COLLECTIVE DECISIONS INVOLVING RISK: A LITERATURE REVIEW

by

ROD DOBELL and TED PARSON

August 1986

<u>IRPP Working Paper</u>

Preface

This working paper was developed in response to a request for an address on the responsibilities of public servants in dealing with risks to the public. It surveys a portion of the growing literature relevant to collective decisions involving risk. It is intended primarily to provide some helpful introductory background to public servants and students of public administration coming to this topic without substantial expertise. Its purpose is therefore expository, and the style descriptive. This review simply attempts to organize some elements of a large and disparate literature, without developing any one analytical approach or institutional process in detail, and with no pretense of being complete or exhaustive.
The address in question, and some further more speculative considerations, are discussed in a companion IRPP working paper titled, "Collective Decisions Involving Risk: Three Essays" (August 1986).

CONTENTS

INTRODUCTION

That deep confusion prevails in debate over public policy on risk is illustrated by the substantial disagreement that exists even over how risk should be defined for public policy purposes. Only the broadest possible definition - risk is the possibility of harm - would receive widespread agreement. In specifying what kind of harm is meant, or by what means or in whose opinion the possibility is determined, there are several competing views. Some define harm as any loss, others argue that a matter is only risky if the possibility of injury, impairment of health, or loss of life are involved. Because the most difficult and controversial decisions are those involving harm to life and limb and because the arguments made in this area can be generalized to include other kinds of harm, we focus here on risks to life and limb.

Concerning the relevant meaning of "possibility", there are at least three schools. Scientists and engineers normally define risk in terms of objectively determined probability of harm, as evinced by either observed prior occurrences or the clear theoretical potential for hazard. Psychologists emphasize individually perceived risk, feelings of threat of harm that may be unrelated to the degree of objectively measured risk, but that nevertheless are the factors that move people to take or avoid risks, or to protest their imposition. Finally, a body of work from a sociological and anthropological perspective emphasizes the consensual nature of risk definition and the role that cultural conditioning plays in determining what risks concern particular social groups and what strategies they consider for coping with them. As an instance of cultural determination of risk over-riding objective factors, Douglas and Wildavsky (1982) cite the case of a people in Zaire who suffer the usual tropical dangers of fever, leprosy, malaria, and parasites, but who select only three risks for worry and intervention-barrenness, bronchitis, and lighting.

In this brief survey, we take the perspective of a public servant with the responsibility to make decisions on matters of risk, and consequently can not choose only one of these approaches. All three have validity, and all define types of pressure with which the public policy process must contend. Therefore, we will deal in some measure with all three.

A practical approach to risk must recognize that not all risks are the same, and that depending on origin, nature of accompanying benefits and many other factors, public perceptions and appropriate policy responses may be greatly different for different risks. A useful taxonomy of risks by source was proposed by the American National Research Council committee on Risk and Decision-making (1982), sorting risks into seven classes:

1) Self-imposed risks, e.g., smoking or hang-gliding;

2) Co-Generated Risks, or risks imposed by the combined actions of two or more parties on one of the parties, e.g., occupational risks, consumer product risks;

3) Risks from Production Externalities, e.g., pollution, transport of hazardous goods;

4) Risks imposed by particular individuals on others, e.g., crime, smoking in public, drunken driving;

5) Risks from Nature, e.g., earthquakes, floods;

6) Risks from Economic conditions, e.g., poverty-induced diseases;

7) Risks from Government Policies, e.g., national energy policies, risk of war.

These categories are not without overlap; for example, risks from nature can be reduced or exacerbated by individual or public policy decisions. Risks of flood may be decreased by building dams but increased by individuals choosing to live in flood plains. Government policy may also create economic conditions that impose risks on some, or encourage the imposition of risks by production process externalities. Clearly government policy is involved in most types of risk, either as initiator,, regulator, or mediator. The merit of taxonomies such as this one is not that they define clear separation of classes, but that they illustrate how people's reactions may be different to different kinds of risk that are similar in magnitude, and how the responsibility of government may be different for different classes of risks.

We consider the issue of risk from several viewpoints. We begin in Chapter 1 with attempts to define and analyze objective risk, and a summary of analytical techniques to make decisions on objective risk. In Chapter 2, we summarize the train of psychological research on individual perceptions of risks that suggests that the analytical approach may be prescribing optimal decisions for mythical beings. Chapter 3 considers other aspects of collective risk decisions that makes them difficult. Chapter 4 suggests an appropriate role for analysis, and describes policy responses that are available. Finally, Chapter 5 takes a broader view and considers some intrinsic limitations on any attempt to deal with risk and some principles that might guide public policy in an unavoidably risky world.

CHAPTER 1: THE MECHANICS OF RISK-BENEFIT ANALYSIS

Introduction

There is an abundance of similar names often used interchangeably to refer to similar but distinct techniques for analytic study of risks: Risk Analysis, Risk-Benefit Analysis, Risk-Cost-Benefit Analysis, Risk Assessment, Quantitative Risk Assessment, and Probabilistic Risk Assessment. When these names refer to different techniques, the techniques normally are used for different classes of decision problems, for example for individual vs collective decisions, or for different stages of the same decision problem. Public policy decisions on risk are necessarily made collectively on behalf of a group of people. Any risk decision has the properties that the outcomes depend on uncertain and uncontrollable factors, and the outcomes include unwanted events -- loss of life, injury, disease, environmental damage, or economic losses. Public policy on risk is particularly concerned with physical risk, outcomes involving harm to life or health. We will refer generically to any analytical technique that addresses problems of this kind - collective decision problems involving the possibility of harm to life and limb - as Risk-Benefit Analysis, or RBA.

While it could be argued that every public policy decision is of this kind, analytic techniques for analyzing risk are potentially most useful when both the uncertainty and the magnitude of possible losses are large. In such cases, the intuitive processes by which risky decisions are made in everyday life are overwhelmed by the magnitude of uncertainty and the importance of the consequences. Any form of Risk-Benefit Analysis represents an attempt to make such obscure decision problems more transparent. This is accomplished by enumerating the possible choices and the outcomes that can result from each choice, and by explicitly using estimates of probability to describe as best we can the dimensions of the uncertainty.

Risk-Benefit Analysis has two heritages - the techniques for analytic study of collective decisions under certainty, known principally as Cost-Benefit Analysis, and the techniques for analytic study of individual decisions under uncertainty, known principally as Decision Analysis.

