EVALUATING TECHNOLOGICAL RISK:
PRESCRIPTIVE AND DESCRIPTIVE PERSPECTIVES*

John W. Lathrop

June 1981
WP-81-78

CONTENTS

EVALUATING TECHNOLOGICAL RISK:
PRESCRIPTIVE AND DESCRIPTIVE PERSPECTIVES

John W. Lathrop

I. INTRODUCTION

Decisions concerning the deployment and management of novel or
hazardous technologies raise several issues involving the evaluation of
their impacts on society. Examples of such decisions include the siting of
a liquefied natural gas facility, the regulation of nuclear energy produc-
tion, and the screening and regulation of toxic chemicals. Each of these
kinds of decisions results in uncertain benefits and costs to society. It
would seem reasonable, then, that such decisions could be aided by any
of several analytic techniques, including cost-benefit analysis, or perhaps
decision analysis, which could include in the evaluation attitudes toward
uncertainty and value tradeoffs between conflicting objectives. However,
there are often special aspects involved in such decisions which can make

standard technical or economic analyses not very useful for aiding political decision-making processes. These aspects include outcomes of the decision having very serious negative consequences with very low probability, inequitable distribution of burden, large scale, novelty and others to be discussed below. Decisions involving such aspects sometimes come to be known as problems in managing social risk. Even though the word risk is currently in wide use in the media, it is often defined or used in different ways by different parties to the decision at hand. In spite of this serious problem, to be discussed at some length below, the need to appraise the risks presented by a new or hazardous technology has led to the development of several analytic techniques often referred to collectively as risk assessment. Yet those techniques generally assume, either implicitly or explicitly, certain prescriptive objective functions that are not sensitive to important societal concerns about potentially hazardous technologies. As a result, such techniques may be useful as inputs to a political decision making process, but are not as helpful as they could be if they used objective functions more descriptive of the relevant social concerns. This point is illustrated by cases where a member of the technical risk assessment community assumes a simple objective function, such as minimizing expected lives lost or life expectancy lost, observes individual or government behavior that does not minimize that objective function, then suggests that therefore something is wrong with the decision makers involved (see e.g., Rothschild 1979). That deduction is not the only one that could be made from the evidence. It could equally well be deduced that the objective function is inadequate.

The technical risk assessment community is not the only set of people addressing problems of social risk management. There is a growing body of research that is developing descriptions and explanations of human behavior that does not minimize narrow technical objective functions (Kunreuther, 1980). Yet the descriptions and explanations most sensitive to individual attitudes are not oriented toward developing broader social objective functions that can be directly used to aid the political decision making process in managing social risk. The development of evaluation models and techniques by the technical and psychological research communities forms a rough spectrum from the narrow but readily applied technical objective function to the broad set of social and psychological considerations that are not readily applied to aiding the decision process. This paper seeks to identify the part of that spectrum that, if further developed, may be the most useful in that it may provide usable decision aids that are also sensitive to the important societal concerns involved.

II. THE BASIS OF THE PROBLEM: TWO PERSPECTIVES OF RISK

While the previous section referred to the range of approaches to risk evaluation as a spectrum, it has underlying it two poles of thought which underly the theme of this paper: The basis of the problem of managing technological risk is that it is not a single concept, but is viewed from two distinct perspectives, the technical and the psychological.

A. Technical Perspective

From the technical perspective, technological risk is some probability distribution over sets of negative effects. Those effects are often limited to health effects, or casualties. For example, Keeney (1980b) defines the risk of an action as the probability distribution over sets of individual probabilities of dying, $\{p_i, i=1 \cdots N\}$, where p_i is the probability that the ith individual will die before the end of the next time period due to the action taken. The probability distribution over those sets is required to represent dependencies between the fatalities. Other analysts use summary measures of that distribution. In the Rasmussen report (Reactor Safety Study) the risk of a nuclear reactor is represented by a reverse cumulative probability distribution over numbers of fatalities per year per reference reactor (USNRC 1975). That distribution is sometimes referred to as a Rasmussen curve. Other risk assessments have used the same measure (Hazelwood and Philipson 1977). Other analysts go one step further and reduce the distribution to the expected number of health effects (e.g., Pate 1978). While it would be perhaps unfair to say that such analysts define risk as expected number of health effects, the fact remains that the negative consequences of a technological option are measured by that index.

There is one definition of risk not discussed here, and that is : "risk is probability times consequence". This is because this discussion is limited to representations of the risk of decision alternatives. Probability times consequence may describe the risk of an event, but could not in general describe the risk of any decision alternative with more than two

possible outcomes.

B. Psychological Perspective

In contrast to the fairly straightforward risk definitions listed above, extensive psychological research in the field of risk perception has suggested many more dimensions to be included in a definition of risk (Slovic et al. 1980, Linnerooth 1978, Kunreuther 1980). While these dimensions are covered in some detail in a later section, they are listed and briefly discussed here as a way of describing technological risk from a psychological perspective. They are listed roughly in order from the most easily adapted to a simple technical measure to the most difficult considerations to measure. While this discussion limits itself to health effects for brevity, there is no particular reason not to include other effects: financial, quality of environment, etc.

1. Expected number of health effects

While this aspect was listed above as part of the technical perspective, that does not mean that it cannot be a part of a psychological perspective also.

2. *Possibility of catastrophe*

Effects can be spread out over time and space, such as car fatalities, or they can be "bunched" into a catastrophe at one place and time, such as a possible major nuclear accident. That bunching can be very important for how society views the risk. Effects can also be bunched by cause, by identifiable population group, or by state of nature, as will be explained later.

3. *Equity*

There are actually two related aspects of equity that are important to perception of risk. The first is the amount of overlap between the populations at risk and at benefit. Clearly if a technology benefits none of those at risk, that is a case of inequity. But there is another, more subtle aspect of inequity that has to do with the ease of identifying the population at risk. If a technology can be identified as negatively affecting 100 people in the U.S., while benefiting all citizens, that may be seen as a necessary evil. But what if those 100 are all poor, or all asbestos workers? That is a matter of greater concern.

