe enough? New York: Plenum, 1980.
- Simon, H. A. Models of man: Social and rational. New York: Wiley, 1957.

- Simon, H. A. The shape of automation for men and management. New York: Harper & Row, 1965.
- Sinclair, C., Marstrand, P & Newick, P. Innovation and human risk. London: Centre for the Study of Industrial Innovation, 1972.
- Slovic, P. Choice between equally-valued alternatives. Journal of Experimental Psychology: Human Perception and Performance, 1975, 1, 280-287.
- Slovic, P., Fischhoff, B. & Lichtenstein, S. Behavioral decision theory. Annual Review of Psychology, 1977, 28, 1-39.
- Slovic, P, Fischhoff, B., Lichtenstein, S., Corrigan, B. & Combs, B. Preference for insuring against probable small losses: Implications for the theory and practice of insurance. Journal of Risk and Insurance, 1977, 44, 237-258.
- Slovic, P., Fischhoff, B. & Lichtenstein, S. Accident probabilities and seat belt usage: A psychological perspective. Accident Analysis and Prevention, 1978, 10, 281-285.
- Slovic, P., Fischhoff, B. & Lichtenstein, S. Rating the risks. Environment, 1979, 21(3), 14-20, 36-39.
- Slovic, P., Lichtenstein, S. & Fischhoff, B. Images of disaster: Perception and acceptance of risks from nuclear power. In G. Goodman and W. Rowe (Eds.), Energy risk management. London: Academic Press, 1979.
- Slovic, P., Fischhoff, B. & Lichtenstein, S. Facts and fears: Understanding perceived risk. In R. Schwing and W. A. Albers, Jr. (Eds.), Societal risk assessment: How safe is safe enough? New York: Plenum, 1980a.
- Slovic, P., Fischhoff, B. & Lichtenstein, S. Informing people about risk. In L. Morris, M. Mazis and I. Barofsky (Eds.), Product labeling and health risks. Banbury Report 6. Cold Spring Harbor, New York: Cold Spring Laboratory, 1980b.
- Slovic, P., Fischhoff, B. & Lichtenstein, S. Perceived risk and quantitative safety goals for nuclear power. Transactions of the American Nuclear Society, 1980c, 35, 400-401.
- Slovic, P., Fischhoff, B. & Lichtenstein, S. Informing the public about the risks of ionizing radiation. Health Physics, 1981, 41, 589-598.
- Slovic, P., Fischhoff, B. & Lichtenstein, S. Response mode, framing, and information-processing effects in risk assessment. In R. Hogarth (Ed.), New directions for methodology of social and behavioral science: Question framing and response consistency. San Francisco, Calif.: Jossey-Bass, 1982.

- Slovic, P., Fischhoff, B. & Lichtenstein, S. Characterizing perceived risk. In R. W. Kates and C. Hohenemser (Eds.), Technological hazard management. Cambridge, Mass.: Oelgeschlager, Gunn & Hain, in press.
- Solomon, K. A., Nelson, P. F. & Salem, S. L. Utilization of risk analysis and risk criteria: A progress report. Working draft 1194-ORNL. Santa Monica, Calif.: RAND Corp., August 1981.
- Sowby, F. D. Radiation and other risks. Health Physics, 1965, 11, 879-887.
- Starr, C. Social benefit versus technological risk. Science, 1969, 165, 1232-1238.
- Svenson, O. Are we all less risky and more skillful than our fellow drivers? Acta Psychologica, 1981, 47, 143-148.
- Tversky, A. Elimination by aspects: A theory of choice. Psychological Review, 1972, 79, 281-299.
- Tversky, A. & Johnson, E. J. Alternative representation of perceived risks. Unpublished manuscript, Department of Psychology, Stanford University, November 1981.
- Tversky, A. & Kahneman, D. Availability: A heuristic for judging frequency and probability. Cognitive Psychology, 1973, 5, 207-232.
- Tversky, A. & Kahneman, D. Judgment under uncertainty: Heuristics and biases. Science, 1974, 185, 1124-1131.
- Tversky, A. & Kahneman, D. The framing of decisions and the psychology of choice. Science, 1981, 211, 1453-1458.
- U.S. Government. Teton Dam disaster. Washington, D. C.: Committee on government operations, 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, 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, 1978.
- U.S. Nuclear Regulatory Commission. An approach to quantitative safety goals for nuclear power plants. NUREG-0739. Washington, D. C.: The Commission, 1980.
- U.S. Nuclear Regulatory Commission. Toward a safety goal: Discussion of preliminary policy considerations. NUREG-0764. Washington, D. C.: The Commission, 1981.