Cost-Benefit Analysis

Cost-Benefit Analysis attempts to measure and compare the economic gains and losses to society as a whole that result from an investment, a program, or a regulation. It bears some resemblance to the accounting techniques that are used in private enterprise to evaluate the profitability of proposed investments, but differs form them in that it attempts to count benefits and costs to all parties, not merely the return to the investors on invested funds. It is consequently concerned with aggregating benefits and costs that accrue to a large number of different people and agents. It is concerned with social totals, but normally ignores the distributive effects of a decision, the issues of who receives the benefits and who bears the costs.

There are two principal differences in procedure that result from this aggregate focus. First, the benefits to society brought about by the provision of a particular good or service can not be adequately represented simply by the revenue that will flow to the provider of the good. This is particularly evident for goods that have no market or that are provided free of charge. Clearly the provision of a new public highway brings some benefits, but if people travel on the highway and receive those benefits at no charge, then the revenue from the highway will be zero. The benefits must be measured, rather, by what people would be willing to pay to receive them if they had to do so. This quantity, the consumer's surplus, can be estimated if the demand curve for the good is known.

Second, the total social costs associated with a project may be different from the costs that would be perceived by a private operator. The provision of a good may create external costs that are imposed on the community but are outside markets. If the smoke from a new pulp mill impairs the health of people in the community and reduces the productivity of nearby vegetable farms, then the mill has created a negative externality. There can also be positive externalities, when the operation of an enterprise creates benefits that accrue to others and for which there is no market. While a private operator would not count these externalities as costs or benefits, they clearly should be counted from a social standpoint, and included in any collective reckoning of the total net desirability of a project.

These two differences can render a proposed project either more or less attractive than would be calculated by a private operator. There are many kinds of projects with large capital costs and low operating costs, such as bridges, tunnels, and railroads for which it may be impossible that a private operator could make enough revenue to cover expenses, but for which the social benefits nevertheless exceed the costs. On the other hand, a project damaging the environment, impairing the health of residents of the community, or imposing other negative externalities might be calculated collectively to be not worth its total costs even though in itself it offers large profits to its private operators.

The treatment of uncertainty in Cost-Benefit Analysis is usually very simple. Normally, quantities that are unknown or are subject to random fluctuation are estimated by a single best guess. When an unknown quantity is of great influence on the result, its impact is normally estimated by a Sensitivity Analysis, in which the analysis is repeated using other possible values for the unknown quantity, values which are felt to span the range of reasonable possibilities. While this procedure can indicate the sensitivity of the outcome to the value of the unknown quantity, it does not allow any for the incorporation of any explicit information about the probability distribution of that quantity.

There is some controversy over what kinds of costs and benefits should be included in the analysis. While all effects that are valued by anyone in society ought in principle to be counted, some - such as preservation of historic sites, scenic beauty, and community pride - are very difficult to value economically. While some analysts argue that to include estimates for such quantities in the analysis obscures the parts that can be estimated accurately (Mishan, 1975), others reply that to exclude non-economic values from quantification is to encourage their omission from consideration in the decision and so defeats much of the intent of a societal level analysis (Shrader-Frechette 1985).

Elementary treatments of CBA can be found in Stokey and Zeckhauser (1978) or Mishan (1975). Somewhat more detailed expositions are in Dasgupta and Pearce (1972), Layard (1972), and Lesourne (1975).

<u>Decision Analysis</u>

Decision Analysis, the other parent of Risk-Benefit Analysis, is a technique to facilitate the explicit consideration of uncertainty in individual decisions. Because it refers to an individual decision-maker, the difficult problems of aggregating benefits and costs across many agents do not arise. It is assumed that the decision-maker has a well-behaved preference ranking of every possible outcome, one that can be represented by a utility function. The utility representation also avoids the difficulty of comparing outcomes that differ in many attributes. The decision-maker is assumed able to weigh all economic, security, ethical, and aesthetic considerations in establishing his or her preference ordering.

Decision Analysis conceptually separates the choices that can be made from the outcomes that result, and often represents this separation by drawing the decision problem in the form of a tree. Beginning with the first decision that must be made, the tree branches each time one of several possible choices must be made, and again each time an uncertain event occurs. This tree formalism posts that it is possible to represent each choice as the selection of one from a finite set of possibilities, and every uncertain event as the occurrence of precisely one of a finite set of possible events. When either the possible choices or the possible events are continuously variable, it may be appropriate to represent the infinite variety of possibilities by a finite set of representative choices or events that spans the relevant range. It is also assumed that to each possible event can be assigned a unique probability (such that at each junction, the sum of all probabilities is one), either objectively obtained or representing the decision-maker's subjective feelings of likelihood. Elementary treatments of Decision Analysis can be found in Raiffa (1968), Schlaiffer (1969), and Howard et al (1976).

The required behaviour of the individual's utility function is based on several fundamental assumptions about individual preferences that are taken to define rational behaviour under uncertainty. Slightly differing statements of these can be found in Von Neumann and Morgenstern (1947), Savage (1972), and Marschak (1950). the fundamental axioms are:

1) Complete Ordering: Given any two prospects (sets of outcomes and probabilities of obtaining them) A and B, the decision-maker can state unambiguously whether A is preferred, B is preferred, or they are indifferent.

2) Transitivity: If A is preferred to B and B is preferred to C, then A must be preferred to C.

3) Continuity: If A is preferred to B which is preferred to C, then there is some probability combination of A and C that is indifferent to B.

4) Independence: If A is indifferent to B, then some probability combination of A, with any other outcome X is indifferent to the same probability combination of B with X.

These axioms are widely accepted as validly representing rational behaviour, although there is some controversy over the validity of the Continuity and Independence axioms for certain extreme outcomes. These axioms are sufficient to permit derivation of the Expected Utility Theorem. It is this theorem that allows the existence of a utility function as required in Decision Analysis. It states that when choosing between risky prospects specified in terms of the probabilities of various outcomes, a rational agent will choose the one with the highest expected utility. The essential elements of this derivation are that only the probability and consequence of each outcome matters, and that preferences are linear in probability, which means that 20% chance of a desired event is exactly twice as desirable as a 10% chance of the same event.