4. *Degree of control*

This aspect is a more general version of the voluntary/involuntary distinction made by Starr (1969). The central concept here is the level of participation of each potential impactee in each of two decisions: to expose himself to the risk, and to deploy the technology in the first place.

5. Attributability

This is an important aspect of social reaction to a risk that is often overlooked because of the cause-specific way risk analyses are done. It could be that generally incompetent engineering in cars kills far more people that a particular gas tank design. Yet the more easily identified cause of accidents, the gas tank, may give rise to a much stronger public reaction.

6. Non-probabilistic factors

Perhaps the most serious mismatch between technical and psychological perspectives lies in the evaluation of risk related to severe outcomes with no regard for the estimated probability of those outcomes. While there may be good reasons for that sort of evaluation, it can cause serious problems in developing consistent decision aids.

7. Non-decision comparisons

The search for criteria for acceptable risk often falls back on comparisons not involved in actual decisions. For example, comparisons are often made between some technological risk and moving to Denver, or smoking an extra cigarette, or driving an extra mile. Yet very few people choose between living near a nuclear reactor or living in Denver. Perhaps even fewer choose between living near a liquefied natural gas terminal and smoking an extra cigarette, though that decision involves some complicating factors.

B. Non- technological factors

Some aspects of a technology having little to do with possible health effects may have a great deal to do with perceived risk. Those aspects include the degree of centralization, the extent to which deploying a technology would infringe on civil liberties, how closely a new technology is linked with a high-consumption lifestyle, etc. While these aspects seem far removed from evaluating the health effect risk of a technology, they should be included in any effort to understand possible opposition to a technology. Such aspects may be as important or more so than any of the other aspects listed above in determining differences in what levels of risk are held to be acceptable.

C. Results of the Two Perspectives

The differences between the two perspectives presented above are very relevant to many problems in technological risk management. Technological risks are typically evaluated using analyses that assume the technical perspective on risk. Yet very often the political process in which the decisions actually get made is responsive to the psychological perspective. As an example of this problem, consider the case of the selection of a site for a liquefied natural gas import terminal in California. A technical risk analysis by a competent technical consulting firm found a proposed site at Oxnard, California, to be very safe, with very low risk to the community. Yet that analysis stated that a maximum credible accident (MCA) could involve up to 70,000 fatalities, though only by a seemingly incredible series of events that could only occur with an extremely low probability. In the political process, however, considera-

tion of such large numbers of fatalities led to a requirement that the terminal be sited remotely, away from any city, precluding the Oxnard site (Ahern, 1980). The applicant made plans according to a technical risk analysis, then had to change plans as a result of the political process.

As the example just presented illustrates, one result of the gap between the technical and psychological perspectives of risk is that technical risk analyses are often not effective in the political process. Conversely, another result is that the political process often does not receive the help it could from the scientific and engineering community.

Another result of the two perspectives on risk is that the community of risk analysis scientists and engineers work with analysis tools spanning a broad spectrum. One end of that spectrum involves prescriptive evaluation models that offer very clear guidance to decision makers, but that are not sensitive to psychological concerns. The other end of the spectrum involves psychologists and other social scientists identifying societal concerns very well, but not in a manner that results in clear guidance to the decision making process. Clearly, what is called for is more development of risk evaluation models that are sensitive to social and psychological concerns, yet are meaningful and useful aids to the political decision-making process. Section IV below spells out several considerations for that model development. Before that, however, the next chapter briefly reviews two existing approaches to risk evaluation that bracket the part of the analysis spectrum proposed for development in Chapter IV.

III. TWO APPROACHES TO RISK EVALUATION

A. Multiattribute Utility Functions

Within the framework of decision analysis, multiattribute utility theory can be used to develop essentially prescriptive risk evaluation models that take into account more than the expected number of health effects. As one example, Ralph Keeney (1980b) sets up a notation where "risk" *ex ante* is represented as a probability distribution over sets of probabilities, $\{p_i\}$, where each p_i is the probability that the ith individual will die due to the action being evaluated before the end of some time period. He represents "risk" *ex post* as a set of status indicators $\{ x_i \}$, where $x_i = 1$ if the ith individual has died due to the evaluated action in the time period, and $x_i = 0$ otherwise. He goes on to postulate von Neumann Morgenstern utility functions U_R on $\{p_i\}$, U_F on $\{ x_i \}$ and u on x, where $x = \sum_i x_i$, the total number of fatalities. He defines a desire for equity as a preference for more equal p_i's, sets up a consistency condition between U_R and U_F, and shows that these conditions lead in his notation to a risk evaluation function as simply u(x) and even shows that it must be convex downward, representing preferences that would prefer a technology characterized by a low-probability-high-loss over one with high-probability-low-loss and the same expected loss. This seems at odds with the general aversion to catastrophe that seems to be found in some public attitudes toward risk. Keeney goes on to add another condition that is consistent with a slightly more complex functional form. While Keeney's evaluation functions incorporate a particular form of equity preference

and an attitude toward uncertainty, the form of the evaluation model follows from more of a prescriptive than descriptive set of assumptions.

Sam Bodily (1980) develops another sort of utility function that accounts for several more descriptive aspects than does Keeney's models. Bodily's model accounts for the number of people sharing the risk, number of casualties per incident, the initial and final states of health involved, changes in individual probabilities of harm or benefit, how voluntary the risk is, and whether lives are being saved or lost. In a set of different examples of risk management alternatives, his calculations show the effective value of life varying by 170% due to differences in factors just listed. The key to Bodily's model is the combination of summed willingness to pay to avoid individual risk and a utility measure of attitude toward group risks in a single measure of social risk. While Bodily's model captures an admirable number of factors, it does not involve some other factors that are important, especially equity considerations. That is more or less an artifact of his particular presentation, however. The model framework Bodily proposes could be extended to include other factors of risk.