Vlek, C. A. J. & Stallen, P. J. Persoonlijke Beoordeling van Risico's. Groningen, The Netherlands: University of Groningen, Institute for Experimental Psychology, 1979.

Vlek, C. A. J. & Stallen, P. J. Rational and personal aspects of risk. Acta Psychologica, 1980a, 45.

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

Weinstein, N. D. Unrealistic optimism about future life events. Journal of Personality and Social Psychology, 1980, 39, 806-820.

Wilson, R. Analyzing the daily risks of life. Technology Review, 1979, 81, 40-46.

Zebroski, E. L. A proposed national nuclear safety goal. Paper presented at the Seventh Energy Technology Conference, Washington, D. C., March, 1980.

Footnotes

1. In his forward to MacKay's remarkable book, Bernard Baruch observed: "Yet I never see a brilliant economic thesis expanding, as though they were geometrical theorems, the mathematics of price movements, that I do not recall Schiller's dictum: 'Anyone taken as an individual is totally sensible and reasonable--as a member of a crowd, he at once becomes a blockhead'" (p.iii).

2. The invariance obtained thus far with factor analytic studies does not imply, however, that approaches based on quite different methods and assumptions would also produce similar results. In fact, Tversky and Johnson (1981) have shown that a very different hazard structure results from representations based on judgments about how similar one hazard is to another with respect to risk. The implications of such differences remain to be determined.

3. In order to limit aggregated risk to society, it is necessary to take into account the size of populations at risk. A safety goal that encompasses the entire population would influence reactor siting policy and measures to mitigate the consequences of a major accident (e.g., emergency response provisions and design of containment structures).

Table 1

Results Obtained in a Nationwide Survey Titled
 "Risk in a Complex Society" Conducted in 1980
 for Marsh & McLennan by Louis Harris & Associates

	Top Corporate Executives (N=401)	Investors/ Lenders (N= 104)	Congress (N=47)	Federal Regulators (N=47)	Public (N=1488)
(a) TODAY'S RISK COMPARED TO THAT OF 20 YEARS AGO					
Q: Thinking about the actual amount of risk facing our society, would you say that people are subject to more risk today than they were 20 years ago, less risk today, or about the same amount of risk today as 20 years ago?					
More risk	38%	60%	55%	43%	78%
Less risk	36	13	26	13	6
Same amount	24	26	19	40	14
Not sure	1	1	--	4	2
(b) EXPECTATIONS FOR THE FUTURE					
Society has only perceived the top of the iceberg with regard to the risks associated with modern technology.					
Agree	19%	20%	47%	38%	62%
Disagree	78	71	51	60	28
Not sure	3%	9%	2%	2%	10
(c) FEELINGS ABOUT AMERICAN SOCIETY'S ATTITUDES TOWARD RISK: OVERLY SENSITIVE VS. MORE AWARE					
Q: Do you think American society is becoming overly sensitive to risk or are we becoming more aware of risk and taking realistic precautions?					
Overly sensitive to risk	56%	48%	26%	13%	15%
More aware and realistic precautions	28	41	57	80	78
Both	15	8	15	4	2
Not sure	1	3	2	2	5

TABLE 2. ORDERING OF PERCEIVED RISK FOR 30 ACTIVITIES
AND TECHNOLOGIES

(The ordering is based on the geometric mean risk ratings within each group.
Rank 1 represents the most risky activity or technology.)

	League of Women Voters	college students	Active Club members	experts
nuclear power	1	1	8	20
motor vehicles	2	5	3	1
handguns	3	2	1	4
smoking	4	3	4	2
motorcycles	5	6	2	6
alcoholic beverages	6	7	5	3
general (private) aviation	7	15	11	12
police work	8	8	7	17
pesticides	9	4	15	8
surgery	10	11	9	5
fire fighting	11	10	6	18
large construction	12	14	13	13
hunting	13	18	10	23
spray cans	14	13	23	26
mountain climbing	15	22	12	29
bicycles	16	24	14	15
commercial aviation	17	16	18	16
electric power (non-nuclear)	18	19	19	9
swimming	19	30	17	10
contraceptives	20	9	22	11
skiing	21	25	16	30
X-rays	22	17	24	7
high school and college football	23	26	21	27
railroads	24	23	20	19
food preservatives	25	12	28	14
food colouring	26	20	30	21
power mowers	27	28	25	28
prescription antibiotics	28	21	26	24
home appliances	29	27	27	22
vaccinations	30	29	29	25




FIGURE 1. Relation between judged frequency and the actual number of deaths per year for 41 causes of death. Source: Lichtenstein *et al.* (1978).