Risk-Benefit Analysis: Probabilities

Risk-Benefit Analysis (RBA) combines these two approaches; it attempts to represent uncertainty explicitly by describing uncertain events in terms of probabilities, and it considers socially aggregated measures of risk, cost, and benefit. Every possible outcome of a decision is described by its probability of occurring and the consequences of its occurring. The process is very complex. In the first place, there is no social utility function corresponding to the individual's in Decision Analysis. Consequently, there are many possible ways to measure risk consequences - total lives lost, days of lost-time injury or illness, economic effects, environmental effects, or some measure that combines two or ore of these. Moreover, there is some ambiguity in the use of the terms "outcome" and "consequences". An outcome can be defined directly in terms of the consequences of interest; an outcome could, for example, be 10 deaths per year. In this case, there is no need for further analysis to describe the consequences. Outcomes are often defined, though, in terms of some observable intermediate events, such as a certain level of public exposure to a hazardous material. In this case, determining the consequences of the event in terms of the public health variables of interest may represent a large amount of further analytical work.

The probability of a particular event occurring is a numerical measure indicating how likely it is to occur. Although modern theories of sub-atomic physics may assert that events at that scale are fundamentally random and unpredictable, we normally regard the events that we observe around us as causal, or predictably determined in principle. Describing an event merely by assigning a probability to it is not then attributing an essential randomness to the event, but merely characterizing out inadequate partial knowledge of the factors leading to the event. Not having enough knowledge to state with certainty whether an event will happen or not, we use a probability measure to delimit our ignorance.

There are three approaches to the measurement of probabilities, each of which is appropriate for different classes of events. The relative frequency approach refers to events which have been or could be repeated a large number of times, with the numbers of occurrences of each possible outcome counted. The probability of any particular outcome is estimated to be equal to the fraction of all past similar events in which the specified outcome has occurred. If during the last 1,000 days of my driving to work on the same highway at 8:05 a.m., there have been 200 days on which the highway has been extremely congested, then by the relative frequency approach I would estimate the probability of congestion today as 20%.

The classical approach to probability estimation involves looking for fundamental symmetries in the random process being examined, and exploiting these to assign equal probabilities to equivalent possible outcomes. This approach is most useful for estimating probabilities in gambling or in certain experimental processes. Through this approach we say that the probability of any one face on a die turning up is one-sixty. Similarly, if we draw (fairly!) a name out of a hat when the hat contains 20 names, we estimate the probability of each name being drawn as one-twentieth.

The subjective approach is the one used when neither of the other approaches is applicable. When neither repeatibility nor symmetry can suggest a clear theoretical or empirical basis for assignment of a probability, a subjective estimate can still be made. I might, for example, estimate the probability of a proposed new business venture turning a profit as 60%. When there is the opportunity to undertake sampling or experiments to refine the estimate of a probability, the subjective approach can be used to estimate starting points which are then revised to incorporate the information contained in the experimental results by a process called Bayesian revision of probabilities, the heart of mainstream decision theory.

In risk analysis, it is principally the relative frequency and subjective approaches that are used to estimate probabilities. The relative frequency approach is used when the event whose probability is required is essentially similar to other events which have been observed many times in the past. For example, if risk calculations are being performed for the construction of a new highway, the probability of a fatal highway accident per passenger-mile travelled might reliably be estimated by examining the record of past fatalities on similar stretches of highway. Similarly, in performing risk calculations for a new proposed steel mill, it may be reasonable to estimate the probability of an occupational death per person-day by referring to experience in similar facilities in the past. In applying this approach it is assumed that the new event being examined is sufficiently similar to the past events for which the frequencies have been observed, and that any trends present in the historical data have been adequately incorporate in the estimate.

When considering events without a large body of historical precedents, the relative frequency approach can still be used in some cases. To estimate the probabilities of failures or accidents in new and complex engineered systems such as LNG terminals or nuclear power plants, the relative frequency approach can be used in conjunction with the rules of combining probabilities to reason from the total system failure in question, for which there may be insufficient historical evidence, to failures in the fundamental components that make up the system, for which historical data of sufficient similarity may be available.

The principal technique used for this estimation process is called Fault-tree analysis. It is used to study the genesis of a specified total-system failure, by using a tree structure to represent the logical relations between the smaller components that make up the total system. One begins by supposing that a particular system failure has happened, and asks what events must have happened in order for the failure to happen; then, what events must have happened in order for these events to happen. Proceeding further and further back through the tree, one is examining an ever-increasing number of ever-simpler failure events and the logical relationships between them. Eventually, the process reaches so-called primary events, events which are not analyzed in more detail but for which reasonable estimates of probabilities are assumed to be directly available. Examples of primary events would be a particular switch failing to open, a valve sticking closed, etc. These events are sufficiently simple that it is assumed that the large amount of operating experience available on similar components in other kinds of systems is sufficient basis for using relative frequency to estimate probabilities. Fig. 1 shows a simple fault tree that could be used to estimate the probability of a fluorescent light fixture failing (after McCormick).

SPACE FOR FIG 1

There are two crucial assumptions in this approach when it is used to estimate the total risk associated with a particular facility. First, it is assumed that the list of total system failures for which fault trees are drawn represents all the significant risks associated with the system. In other words, it is assumed that every major accident that can happen has been identified. Second, it is assumed that the probabilities of failure of the primary events are independent of each other: the occurrence of one event has no effect on the probability of the other. If two pumps are in different factories and are of different manufacture, then their probabilities of failure are likely independent. If two pumps operate in series on the same system, then their probabilities of failure may well not be independent, for the failure of one may overload the other and render it more likely to fail. There is some opportunity to incorporate non-independent failure events into a fault-tree analysis through the use of so-called common mode failures. These are events such as fires, explosions, or operator errors that may cause several simultaneous failures or alter the probabilities of failure of several components. Because of the computational complexity involved in the alternative, the assumption of independence is usually made in practice. Introductions to Fault-Tree Analysis can be found in McCormick (1981), Rasmussen (1981), and the University of Waterloo (1984).

If the event whose probability must be estimated does not have enough historical antecedents to use the relative frequency approach and can not be decomposed into a collection of more basic technical events, then there is usually no alternative to using subjective estimates of probability. such events would include a terrorist attack on a nuclear reactor, an attack in war on a LNG terminal, or the establishment of an effective cartel in an essential natural resource. As noted above, it may be possible to begin with a completely subjective estimate of a probability and refine it through evidence or experiment. It may also be possible to obtain more reliable subjective probability estimates by pooling the estimates of several knowledgeable people rather than using just one individual's estimate.