The evaluation models proposed by Bodily and Keeney represent promising attempts to model preferences concerning social risk that are not simply linear in numbers of health effects. Their models are basically applications of multiattribute utility theory, with the value elicitation questions effectively designed by the model. That is, the model is developed, with some consideration for what is important, then the elicitation questions are derived by what is necessary to derive the values of the parameters used in the model. This process leads to elicitation

questions that are very relevant to the issues involved, but that are quite unusual and do not occur to people in their normal experience. Approaches for improving this situation are discussed in the next chapter.

B. Psychological Considerations

Research on risk evaluation does not always entail a quantitative model that yields a scalar index of risk. In this section we will discuss psychological research that simply identifies and lists the considerations to be kept in mind in evaluating risk. The research selected as an example is the work of Paul Slovic, Baruch Fischhoff and Sarah Lichtenstein (Slovic et al. 1980, 1981; Fischhoff et al. 1978, 1980; Lichtenstein et al. 1978). They draw on a large body of experimental work of their own and others to enumerate the ways in which people are bad at probabilistic thinking, including the several biases and heuristics people have been found to use in choosing among alternatives with uncertain outcomes. They then develop implications of those biases and heuristics for problems in managing social risk. As they have done much of their work in problems with acceptance of nuclear power, they list the qualitative aspects of that technology that help to explain the lack of acceptance, as derived from experimental work. Those aspects include lack of control over the technology, dread, lethality of effects, potential for catastrophe, potential for unknown effects, and novelty of the technology. Each of these aspects is defined more precisely in their papers. It should be clear from this list that the work of Slovic, Fischhoff and Lichtenstein was a primary source for the list of aspects of the psychological perspective on risk presented in this paper. They go on to list problems in evaluating

social risk and recommendations for paths to acceptance of technological risk. They do a remarkable job of listing considerations to be kept in mind in evaluating and managing social risk, but do not develop any equations for an evaluation index for risk. In the most broad sense, it could be maintained that they do present a model of social preference concerning risk, in that their lists of biases and considerations explain observed phenomena of societal acceptance and rejection of certain risks. It could also be maintained that their work results in decision aids, in that they list recommendations for decision makers involved in managing social risk. However, they do not develop an evaluation model, if such a model is defined in the narrow sense as a set of equations resulting in a single evaluation index.

C. Limitations of the Approaches

The multiattribute utility approach does not provide the actual utility function to be used, but only the parameterized form of the function along with a protocol of questions to ask somebody whose answers can be used to calculate those parameter values. In other words, the utility function approach does not actually provide any answers, but only rephrases the questions into simpler, more understandable forms that are nevertheless just as difficult to answer. But this approach has even greater difficulties in that there is seldom such a consensus on value tradeoffs that it does not matter from whom you elicit the values. That gives rise to a whole host of problems concerning whose values to elicit and how to aggregate different values into a single risk evaluation model

(see Arrow 1977 for a review of social choice theory). The second approach discussed, labeled psychological considerations, is the most sensitive to the social value aspects of risk evaluation. However, because it does not provide a risk evaluation index, it is not directly useful as a decision aid.

There are two limitations shared by both of the approaches discussed. The first of these is that equity considerations may simply not be appropriately incorporated into any evaluation approach. Such aspects may only be approachable within the bargaining procedure of the political decision making process. An evaluation approach may hope to aid one or more parties in that bargaining, but any attempts to replace that bargaining with a model could be considered inappropriate.

The second shared limitation is perhaps the most serious. That is that the institutional structures making the risk management decision may not be compatible with the decision structure assumed by the evaluation approach. For example, the most basic assumption made by the utility approach is that there is some single self-aware process somewhere where the risks and benefits of the actual decision alternatives are compared. In fact, it is often the case that a regulatory agency, such as the Nuclear Regulatory Commission (NRC), is faced with what appear to be "yes/no" decisions on a single alternative that involves social risk. Of course, the NRC is actually participating in an unorganized way with several other government agencies and private companies in a set of actions which results in a selection of one of the risky alternatives. But where in that set of actions is the single self-aware process that weighs the risks and benefits of the different alternatives? Whose value

assumptions and tradeoffs should be used in the risk evaluation? Where should the results of the risk evaluation be delivered? It is entirely possible that the biggest limitation in the usefulness of a risk evaluation is the decision making process it is intended to serve (Kunreuther 1980). It could be that any improvement to be made in risk evaluation models would not be as important as it would be to organize the decision making process so that the risks of the actual decision alternatives are compared directly. That type of organization will be referred to in this paper as decision-focused. Yet a development of a risk evaluation model that is demonstrably sensitive to the aspects of risk management decisions that the political process is sensitive to could go a long way toward encouraging the political process to adopt that desirable decision-focused organization. The development of a risk evaluation model that could achieve that end would be an ambitious undertaking. The next section explores the initial considerations to be kept in mind as one embarks on such an adventure.

IV. TOWARD A MORE BROAD APPROACH

A. General Perspective

The previous sections have discussed research resulting in catalogs of social and psychological considerations, and contrasted that with multiattribute utility risk evaluation models, which provide a risk evaluation index. The previous sections have also explained the need for a risk evaluation model that lies between the index-producing models and the catalog approach. Such a model would provide a risk evaluation index,

and so be more directly useful a decision aid than the catalog approach, yet it would also be more sensitive to societal concerns regarding risk than the index-producing models reviewed above. This section will outline the fundamental considerations necessary to the development of such a model.

The discussion in this section will be based upon two major themes. The first is that a risk evaluation model should not be limited to providing only one of many inputs into a decision making process, but should be a direct aid to that process. That is, the model should not leave entirely to the decision makers the difficult job of putting together a large set of seemingly incommensurate pieces of data. Rather, it should deliver results in a form that aids the members of the process to integrate those results with the other considerations that enter in. The model should be designed to elevate the level of debate by providing a structured framework of reasoned evaluation within which the decision can be made.