A civil defense committee in a large metropolitan area met recently to discuss contingency plans in the event of various emergencies. One emergency threat under discussion posed two options, both involving some loss of life.

Option A: Carries with it a .5 probability of containing the threat with a loss of 40 lives and a .5 probability of losing 60 lives. It is like taking the gamble:

.5 lose 40 lives

.5 lose 60 lives

Option B: Would result in the loss of 50 lives:

lose 50 lives


These options can be presented under three different frames:

I. This is a choice between a 50-50 gamble (lose 40 or lose 60 lives) and a sure thing (the loss of 50 lives).

II. Whatever is done at least 40 lives will be lost. This is a choice between a gamble with a 50-50 chance of either losing no additional lives or losing 20 additional lives (A) and the sure loss of 10 additional lives (B).


III. Option B produces a loss of 50 lives. Taking Option A would mean accepting a gamble with a .5 chance to save 10 lives and a .5 chance to lose 10 additional lives.

Figure 2. Decision framing: Three perspectives on a civil defense problem.




of lives lost
at one time

Curve 1



of lives lost
at one time

Curve 2



of lives lost
at one time

Curve 3

Figure 3. Task 2: The impact of catastrophic events. Subjects were asked to rate their agreement with the principles embodied in each of the above three proposals. (Two pages of instructions explaining the meaning of the curves preceded the task.)




Figure 4. Judgments of perceived risk for experts (top) and laypeople (bottom) plotted against the best technical estimates of annual fatalities for 25 technologies and activities. Each point represents the average responses of the participants. The dashed lines are the straight lines that best fit the points. The experts' risk judgments are seen to be more closely associated with annual fatality rates than are the lay judgments.