Consequences

If the event being studied is a specified release of a toxic substance, then estimating the consequences requires modelling the transport and dispersion of the hazard through air, water, and soil. The dispersion of a hazardous substance depends on the speed and direction of the wind or current, the chemical and physical properties of the substance, and the quantity emitted. For both air and water dispersions, the most common model used is the Gaussian plume model (Lamarsh 1975, Ch 11). In order to estimate human exposure, it is also necessary to estimate the number of people at various distances from the source of the hazard, and how sheltered they will be.

If the event being studied is a particular level of human exposure to a hazard, the health consequences are still uncertain; they depend on the mechanism of exposure and the often unknown toxicity or carcinogenicity of the hazard. For some materials such as coal dust or asbestos, there is enough historical evidence of a wide range of human exposure levels that the consequences of a particular level of exposure can be estimated quite reliably. But for most substances of potential hazard, there is (fortunately) not so much evidence of past exposure available. For such substances, it is necessary to assemble piecemeal evidence from a number of sources, and assume a particular shape for the graph of the relationship between exposure and response.

The most frequent shape assumed for the dose-response curve is linearity at low dose levels; an exposure half as great carries half the risk of harm. It is widely agreed that this assumption is conservative, in that it over-estimates the magnitude of harm. The computational advantage of this approach is that it vastly simplifies the calculation of total risk when the risk is spread over a large number of people. If the probability of harm is linear in the exposure, then a specified amount of the hazardous substance will cause the same number of illnesses or deaths whether spread uniformly or non-uniformly over a large number of people.

The normal alternative to the linearity hypothesis is that there exists a threshold level of exposure-a non-zero level below which the probability of harm is zero. In this case, the total consequences of a specified amount of hazard depend on the dispersion, for if it is dispersed finely enough the total resultant harm will be zero.

The data available with which to calibrate the dose-response curve, and so put a numerical estimate on total risk, come principally from three sources-observed or theoretical biochemical effect, animal bio-assay experiments, and human epidemiological studies.

When a substance can be observed in the laboratory to disrupt important human biochemical pathways, then that constitutes clear evidence that it is a risk to health. Such clear evidence, though, is normally only available for the most severe hazards. A slightly less crude level of testing s available from observing the mutagenic effects of a substance on simple, rapidly-reproducing organisms such as bacteria, or cultured cells. It is generally assumed that substances which are mutagenic are also carcinogenic. Such tests are fast and inexpensive to perform, but still provide a rather rough indication of the existence or degree of hazard to humans.

Information about human health hazards is most often obtained through animal bio-assay experiments, in which large numbers of animals are exposed to high doses of substances of suspected hazard. Such experiments are costly, difficult, and time-consuming to perform, but yield ore reliable information than experiments on simple organisms. A typical experiment would involve obtaining two large groups of animals (often rats) of nearly identical stock, and exposing one group to doses as high as possible of the suspected hazard, beginning exposure in utero when possible. The animals are then dissected at death to determine cause of death, and to seeks any significant differences in abnormalities between the two groups such as tumors or organ damage. (Lave 1982 Ch 2, Tuohy and Trebilcock 1982 Ch 2).

Animal bio-assay experiments are often assumed to be quite reliable indicators of the possibility of human health effects. The US FDA regulations prohibit the licensing of any food additive that is shown to cause cancer in animals. The difficulties with this source of information involve the extrapolation from large doses to small ones, and the extrapolation from animals to humans. Exposing the animals to enormous doses is necessary in order to observe statistically significant effects in a reasonable time, but the doses are truly extreme in terms of what a human would have to ingest to obtain an equivalent dose. Consequently, the inference to the doses normally received by humans depends very strongly on the assumed shape of the dose-response curve in the low-dose regions. If linearity is assumed when a substance in fact has a rather high threshold, then a major human health risk could be inferred when in fact there is none at the prevailing dose levels. The assumption of linearity, on the other hand, is considered a responsible over-estimation of hazard, particularly when considering substances that may show long-term or cumulative effects that can not be detected by the bio-assay experiments.

The generalization of results from animals to humans is also a cause of major controversy. When an experiment on animals finds a significant hazard, as in the Canadian experiments that detected carcinogenicity of saccharin in rats, it is often claimed that humans are less sensitive to the hazard than the experimental animals. When, on the other hand, a major bio-assay experiment fails to detect a significant hazard in animals, there are always those ready to argue that humans are more sensitive than the animals.

When the risk being studied is one to which people have had varying degrees of historical exposure, it is possible to estimate the magnitude of the hazard by sophisticated epidemiological studies that seek statistically significant differences in the incidence of health damage between two groups of people who have been exposed to different levels of the hazard. Ideally, for such a study to yield significant results, we would wish to have two large groups of people, identical in every significant respect except for their exposure. This situation rarely occurs naturally, and to create it artificially by intentional exposure of one group to the suspected hazard would be a violation of basic ethical principles. Consequently, most epidemiological studies are thwarted by noise in the data. It is usually very difficult to find two large groups whose only significant difference is their exposure to one particular hazard, and, in the case of long-term exposure, it can also be most difficult to determine after the fact what people's exposure levels were. Consequently, epidemiological studies are often plagued with methodological difficulties and subject to endless challenge. Their great strength is in detecting hazards of great specificity and large impact. The effects of asbestos in causing the otherwise rare lung-lining cancer mesothelioma were detect3ed by epidemiological studies of asbestos workers. The now universally accepted studies implicating smoking in lunch cancer and heart disease are epidemiological. (Miller, 1982).

The two stages outlined thus far-estimation of probabilities of events and estimation of consequences of events- will be common to any attempt to analyze collective risk; the relative importance and difficulty of each stage will vary greatly depending on the issue being addressed, however. For example, a risk analysis to be used in selecting a regulated exposure limit for workers in a chemical factory will focus almost entirely on the consequence of various proposed exposure levels, while a study no the risk to the surrounding community of a pipeline pumping station will focus primarily on the probabilities of various sizes of explosions and fires.

Beyond this point, there is no standard form for the analysis or presentation of results. Several key decisions must be made, and will depend strongly on the type of question that is being addressed.