The second theme which forms the basis of the following discussion is that the risk evaluation model should be built around risk considerations that drive the political decision making process. The order of development of the model should be from the primary concerns of the process to the value elicitations that would capture those concerns to the form of the model that would be logically fit to those elicitations. This is a subtle theme, at best, and one not incompatible with the models of Keeney and Bodily. The basic idea is to start from the observables that people react to in a risk management problem, and build the model around those.

B. General Methodology

The general methodology adopted here is based on multiattribute utility theory. That discipline is only briefly described here. The reader is referred to a basic text for a complete description (e.g., Keeney and Raiffa 1976). The methodology is based on the fitting of a multiattribute utility function (i.e., a multi-argument von Neumann Morgenstern utility function), to the answers a person gives to a structured set of questions. The basic aspects of preference captured by the function are tradeoffs between conflicting objectives (i.e., slopes of indifference curves) and attitudes toward uncertainty. The key development of the methodology is a set of theorems relating various plausible and testable assumptions about preferences to corresponding simple forms of multiattribute utility functions (MAUF). The fitting of the MAUF in a value elicitation interview then amounts to a set of assumption tests to identify the most simple acceptable form of the function, followed by the set of judgments neces- sary to fit the parameters of that functional form. Those judgments include indifference map comparisons and preference comparisons of simple uncertain alternatives. The fundamental concept underlying the methodology is that individuals are not good at choices between complex alternatives (such as actions involving multidimensional and uncertain outcomes), but can make choices between simple alternatives in a manner consistent with their underlying preferences. The methodology provides a mathematical system which can take alternatives too complex for consistent intuitive judgment, and evaluate them in a manner con- sistent with a person's judgments concerning simple alternatives, where

his judgments are more apt to accurately reflect his underlying preferences.

In typical expositions of the decision analysis methodology, the first steps in building the evaluation function are to identify the objectives of the decision maker, and describe an observable measure for each objective that represents the degree to which each alternative satisfies that objective. This section identifies typical objectives of the decision making process concerned with social risk·management, taking the psychological perspective of risk described earlier. Each objective is identified and characterized in a way that suggests possible observable measures that could be used to develop the evaluation function itself and its corresponding elicitation procedure. As this section represents only a first effort at the development of a comprehensive risk evaluation model, those further steps in model development are left to future research.

C. Objectives of Risk Management from a Psychological Perspective

In the objectives listed below, it is assumed that some single-dimensional measure of social cost has been decided upon that evaluates all the health and environmental effects upon an individual in some standard unit. There is no intention to suggest that such a measure would be easily derived, as it would require answers to such questions as: How many colds is a cancer worth? However, as the objectives listed involve considerations above and beyond the measure of the severity of a deleterious effect upon an individual, it clarifies the discussion considerably to assume such a measure exists.

As a second concession to clarity, the following discussions focus on the evaluation model itself, and so do not consider such frequently dominant problems as disagreements on the probabilities to be used in calculating the evaluation measure.

1. Expected number of health effects

While the expected number of health effects is a common measure of risk from a technological perspective, it can also be considered relevant to the psychological perspective. It is most relevant when the health effects are lives lost or saved, and expected lives lost or saved are considered equivalent to actual lives lost or saved. In that sense the expected fatalities measure has an ethical basis that leads to a troubling dilemma concerning the addition of other objectives to an evaluation function. A risk management strategy using simply expected fatalities as an evaluation function will, if successful, minimize expected lives lost. Any strategy using a more comprehensive evaluation function will allow more expected lives lost in any nondegenerate case. Thus that comprehensiveness in evaluation has "cost" some increment in expected lives lost (or decrement in expected lives saved). While that increment would not actually be a cost in terms of the social welfare represented by the evaluation function, it would nevertheless be the case that any attempt to use a more comprehensive evaluation model would be open to the attack that its use would cost human lives, expectationally.

This first objective is exceptional in that the measure of the objective to be used in the evaluation model is the objective itself. The major

measurement problems involved in this objective are hidden in the clari-
fying conventions presented at the beginning of this subsection (IV.C.).
The remaining objectives listed below do not have such obvious measures
associated with them.

2. Possibility of catastrophe

While the idea of possibility of catastrophe may be intuitively clear,
closer examination reveals a complex concept. As has often been men-
tioned in comparisons of coal and nuclear power as sources of electricity,
while expected fatalities per megawatt-year (electric) may be estimated
as less for nuclear than for coal, the fact that the nuclear fatalities occur
in high-fatality, low-probability accidents explains public resistance to
nuclear power (Barrager et al. 1976). This comparison will be mentioned
again below, since in fact there are a number of reasons that could
explain that resistance aside from the possibility of catastrophe. How-
ever, for this discussion the example is useful in considering just what
catastrophe means. It seems clear that a catastrophe is a "bunching"
together of fatalities, but on what dimension does that bunching occur?
Fatalities can be bunched together in one place, or at one time, or in one
state of nature, or by one identified cause, or in one previously identifi-
able group of people. These various types of bunching may correspond to
different types of preference mechanisms.

The risk of a nuclear reactor involves bunching on all five of the
dimensions just listed: fatalities are concentrated in states of nature that
correspond to an accident, are grouped downwind of the facility at the

time of the accident, are bunched by single cause, and impact the neighbors of the reactor. A coal plant induces fatalities that occur with some degree of certainty and so are distributed over many states of nature, and are distributed in space, time, and groups of people. Coal plant fatalities are not even bunched by identifiable cause, in that respiratory ailments caused by the coal plant cannot be distinguished from ailments with other causes. Bunching is not always so highly correlated across causes. For example, fatalities due to nuclear waste may be bunched by state of nature (failure of casings, misunderstanding of physical processes), and bunched by location, but not bunched by identifiable cause, and distributed over very long periods of time and groups of people (generations). Risk due to train derailments of toxic chemicals may be concentrated in the poor segments of the population who live near railroad tracks, though it is distributed over several accidents (in different states of nature, places, times) each with very few fatalities.