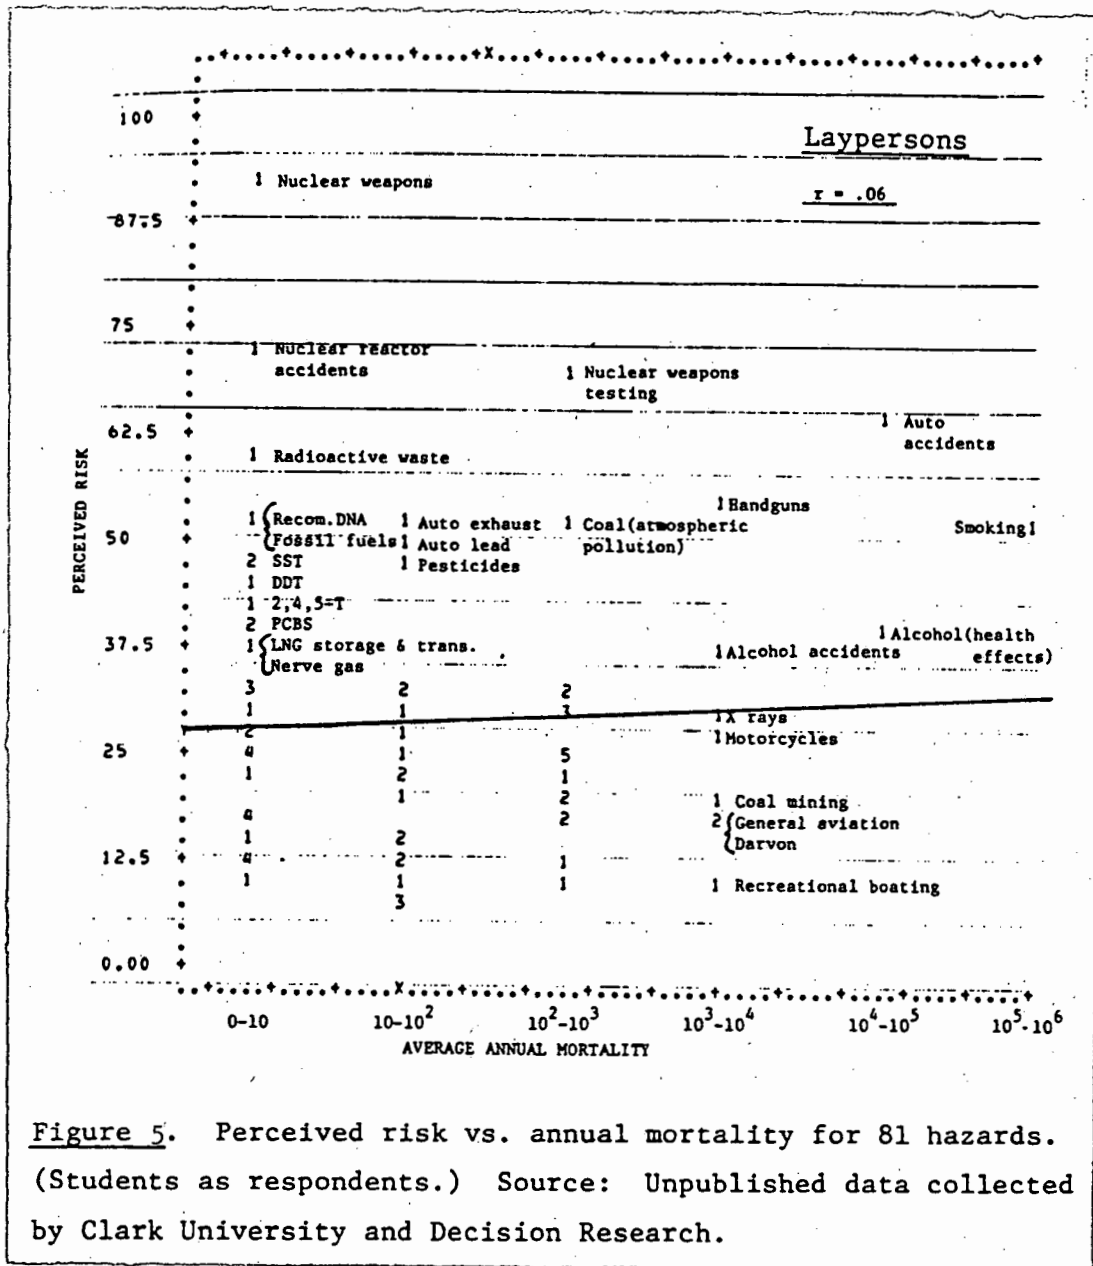


Figure 5. Perceived risk vs. annual mortality for 81 hazards. (Students as respondents.) Source: Unpublished data collected by Clark University and Decision Research.