The group of people for whom risks are added up must be determined. This will normally be all people affected by the decision, but in certain instances separate risk calculations will be performed for sub-groups, e.g., for workers exposed to occupational risks higher than those to which the general public is exposed.

When a risk analysis is used to compare various ways of achieving a given social end, crucial decisions must be made about the basis of comparison of the different options. These decisions will be more controversial and more subject to different views as the options considered become more dissimilar. When comparing a nuclear thermal generating station to a coal-fired generating station, the relative risk of the two may depend strongly on the assumptions made regarding disposal of radioactive wastes, but the two options are sufficiently similar to permit some meaningful comparison of risks. In considering more dissimilar technical options, particularly when the choice of one or other of the technical options may imply other significant technical or economic changes in other sectors of the economy, then establishing a reasonable basis of comparison becomes even more difficult. A striking example is the controversy over the Inhaber report's treatment of the backup and storage requirements of non-conventional energy systems. (Inhaber 1982, Holdren et al 1979).

Several different levels of presentation of the results of a risk analysis are possible. The simplest level would be to present only the total estimated loss of life or health. A more comprehensive analysis could attempt to estimate other effects-effects on property value, and economic benefits and costs- and present these along with the calculated risks. If economic effects are calculated along with risks, then the two kinds of effects can either be presented separately, or an attempt can be made to aggregate them onto a common scale. While such an approach makes the results more compact and easier to compare, it requires establishing a constant of proportionality between dollars, days of injury or illness, and lives lost. This is equivalent to assigning a monetary value to a human lie, a controversial activity which will be discussed in Chapter 3.

Introductions to RBA can be found in McCormick (1981), Crouch and Wilson (1982), and University of Waterloo (1984). Discussions of the basis and role of RBA are included in Starr et al (1976), Shrecker (1984), Weinberg (1981), and Starr and Whipple (1981). Lave (1982), Crouch and Wilson (1982), Burke et al (1982), and Kunreuther et al (1982) present case studies of Risk-Benefit Analyses. Criticisms of the technique can be found in Lovins (1979), Hapgood (1979), and Green (1975).

CHAPTER 2: RISK-TAKING

The attractions of the analytic approach to risk decisions are that it provides a consistent transparent framework for the integration f all the information available into an estimated value for total risk. Its use is now officially supported for policy decisions in both the Canadian and American governments. Unfortunately, there is abundant evidence that people do not act as this theory says they should. In cognitive psychology there is a research tradition that begins with observations of real people assessing uncertainty and making risky choices, rather than with axioms about how ideally rational people ought to do so; this work has led to a body of theory significantly at variance with the axioms at the root of RBA. The ways that people appear to violate this code of behaviour are not minor and haphazard, but large and systematic.

The most prominent workers in this field of research have been the psychologists Daniel Kahneman of UBC and Amos Tversky of Stanford. They have performed a large number of experiments in which they ask people to estimate probabilities or to make choices between risky propositions when the probabilities are given. Again and again they have found that even the most basic requirements of rational decision-making on risk are violated by people's real choices.

Most crushing for the normative theory has been their observation that preferences depend strongly on the "framing" of the problem, the way in which risky prospects are expressed. By altering the language used to describe prospects, whether they are expressed in terms of chance of winning or chance of losing, or the implied context, people can be induced to reverse their choices between two prospects unchanged in substance. For example, they find that most people, if forced to choose between a certain loss of $50 and a 25% chance of losing $200, prefer to chance the $200 loss; but when the certain loss of $50 is called an insurance premium, most people prefer it.

In another experiment more directly related to risks to life and limb, they asked a large number of physicians whether they would recommend treating lung cancer in a 45-year-old man with radiation or with surgery. To one group they showed the probabilities of dying during treatment and of dying within 5 years-e.g., with radiation you have a 0% chance of dying during treatment but a 78% chance of dying within 5 years, while with surgery you have a 10% chance of dying on the operating table but only a 66% chance of dying within 5 years. To the other group they showed the same results expressed in terms of probability of survival, e.g., 100% chance of surviving radiation treatment, etc. The shocking result was that the fraction of physicians preferring radiation to surgery was strongly dependent on which of these two forms of expression was used. Only 16% of the doctors recommended radiation when the figures were expressed in terms of survival, while 50% recommended it when given mortality figures. This influence of framing is particularly remarkable when it is considered that these are experts making a professional judgement on a matter of life and death. A summary of their findings on framing can be found in Tversky and Kahneman (1981).

Another area in which experimental evidence indicates systematic departures from rationality is the evaluation of probabilities. It seems that people cannot handle small probabilities, say of one in a thousand or smaller. Appraisals involving such small probabilities appear to show a strong threshold effect. Until something brings it to our attention, we tend to ignore an event with such a low probability, that is we treat it as if its probability were zero. But once we have acknowledged it as possible, we then greatly over-estimate its probability. An example related to risk is the widely observed tendency of people to over-estimate the probability of certain rare and dramatic causes of death, such as botulims or radiation exposure. This phenomenon has been demonstrated by Slovic, Fischhoff and Lichtenstein (1979), who asked people to estimate the number of fatalities per year in the US due to various causes, after informing them that 50,000 per year die in auto accidents. Their results are shown in Fig. 2; rare deaths are over-estimated by factors as high as several hundred-botulims, for example, which kills 2 people per year in the US is estimated to kill 500- while the most frequent causes of death are all underestimated. The same bias in a different context has been found by Bill Ziemba of UBC in his studies of betting at racetracks. He finds that people bet too often on long-shots and not often enough on favourites, and has exploited, it is said, this and other inefficiencies in the market to develop an effective betting system.

FIG.2

The Anchoring heuristic is a tendency for people to select an initial value for their estimate of a probability or other unknown quantity, and then incorporate further information by making small adjustments to their initial value. The initial guess continues to exercise a significant influence on the final estimate, even in the presence of strong opposing evidence, so the estimate is said to be “anchored”. Anchoring can be demonstrated experimentally by giving various hints or initial estimates to people and observing how their final guess is influenced by the starting point.

The anchoring heuristic can be regarded as a tendency to place too much confidence in the first guess, and consequently is related to the often-observed over-confidence that lay subjects and experts alike place in their estimates of unknown quantities. In an experimental demonstration of this overconfidence, researchers at MIT gave seven engineers all the relevant technical information about an earthen embankment, then asked them to estimate the height at which it would fail. Two of the engineers substantially over-estimated and five underestimated the height of failure, but more remarkably, not one gave a confidence interval wide enough to contain the true (experimentally measured) value.