Bunching by identifiable cause is evaluated in the attributability objective discussed later. Bunching by identifiable group is evaluated as part of the concept of equity also discussed below. That leaves catastrophe defined here as a bunching in space, time, or state of nature. Another way of defining a catastrophe is as a number of fatalities linked in any way. By either definition, the problem remains to evaluate attitude toward catastrophe. The measure could be a set of probability distributions over numbers of health effects in each different type of catastrophe, if more than one type is possible.

Efforts thus far in the evaluation of potential for catastrophe have taken three paths. First, a nonlinear value function over number of fatalities has been suggested, either as an ad hoc function or as a von Neumann Morgenstern utility function (Keeney 1980a). In the cited reference Keeney speculates that such a curve would be convex downward due to riskless preference effects (diminishing marginal disutility for health effects). That convexity represents a preference for bunched over distributed fatalities. Second, a multiattribute utility function has been suggested, where the health status of each person is an attribute (Bodily 1980). With such a function, the interaction terms represent attitudes toward catastrophe. With both of these evaluation strategies, if only the suggested utility function is used the evaluation is only sensitive to bunching in states of nature, and does not account for differences in bunching in space and time. Neither model has the required argument structure or elicitation protocol to be sensitive to bunching in space and time that is not bunched by state of nature. While the required extension of the model's notation would be relatively straightforward, attempts to extend the elicitation protocols as necessary reveal some fundamental problems to be discussed in a subsequent paper.

3. Equity

Just as discussions concerning risk often suffer from lack of a definition of risk, discussions concerning equity may involve as many as three distinctly different concepts going by the same name. Each concept is discussed in turn here.

a) *Correspondence between populations at risk and at benefit.*

Strip mining may involve risk externalities imposed on people of the mining region so that people in a distant metropolis can run air conditioners. This form of the equity problem has been addressed at great length in the social choice literature, and so will not be discussed at any depth here. It should be noted, however, that no consensus exists as to the best method for evaluating a situation that is inequitable in this sense (Arrow 1977). After risk evaluations for each of the populations whose members are roughly equally affected, the use of those measures in a decision may be one aspect of risk evaluation best left to the political process.

b) *Ease of identifying people at risk: size of individual probabilities.*

While the size of an individual's probability for a health effect is of interest for willingness-to-pay calculations, in this discussion that number is used to represent a different concern: how much society knows about who and how many will be victims. Suppose a technology is estimated to cause one expected fatality (EF). The size of the population bearing that risk would be of great importance to its evaluation. Consider that 1 EF risk to one person is very different from 1 EF risk equally shared by 100 people, or 1 EF risk equally shared by 10^8 people. The difference to society between these cases is in its level of knowledge as to *who* and *how many* may die. At a fixed number of expected fatalities, the variance over numbers of fatalities rises with the number of people equally sharing the risk, though only very slightly once there are more than 20 people.

While this sensitivity to the size of the group sharing one EF of risk could be confounded with sensitivity toward attributability, discussed below, it could not be confounded with sensitivity toward catastrophe, as there is no consideration for intercorrelation of the occurrence of the fatalities.

The relationship between differences in individual probabilities and variance in number of fatalities is that for a given level of expected fatalities, the less different the individual probabilities, the greater the variance in number of fatalities. This relationship is interesting in that it explains a problem presented by Keeney in his 1980b paper. Keeney establishes an evaluation of numbers of fatalities by a von Neumann Morgenstern utility function over that number. He points out that an aversion to catastrophe would correspond to a concave downward, or uncertainty averse, utility. He then defines a desire for equity as a desire to make any two individual probabilities of dying less different. Finally, he shows that such a desire for equity corresponds to a convex downward utility function, and so is incompatible, in his model, with aversion to catastrophe. The significance of this conclusion becomes more clear when one realizes that any change to increase equity (his definition) also increases the variance in number of fatalities. That is, Keeney has basically made the statement that "equity is uncertainty". It is from that statement that his conclusion flows most directly. This matter is an excellent example of how surprising conclusions can be drawn from simple evaluation models and simple definitions of measures of social concerns, such as equity and potential for catastrophe.

c)   *Ease of identifying people at risk: ease of defining the impacted group.*

As with the previous concept, the key to this concept of equity is knowledge, but in this case in a more subtle way. Suppose some advance in science enables the certain identification of a victim before he is killed by a technology. Even though society now has perfect knowledge as to who the victim will be, there may not be a perception of inequity if the victim cannot be easily described. If there is nothing in common among the victims except that they are American, say, then the impact of the technology may be seen as equitable. If, however, all victims are black, or all asbestos workers, or all live downwind of a nuclear reactor, or all are poor, the impacts would be seen as inequitable, regardless of how large or small the group sharing the risk is.

The three general considerations described so far (expected number of effects, catastrophe, and equity) are three distinct measures of social risk. A three-dimensional table could be laid out, with each cell filled with an example representing a different combination of levels of the three measures of social risk. Each of those examples could be described by a probability distribution over various dimensions, such as numbers of health effects, classes of people, space, and time. The remaining considerations, discussed below, are not so amenable to quantification.

4. Degree of control

The concept of degree of control has two different aspects: an impacted individual's participation in the decision to expose himself to the risk, and an individual's participation in the decision to deploy the technology. The former aspect coincides with the voluntary vs involuntary risk distinction made by Starr (1969) and discussed by others (Otway and Cohen 1975). Starr presented evidence which he interpreted as indicating that society has a much higher threshold of acceptability for risks involving voluntary exposure than for risks incurred involuntarily. Arguments against that hypothesis have appeared in the literature, most recently in the paper by Slovic and Fischhoff (1981), which concludes that apparent aversion to involuntary risk can be better explained by the higher potential for catastrophe and inequity that often accompany that type of risk. However, another paper involving the same two authors stresses the importance of public participation in the second of the two decisions listed above, concerning the deployment of the technology (Fischhoff, et al 1980). That second decision involves an aspect that is basically different from any consideration mentioned so far, in that it has little to do with the physical source or impactee, but addresses the process which generated that source and impactee. As some have pointed out (Green 1981) one determinant of that elusive concept called acceptable risk is the acceptability of the process that generated the risk.