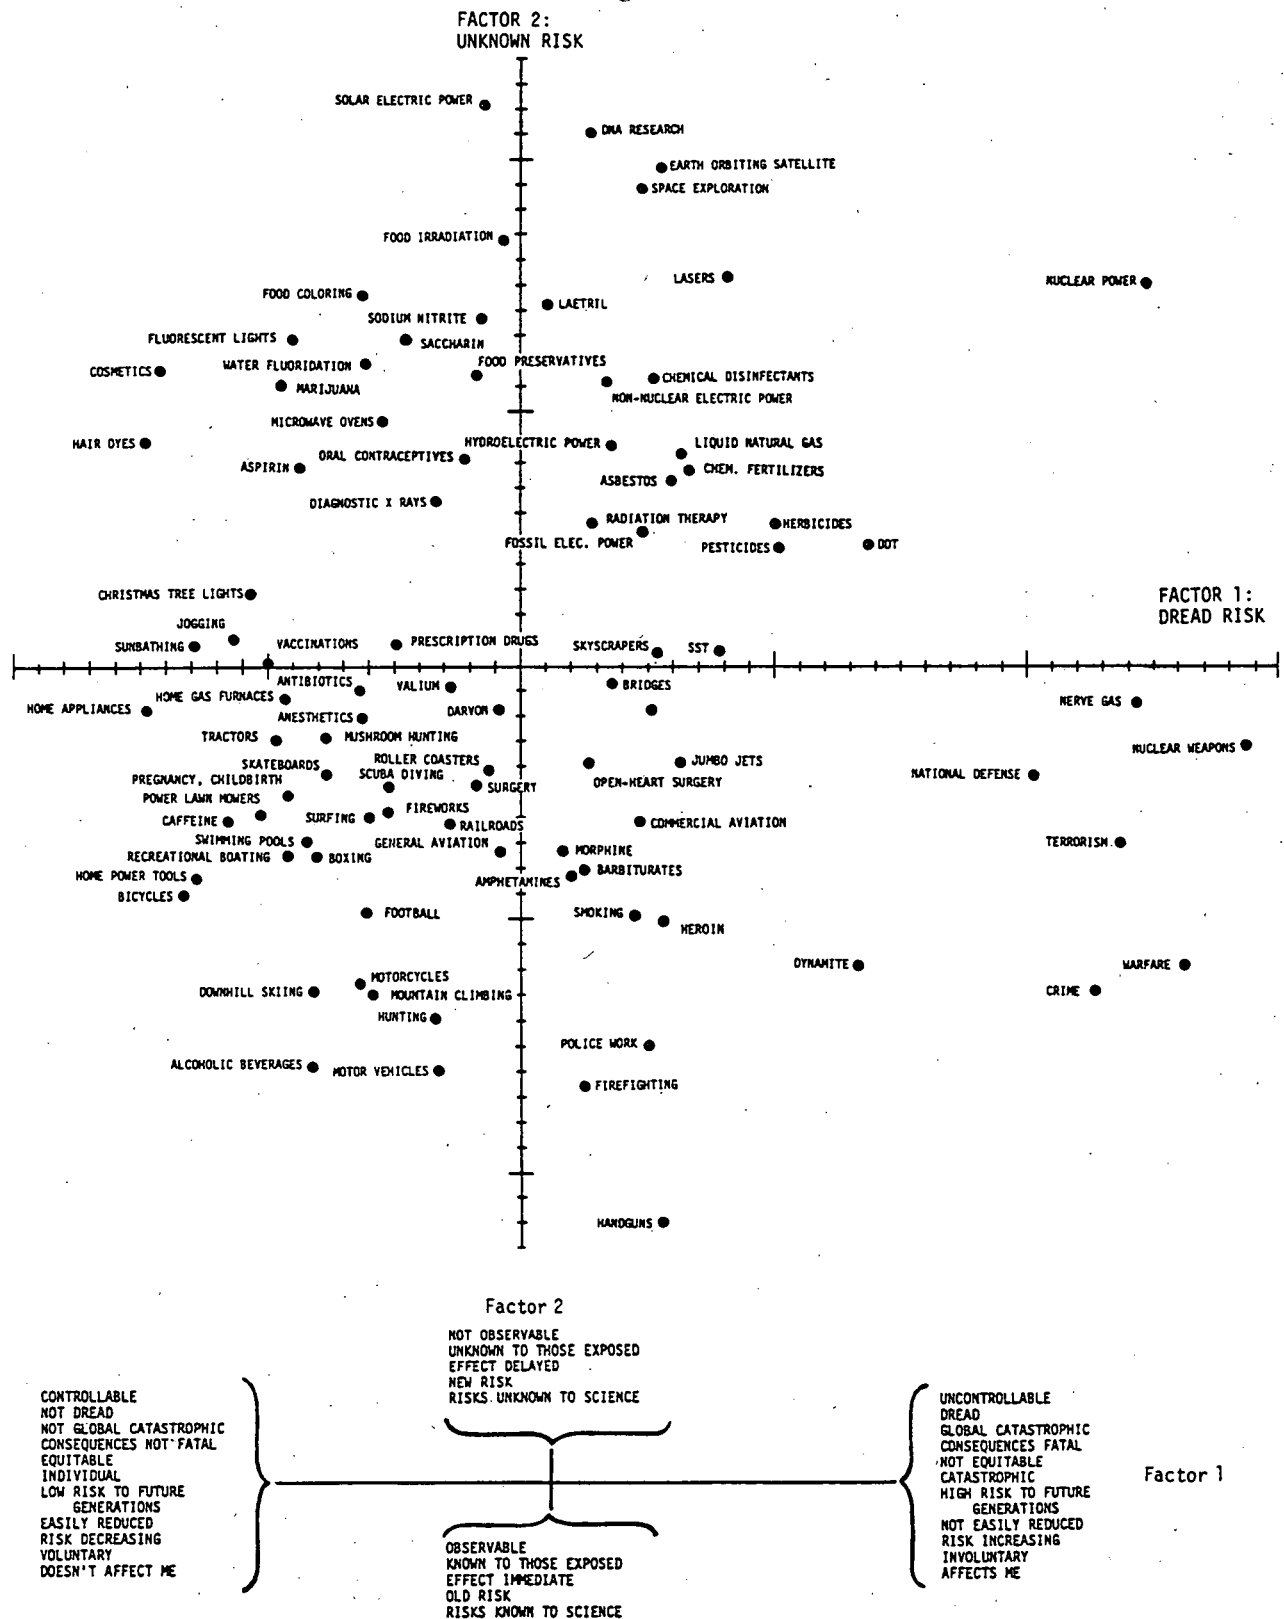


Figure 6. Hazard locations on Factors 1 and 2 of the three-dimensional structure derived from the interrelationships among 18 risk characteristics. Factor 3 (not shown) reflects the number of people exposed to the hazard and the degree of one's personal exposure. The diagram beneath the figure illustrates the characteristics that comprise the two factors.




Figure 7. Attitudes towards regulation of the hazards shown in Figure 6. The larger the dot, the greater the desire for strict regulation to reduce risk.

Source: Slovic, Fischhoff, & Lichtenstein (1981).