It is also related to the illusion of personal immunity that influences people's estimates of the probability of being injured or killed in automobile accidents or by power tools. Svenson (1978) reports that 80% of all subjects report themselves to be less likely than average to be injured in automobile accidents, while Rethans (1979) reports similar results in people's estimates of the likelihood of being injured by the tools and appliances they use. Availability may also be at work in this illusion, because the incidence of automobile accidents is sufficiently low individually that even the worst of drivers have abundant available evidence of their driving without accidents, despite tailgating, speeding, and other driving habits that may in fact drastically increase their risk.

Kahneman and Tversky have developed a consistent descriptive theory of people's choices that accommodates these results, called Prospect Theory. This theory, although not based on any axioms of rationality, is a powerful predictor of people's choices. The two theories of people's decision-making under risk, the normative Bayesian decision analysis and the predictive prospect theory, are unquestionably in conflict. People do not act as the theory of rational choice says they should. Furthermore, it is unlikely that any other theory will be found to reconcile the two, for any theory to be descriptively valid must reflect the dependence of people's preference s on framing, but surely no normative theory could be susceptible to framing. That is, no rational decision process could let the decision rest on the way it is described rather than on the objective outcome.

Economists may argue that the foregoing results only suffice to demonstrate that individuals do not choose rationally, while the validity of policy prescriptions based on assumptions of rationality only depends on collective or market rationality. Further, they may then argue that markets are rational even if the individuals who make up the markets are not, either because individual irrationalities are randomly distributed and sum to zero, or because the more rational individuals dominate markets by more aggressive trading. These premises have recently been subject to some empirical examination, and the results are not encouraging. Arrow (1982) reports that the prices of long-term and short-term securities do not exhibit the relationships that would be required for inter-temporal rationality. The results of Ziemba in horse racing are similarly based on the detection of significant inefficiencies or irrationalities in markets. More recently, Camerer (1986) has conducted experiments in which groups of people play at making a market in a commodity whose value depends on the outcome of an uncertain event. The experiment is structured so as to generate a different equilibrium price depending on whether the representativeness or a rational Bayesian approach is adopted to estimate the probability.

The results strongly support the contention that the representativeness heuristic operates as strongly at the market level as at the individual level.

This conflict between rationality and people's observed behaviour must, however, be treated with caution, particularly if one's reaction is to conclude that people do not know what is best for them and need to have their risk choices made for them by scientists and experts. The first caution is that scientists and experts appear to be as prone to these errors of judgement as anyone else. Much of Tversky and Kahneman's experimental work was performed with scientists and physicians as subjects, and the biases and heuristic devices described here have been observed to be employed by lay people, by scientists making judgements outside their area of expertise, and by scientists in their own fields when the data presented are insufficient to calculate a precise solution. Moreover, the problem of over-confidence in predictions has been observed to be the most pronounced in the judgements of experts. We are clearly all prey to these biases, errors, and misperceptions.

At the individual level, the conflict between these two models of choice merely suggests that if we wish to act like economically rational people, we must police our own personal decisions on risky matters very carefully. But for the public servant with official responsibility for making decisions that will impose risks on others, the conflict makes for a profound dilemma: should you choose for others as they would choose for themselves, knowing that such a choice does not represent their best interests, or should you choose as an ideally rational person would, knowing that the affected people may oppose the choice that you are making on their behalf?

Amos Tversky (1986), one of the architects of Prospect Theory mentioned above, insists that this theory should not be taken as a guide for decisions on behalf of others, but that in such contexts only the normative economic approach makes any sense. He compares his theory to the theories of perception that describe how people are susceptible to certain kinds of visual optical illusions, and argues that we should no more make prospect theory the basis of public policy on risk than we should make the theory of visual illusions the basis of our system of measurement of lengths.

This view is in some measure supported by experimental evidence that when people receive feedback and instruction in response to their naïve choices, they tend to raise their choices and estimates to resemble more closely those of a Bayesian decision-maker. The more transparent the decision problem is, the more likely are people to follow the axioms of rationality; the more obscure it is, the more likely they are to apply heuristics. (Kahneman, Slovic, Tversky, 1981 pp. 28-34). This suggests that there may be informed support for public policy measures that reflect more strongly the normative Bayesian approach to risk than people's observed naïve preferences and estimations.

The importance of the framing of alternatives and the apparent correction of biases and mis-estimations that results from appropriate information feedback both stress the crucial importance of the public servant's role as educator. A principal challenge is to present risky policy prospects in as transparent a fashion as possible, in order that people's preferences reflect as closely as possible their best interests. Provision of information about the likely consequences of expressed preferences will also assist in revising preferences in the direction of rationality. The difficult issue that this argument does not address, though, is the relationship between educating the public and manipulating them. When there is conflict between the public's expressed preferences and the apparently rational estimates of the experts, it is first necessary to make sure that the experts are not in turn succumbing to biases or heuristics or over-confidence. How such conflicts can or ought to be resolved when they persist is a question of grave difficulty; the next section turns to a few observations on the subject.

CHAPTER 3-PROBLEMS OF COLLECTIVE CHOICE

The discrepancy between the perspectives described in Chapters 1 and 2 respectively is the first factor to give one pause in deciding how to approach a decision involving public risk. A straightforward analysis of total risks, based on the best information available at present, may suggest conclusions strongly opposed to those reflecting people's expressed preferences. But where risk is involved, there are other, equally formidable challenges to be overcome in forming sound public policy. None of these is a weakness unique to the analytical approach, but all give reason to doubt that a straightforward analysis will necessarily generate an optimal or even an acceptable decision. Some are intrinsic difficulties of risk decisions that any proposed decision-making technique must address, while some are inevitable accompaniments of any political decision on any subject.