5. *Attributability*

Attributability of cause is the first in a series of considerations that involve very large steps away from a model that could be easily defended on prescriptive grounds, steps toward a descriptive model. In some ways, this consideration is related to equity: while some aspects of equity have to do with ease of identifying an impactee, attributability concerns the ease of identifying the source of the effect. Attributability is also highly correlated with catastrophe, as there is usually little doubt as to the immediate cause of a catastrophe, while it would be much harder to identify the cause of a more distributed set of effects. If a set of effects is clearly attributable to a single cause, that cause is more apt to be discovered and to give rise to societal corrective mechanisms than is a cause with more subtly distributed effects. A gas tank design feature that makes a car go up like a torch when hit is much more apt to be reacted against than a steering design flaw that leads to crashes erroneously attributed to driver error.

The last example suggests that attributability, like equity, is an aspect very dependent on level of information. As epidemiological studies and national medical reporting systems become more effective, diseases that would have gone unnoticed and unfeared a few years ago become centers of attention and concern. Effects of pesticides, defoliants, and other widely dispersed chemicals cause concern only when they are recognized as coming from controllable chemicals, yet those effects may only be able to be detected by very sophisticated techniques, if at all. Attributability has more than the two levels: recognized and unrecognized. It concerns how easily the cause can be identified. Of course, a

society will only react to risks that are recognized, but beyond that, the ease with which the cause can be described will affect the ease of ideation of the risk as a menace. To return to the car gas tank design example, the image of an engineering staff making a fairly well-defined decision not to incorporate a safety feature is a very clear target for public reaction. It could be that a general lack of competence of an automotive engineering staff could contribute much more to driving risk than any single decision about gas tanks, but general lack of competence is much harder to identify, and even if identified, forms a much more diffuse image for public reaction to try to focus on. While these aspects are part of a descriptive basis for attributability as a dimension of risk, they also suggest a prescriptive basis, in that the ease of identification of a cause is apt to reflect very closely the ease of management of the risk.

6. Non- probabilistic factors

Perhaps the most important dimension of social risk from the psychological perspective is the size of the maximum potential catastrophe, considered *without* any weighting by its probability. The idea of evaluating an alternative by a possible outcome without considering its probability is a very large step away from a prescriptive evaluation model and toward a descriptive one. It can be extremely troublesome for the analyst, not only because it doesn't seem to make sense to someone used to probabilistic models, but also because the description of a maximum potential catastrophe is dependent on the imagination of the analyst, and on an ill-defined notion of what "potential" is. Yet the political process

typically does not share the probabilistic perspective that is so funda-
mental to risk assessment. One example involves the attempt to site a
liquefied natural gas (LNG) terminal at Oxnard, California. While risk ana-
lyses found the site to be safe, maximum credible accident (MCA)
scenarios involving very large numbers of fatalities formed a large part of
the political debate which led to the effective rejection of the site. That
debate did not seem very sensitive to the extremely low probabilities
that could be assigned to the MCA's (Ahern 1980).

There are a number of explanations that could be invoked for this
non-probabilistic evaluation. First, it could be that members of the politi-
cal process do not know how to think probabilistically, or do not have a
feel for what a very low probability means. Second, members of the pol-
itical process may simply doubt the low probabilities. That doubt may be
well founded, as it is impossible in any actual case to guarantee the com-
pleteness of any risk analysis, and major accidents can be identified that
were pronounced impossible before their occurrence (e.g., the Titanic),
or that were not effectively described by a preceding risk analysis. In a
particularly strict sense, no absolute probabilities are ever calculated by
a risk analysis. The probabilities reported by such analyses are condi-
tional on the validity of the assumptions of the analysis, including those
concerning its completeness and descriptive validity. It may be very rea-
sonable for members of the political process to doubt that an extremely
low probability is in fact that low, given its conditional nature.

A third reason for non-probabilistic evaluation is a desire for resilient
social support systems. Given the extreme uncertainty in what future
demands will be made on, for example, an energy supply system, and the

uncertainty in the behavior of that system, it may make sense to design it to be as resilient as possible. One way to promote resilience is to limit the maximum potential catastrophe or MCA.

A fourth reason for non-probabilistic evaluation is a sensitivity to the concept of dread as a social cost. If a large plant is erected that has an MCA involving the deaths of 40,000 people living downwind, those people are going to live their lives with a sense of dread that is not very directly related to the very low probability a risk analysis might assign to the MCA. That sense of dread may be a very real decrement in their quality of life, regardless of how safe the plant is as measured by any risk analysis.

It may be helpful to give some meaning to the idea of limiting the MCA of an option with a brief example. The most dramatic MCA's involve accidents at large scale facilities that send poisons or flammable gasses into the air that can cause large numbers of fatalities downwind. Examples include nuclear reactors, LNG terminals, and chemical plants. In any of these cases, the MCA can be limited by requiring that the facility be sited remotely, but only at a financial and perhaps environmental cost.

7. Non- decision comparisons

All the factors discussed above could at least be fit into a decision analytic framework, in that each could be involved in a comparative evaluation of at least two alternatives in a decision. However, there are comparisons made in some risk analyses that have nothing directly to do with the decision being considered (Cohen and Lee 1979). It makes immediate sense to compare the risks of a coal electricity-generating

plant with a nuclear one, as those are two realistic alternatives considered by a utility in expanding its capacity. It makes much less immediate sense to compare the risks of living near a nuclear reactor with the risk of moving to Denver, driving a car an extra three miles, or smoking an extra pack of cigarettes. The last three comparisons do not relate very directly to any decision, as few people actually do choose between a reactor neighborhood and Denver, or between living near a reactor and smoking more.

The intent of the non-decision comparisons listed above is to fit some technological risk into the same scale with risks people normally accept, then go on to draw conclusions regarding the acceptability of the technological risk. There are a number of problems with this approach as a way to establish the acceptability of a risk. However, non-decision comparisons are based on an idea that could be incorporated in a risk evaluation model: people make choices every day that involve risk, and those typical risks may form benchmarks on an individual's scale of perceived risk. There are problems in that the various risks considered are evaluated differently (nuclear power, background radiation, driving, smoking), but once those differences are accounted for within an evaluation function, routine decisions involving risk could be used as meaningful points of reference on the scale of the risk evaluation index (see Lichtenstein et al. 1978).

8. Non-technological factors

While the set of factors described so far has spanned a wide range of characteristics, all of them have been concerned with the deleterious health and environmental effects of the evaluated technology. In this final step in the progression of factors, even that one common thread is abandoned as non-technological factors, factors other than health and environmental effects, are considered. A number of analysts have observed that resistance to nuclear power may be based in part on a resistance to the political side effects of the technology: the centralization of power that might accompany a centralized generation of electricity; the loss of civil liberties that might result from efforts to prevent diversion of nuclear material, etc. There might also be resistance to perceived effects on lifestyle: the more materialistic, growth-based, exploitative society made possible by nuclear power; the less natural, less simple, less human society associated with the high technology of nuclear power. While the attribution of such far-ranging effects to nuclear power is certainly arguable, the fact remains that there is a body of survey evidence establishing that segments of the population do ascribe these effects to nuclear power (Otway and Pahner 1976).

It is one thing to establish that non-technological effects are important to some people's evaluation of a technological alternative. It is quite another problem to decide how that fact would be used in developing a risk evaluation model. The incorporation of non-technological effects would be required in a risk evaluation model intended to predict political resistance to the deployment of a technology. Yet that incorporation would expand such a model beyond the role of evaluating risk. It seems

that the decision whether or not to attempt to include non-technological effects in the model is a decision on model scope. Is the model to be a risk evaluation model or is it to be a resistance prediction model? The discussions in this paper are aimed at the development of a risk evaluation model sensitive to societal concerns. That model does not have to be a resistance prediction model in order to aid the political decision making process in evaluating social risk. It follows, then, that while non-technological factors are a useful end-point in a progression of social concerns, they are not necessarily to be included in a risk evaluation model. Perhaps the most useful role such factors play in the discussions of this paper is in defining the limits to which a risk evaluation model should go in the progression from technical risk evaluation to aiding the political decision making process.

D.   Incorporating Risk Evaluation Factors into a Risk Evaluation Function

The previous section listed eight factors involved in the evaluation of social risk from a psychological perspective. While it has been argued that the eighth factor is not appropriate for inclusion in a risk evaluation model, the other seven are. The development of a risk evaluation model is then defined as the development of measures of each of the seven factors, then the combination of those measures into a single index. While these two tasks are the subject of future research, this section briefly comments on some aspects of the two tasks apparent from an examination of current risk evaluation work.

The first thing to note about the problem of incorporating the different factors into a single evaluation model is that each of the existing measures incorporates only a few of the various factors. Rasmussen curves, for example, measure the potential for catastrophe, could be used to define a maximum credible accident as a non-probabilistic measure, and have been used for non-decision comparisons of social risk (e.g., reactor risk vs. meteorite risk). However, such curves do not reflect any of the other factors. Descriptions of risk involving the set of individual probabilities of health effects, $\{p_i\}$, lend themselves to reflecting concerns about equity, and are appropriate for non-decision comparisons at an individual level, but do not reflect the other factors. Of course, Rasmussen curves could be calculated for each of several subgroups of the population to address concerns about equity, and the work of Keeney (1980b) and Bodily (1980) extends evaluations of $\{p_i\}$ to include concerns about catastrophe, but the fact remains that no existing evaluation measure or model addresses all seven of the factors described.

Sometimes the assumptions involved in a risk evaluation model can restrict in a very subtle way which factors are addressed. For example, Keeney (1980b) sets up a notation and a basic model in such a way that assumptions about treating everyone equally and correspondence between *ex ante* and *ex post* utilities lead to a single-attribute utility function on number of fatalities as a risk evaluation measure. While his notation, basic model, and assumptions are intuitively compelling, the resulting measure is sensitive only to expected numbers of fatalities, to bunching by state of nature and to equity in an equally-shared p_i sense. The other factors are missed not because of obvious features of the model,

but because of a lack of dimensionality in the original situation to be evaluated.

The example of Keeney's model is given because it illustrates the importance of the first step in developing a comprehensive risk evaluation model: the description of the effect space to be evaluated. Keeney's description of the effect space as a probability distribution over $\{p_i\}$ accounts in large part for the nature of his resulting risk evaluation measure. The same information cast into a catastrophe-by-equity-by-expected fatality space could lead to a very different risk evaluation measure, but one still limited in the factors it could address. In order to address all factors, an evaluation model must begin with an effect space in which all factors are represented. Because the seven factors involve such different aspects as bunching of effects and the participatory nature of the process, there are several different ways to set up the effect space. Clearly, the first step in future research on the development of risk evaluation models is an examination of those different effect spaces to select the one most appropriate for a meaningful value representation and elicitation procedure.

V. SUMMARY

This paper has discussed the problem of developing better risk evaluation models. After an examination of the central problem of risk evaluation, defining what risk is, two risk evaluation approaches were presented that bracket a particular part of the analysis spectrum in need of development. That approach involves the expansion of multiattribute

utility risk evaluation functions to be more sensitive to social and political concerns regarding risk management. Initial considerations for research in that direction were laid down. That approach appears to be feasible, and promises a more useful role for analyses in the process of social risk management.