Insufficient Information

Massive uncertainty prevails on important issues of fact relevant to many risk decisions. As discussed above, the total risk associated with a choice may depend on the probability of catastrophic accidents at a facility,, probabilities which are known to be small but for which there is no reliable way of telling how small-say, one in ten thousand or one in a million. But this difference can be of enormous consequence in estimating the total risk. Conversely, the risk may depend on the toxicity or carcinogenicity of a new chemical for which only crude estimates are available. Significant improvements n the information available may not be attainable, because one cannot experiment directly on people, and the limits in the reliability of bio-assay or epidemiological research are intrinsic to the process and can only be overcome by mounting experiments of staggering length, complexity, and expense. If animal rights activists prevail, even the imperfect source of information represented by bio-assay may cease to be available.

Perhaps, as Weinberg (1972) pointed out, some of the dilemmas we face reflect questions that can be formulated in scientific terms, but for which no forseeable scientific procedure can give an answer. In this case, the crucial public policy decision will be how to employ incomplete or partial evidence. This is often the case now, for it frequently happens that we have to make decisions before the evidence we need is in; Weinberg simply points out that for some questions we may never be able to obtain the evidence we desire. Decisions based on partial knowledge raise the issues of where the burden of proof should lie, and what standards of proof short of conclusiveness are required. These are traditionally questions that have been addressed by the judicial tradition, and Shrecker (1984, pp. 25-31) argues that the growing tendency for matters of risk to be shunted over to the legal system is resulting in an inappropriate application of these standards. He argue that traditionally the presumed innocence of a person charged of a crime has been applied to people or organizations desiring to impose risks or charged with imposing risks, and that this practice has resulted in insufficient caution being exercised in determining what risks are allowed. Decisions on burden of proof and standard of proof must consider the balancing of the two possible kinds of errors recognized in statistical hypothesis testing- type 1 errors and type 2 errors, or convicting the innocent and acquitting the guilty, respectively. Shrecker argues that the traditional rules of procedure correctly regard the conviction of an innocent person as the more grievous error, but that the severities of the two kinds of error can be so different in matters of risk that different rules are needed. (The search for new rules of procedure that employ the strengths of the judicial system but reflect the different decision environment is a theme of the Science Court proposals discussed in Chapter 4 below.)

When vast factual uncertainty prevails, there is ample room for experts in the field to disagree. Consequently, virtually any side of risk debate can find credible expert support, but non-scientists are likely to become disillusioned and confused by the sight of scientists arguing acrimoniously over issues of fact relating to nuclear power, toxic waste, etc. Manne and Richels (1980) asked a large number of American scientists their opinions on a variety of questions related to energy. They found that very strong clustering occurred between opinions on issues that appeared to be unrelated. For example, there was strong clustering between expectations of high future demand for electricity, high future uranium prices, and high prices for solar photovoltaic cells. The researchers were puzzled by this clustering until they found that these scientists generally supported the breeder reactor, whose feasibility depends on high electricity demand growth and high prices for the competing alternatives, including conventional uranium-fuelled nuclear power and solar photovoltaics. The scientists who opposed the breeder had a correspondingly strong clustering around the opposite extremes. In each case, a prior opinion on a political issue appears to have organized a cluster of factual opinions to support it.

Such phenomena can have two explanations. It could be that the vagueness of the evidence allows values to creep unnoticed into the process of forming scientific opinions. such an effect could occur in all innocence, and exert substantial unperceived bias on both opinion and scientific work. Alternatively, conscius knowledge of the important role expert voices can play in swaying political debate may lead one’s scientific opinion to come to be perceived as a means to further one’s advocacy of political goals.

If Weinberg is correct that the evidence necessary to reconcile expert dispute will not be forthcoming, then we can expect this dispute to continue. The anthropologists ((Wynne (1981), Thompson (1981 and 1982), Douglas & Wildavsky (1982)) go further still, and argue that people, whether expert or not, choose their sides on matters of public risk on the basis of culturally determined attributes of the social group to which they belong. Consequently, even were definitive evidence in fact to become available, the disputants likely would not be reconciled; rather, they would further elaborate their respective theories and ideologies to embrace the new evidence, and turn its refutation into support.

A final matter of great practical importance in determining magnitudes of risk is that the size of risk to which we are exposed normally depends on much more than issues of scientific fact. There is almost always human action or omission involved in realizing or preventing the realization of a risk. This is true of nuclear plant operators, railroad engineers, pilots and air traffic controllers, and also of private individuals who have the power to impose risks on us in their capacities as drivers, cyclists, victims of infectious diseases, and so on.

When such factors must be estimated in a risk calculation, a historical statistical measure is normally used. This may be valid if the sample (the number of people whose performance determines the risk) is large, but will certainly be subject to large fluctuation if the sample is small. Also, the use of such estimates for the frequency of human errors or a description of potential malice neglects the possibility of enhancing performance by providing clear information and feedback, designing jobs and environments so as to encourage alertness, etc. (A Marxist perspective might perhaps go even further, arguing that certain socio-political environments encourage aggression, depression, and irresponsibility, while others encourage alertness and responsibility.)

In a broad sense, the assumptions that are made about people's responsibility may be self-fulfilling. Starr (1985) makes this suggestion regarding people's political behaviour, but it can also be observed to operate in any workplace or school.

Comparing the Incomparables

Because the consequences of decisions are of several kinds-economic, hedonistic, risky, distributive-any decision, individual or collective, implies a trade-off involving disparate values. Specifically, any decision that weighs economic considerations and matters of risk implies a value for human life. It is possible to examine decisions in retrospect and estimate the value of life that they implied; studies such as this on public policy decisions normally find enormous discrepancies in the implied values, from thousands to billions.

Supporters of the analytic approach to risk decisions suggest that these discrepancies are unacceptable, and that a single value or a specified range of values ought to be used in all decisions. A number of writers, including Kluge (1986) and Hapgood (1979) oppose this view. They argue that it is irrelevant that an implied value can be attributed retrospectively, and that the explicit consideration of life in economic terms is repugnant, and must not be practiced. Hapgood (1979, p. 36) cites the case of a city's decision to close its playgrounds after a child dies in a fall from a swing. He argues that a strong official could legitimately resist pressure to close the playgrounds by saying "We're not going to make every child in the city suffer because once in a blue moon a tragedy happens.", but that an official who responds that the benefits of the playgrounds to all the city's residents justify the risks as long as less than, say, one child per year dies, is so detached from human feelings and values as to be unworthy of a position of public trust. This argument seems to reflect a longing for past innocence. It says in effect that decisions that balance life against economic factors can be tolerated only as long as the balancing is done implicitly; it is better that we do not know the precise tradeoffs involving human life that our decisions entail.