The most unusual feature of the broad approach proposed in this paper is that it seeks to explicitly recognize some of the political realities involved in the social risk management process. An attempt is proposed to develop an evaluation model that deviates from prescriptive considerations to account for preferences expressed in the political decision making process. This attempt is marked by several difficulties, perhaps the most notable among them being the ethically troublesome result mentioned earlier: evaluation models sensitive to any feature in addition to or other than expected fatalities will generally advise the selection of policies which do not maximize expected number of lives saved. Policy makers using such models could be accused of "sacrificing lives, expectationally" for the sake of being responsive to other societal concerns. Of course, expected lives are allowed to be lost to gain on other dimensions all the time in public policy decisions, but it may be difficult to do so deliberately, as an explicit step in the policy making process. While this example is given as the most troublesome one, it is clearly only one of many difficulties to be encountered in explicitly incorporating sensitivity to societal concerns in the political decision making process.

The main potential benefit of the proposed broad evaluation model is that it seeks to avoid the problem of a technically correct risk analysis advising the promotion of a project that is eventually rejected due to

popular opposition. Another benefit is its general decision aid orienta-tion, an improvement over standard risk analyses which are intended only as inputs into the decision making process.

There are two major limitations to the broad approach proposed. First, the more explicit use of subjective judgments involved in the value elicitations can cause problems with the defensibility of the approach. This is in turn a problem with the political decision making process, since those subjective judgments generally must be made somewhere in the decision making. They are either hidden in the assumptions of the model, or they are explicit. But the fact remains that a government agency may have more trouble defending a decision based on an analysis with expli-citly subjective judgments than one based on an apparently objective analysis.

The second problem is more basic, but less specific to the broader approach proposed. That problem is that better evaluation models may not help social risk management as much as would better consideration of actual decision alternatives. As was mentioned in an earlier section, risk analyses are typically not part of a political decision making process that recognizes the choice faced between several risky alternatives. As mentioned before, the Nuclear Regulatory Commission seems to be faced continually with decisions that have the appearance of being "yes/no" decisions on a risky option, when it is actually participating in an unor-ganized way with several other government agencies and private com-panies in a process that makes "which" decisions, choosing one of several risky alternatives. A political process that explicitly faced such "which" decisions would be able to make better use of comparative risk analyses

(Ahern 1980). But perhaps this identifies another potential advantage of the proposed approach. A risk evaluation model that is oriented toward dimensions of concern to the political process could be designed in such a way that comparison with the actual alternative is an intrinsic part of the evaluation.

REFERENCES

Ahern, W.   1980.   "California Meets the LNG Terminal", *Coastal Zone Management Journal 7*, p. 185-221.

Arrow, K.   1977.   "Current Developments in the Theory of Social Choice, *Social Research 44* (4).

Barrager, S.M., B.R. Judd and D.W. North   1976.   The Economic and Social Costs of Coal and Nuclear Electric Generation.   Washington D.C.: National Science Foundation.

Bodily, S.E.   1980.   "Analysis of Risk to Life and Limb", *Operations Research* 28 (1), Jan.-Feb.

Cohen, B. and I. Lee.   1979.   "A Catalog of Risks", *Health Physics*.

Fischhoff, B., S. Lichtenstein, P. Slovic, R. Keeney, and S. Derby.   1980. *Approaches to Acceptable Risk: A Critical Guide*.   NUREG/CR-1614. Washington, DC: U.S. Nuclear Regulatory Commission.

Fischhoff, B., P. Slovic, S. Lichtenstein, S. Read, and B. Combs.   1978. "How Safe is Safe Enough? A Psychometric Study", *Policy Sciences* 8.

Green, H.   1981.   "Implications of Nuclear Accident Preparedness for Broader Nuclear Policy.   To appear in: J. Lathrop (ed.) *Planning for Rare Events: Nuclear Accident Preparedness and Management*, Oxford:Pergamon Press.

Hazelwood, R. N. and L.L. Philipson.   1977.   Survey of LNG Risk Assessment, Report for the California Public Utilities Commission, Socio-Economic Systems, Inc., Los Angeles.

Keeney, R. L.   1980a.   "Evaluating Alternatives Involving Potential Fatalities", *Operations research*, 28 (1), Jan.-Feb.

Keeney, R. L. 1980b. "Equity and Public Risk", *Operations Research* 28 (3), May-June.

Keeney, R. L. and H. Raiffa. 1976. *Decisions with Multiple Objectives*. New York: Wiley and Sons.

Kunreuther, H. 1980. "Societal Decision making for Low Probability Events: Descriptive and Prescriptive Aspects," WP-80-164.

Lichtenstein, S., P. Slovic, B. Fischhoff, M. Layman, and B. Combs. 1978.. "Judged Frequency of Lethal Events", *Journal of Experimental Psychology: Human Learning and Memory*, 4 (6).

Linnerooth, J. 1978. "Reevaluating the Value of Life: Practical Considerations", invited paper presented at ORSA/TIMS Conference, Los Angeles, November.

Otway, H. J. and J. J. Cohen. 1975. Revealed Preferences: Comments on the Starr Benefit-Risk Relationships. Research Memoranda RM-75-5, IIASA, Laxenburg, Austria.

Otway, H. J. and P. D. Pahner. 1976. "Risk Assessment", *Futures*, 8 (2).

Pate, M.-E. 1978. "Public Policy in Earthquake Effects Mitigation", Technical Report 30. Blume Earthquake Engineering Center, Stanford.

Rothschild, L. 1979. "Risk," *Atom 268*, February.

Slovic, P., and B. Fischhoff. 1981. "How Safe is Safe Enough? Determinants of Perceived and Acceptable Risk", in L. Gould and C. A. Walker (Eds.), *Too Hot to Handle: Social and Policy Issues in the Management of Radioactive Wastes*, New Haven: Yale University Press.

Slovic P., B. Fischhoff and S. Lichtenstein. 1980. "Facts and Fears: Understanding Perceived Risk", in R. C. Schwing and W. A. Albers (Eds.), *Societal Risk Assessment: How Safe is Safe Enough?* New York: Plenum Press.

Starr, C. 1969. "Social Benefit Versus Technological Risk", *Science*, 165: 1232-1238.

USNRC. 1975. *Reactor Safety Study*, ("Rasmussen Report"), WASH-1400, Washington,D.C.: U.S. Nuclear Regulatory Commission.