In reply to these arguments, an economist might say that the existence of a discrepancy in the monetary values assigned to human life means that society could move its safety expenditures around from one activity to another, and at no cost save lives that are lost at present. Surely the costless saving of lives is worth enduring some morally difficult deliberations. Moreover, even Hapgood's playground official must respond to the numerical evidence; there must be a threshold number of deaths at which he or she would reconsider.

If the notion of assigning a value to human life is accepted in principle, then the operative questions are what value is appropriate, and how to determine it. One earlier approach attempted to calculate the net present value of a person's future earnings. This approach is out of favour now, for several reasons; it seems to calculate the value of a person's productive services to society, not of one's own life to oneself; surely life is valuable beyond simple economic production. Moreover, this approach yields values that are very small or zero for people not producing in the economy; this blatantly unacceptable outcome provides a direct signal of the intellectual problems with the above procedure.

A preferred approach now assigns a value based on people's observed willingness to pay for small reductions in risk to themselves. Details can be found in Zeckhauser (1975), Schelling (1968), and Linnerooth (1982).

Empirical attempts to calculate values on this basis have examined labour markets, in which people accept riskier work for higher wages, and markets for consumer products, in which people have the option of spending more money to obtain safer products or safety devices. Studies in Thaler and Rosen (1976) and Rappaport (1981) used these methods to find estimated values of life ranging from $200,000 to $2,000,000.

While these studies correctly attempt to infer the value that individuals place on their own lives when deciding whether or not to take small risks, two factors call the results into question in more general applications. First, it is not clear that the labour market pays a premium for riskier jobs. While within some industries such as logging the riskiest jobs are the highest paid, there is no such correlation evident in the market as a whole; and even within a single industry it is difficult to separate the effect of risk on wages from that of skill. In many cases, workers may not have a broad enough range of employment options offered to them to express their risk/wage preference. Second, the evidence cited above about people's misperceptions of risky situations suggests that workers may not perceive the trade-offs they face clearly enough to make a choice that represents their true valuation of the risks they face.

A final difficult question in the valuation of lives is whether or not every life should be valued equally. Perhaps a young life should be valued more than an old one. Schelling (1968) argues that a strict willingness-to-pay approach would assign a higher value to a wealthy life than to a poor one. Or a Rawlsian ethical perspective might suggest that the lives most at risk should be valued the highest. This latter view clearly has much support in our feelings and current practice; our willingness to pay to reduce particular threats to our own lives becomes much greater as the magnitude of the threat we face increases, and collectively we are willing to pay much more-or even put additional lives at risk - to rescue specific individuals from peril than we are to avert in advance threats that would fall onto unknown individuals. A discussion of analytical issues involved in saving identifiable lives vs saving statistical lives can be found in Linnerrooth (1982).

Distribution and Equity Issues

Distributional effects are present in virtually all public policy decisions, and those involving risks are no exception. It is virtually never the case that the risks, costs, and benefits resulting from a decision are distributed perfectly equally, and consequently an aggregated investigation of total social risks and benefits does not answer all the necessary questions. We must look at the disaggregated effects and ask, should these people be made less safe so that those can be made safer? Or, should these people be made less safe so that those can be made wealthier?

These questions are similar to those that must be asked of any policy decision, but are more poignant because loss of life is involved. But because they are phrased in terms of inequity in prior risk, or probability of harm, they do not yet capture the full inequity of risk decisions. Whether prior probabilities are equal or unequal, the realized harm comes in large doses to only a few of those at risk. The probabilistic nature of risk thus ensures that the consequences of risk decisions are the most inequitable of any policy decisions.

If the clear best interest of all requires the imposition of risks that fall especially heavily on a few, or alternatively if the reduction of risks to all requires that costs or curtailment of rights be imposed on a few, then we must ask whose rights have a priority, those of society or those of the threatened individuals. Decision problems of both these kinds arise frequently. Examples of the former are the siting of dangerous facilities such as hazardous waste disposal sites, for such facilities are clearly needed by society, but are only located in one community, or the transportation of chlorine, ammonia, and propane by rail, for which the people who live near the tracks bear most of the risk. Examples of the latter are the curtailment of individual rights to privacy entailed by the compulsory reporting of medical conditions that represent risks to the public, as exists routinely for pilots of aircraft, and is now the subject of controversy for victims of AIDS.

If in a particular case the power to impose risks lies with society, and the risk is to be imposed on some individuals, then competition among “not in my backyard” interests may ensure that it is imposed on those ignorant of the risks or on the politically powerless. If the power to refuse risks lies with individuals, then the risk may well not be imposed and society as a whole will be worse off. But if the total social benefits really do exceed the costs and risks, then in some fashion it ought to be possible to strike a deal. In both of these cases, compensation can help; in the former case it is a matter of fairness, in the latter a matter of necessity if the risk is to be accepted. Some issues in compensation for imposed risks are discussed in Chapter 4.

The Decision Environment

It should finally be noted that public policy decisions on risk of course are made in the same political environment as any other public policy decisions. Consequently they are temporally complex, with many urgent items competing for space on the decision agenda of busy people; they must be based on pragmatic criteria such as organizational feasibility and implementability, as well as simple public risk and benefit; and they involve many actors who will compete with all of the means at their disposal to influence decisions in directions favourable to their own interests.

These familiar challenges are likely to be increased in intensity when the decision in question involves risk, however. The magnitude imposed inequity associated with risks means that the prospective winners and losers will press their cases with great vigour. The paucity of clear factual information that normally prevails means that it is possible for each contending faction to have distinguished scientific experts arguing on its side. And the fact that possible loss of life, and the weighing of lives with economic considerations are involved, means that the conflicts are both volatile and dramatic; in the terms of Chapter 2, they have high availability, and consequently attract great attention. There is normally prominent press coverage, and careers are made or broken on the outcomes. How to deal with it all?

CHAPTER 4: POLICY RESPONSES TO RISK ISSUES

The Role of Analysis

The challenges that beset the attempt to develop sound policy on matters of public risk are formidable indeed, but decisions must nevertheless be made. What guidance can be offered on how to approach decisions on matters of risk in a responsible manner, and what policy instruments can be employed to help lessen the impact of the problem discussed above?

Although the use of Risk-Benefit Analysis is subject to several forceful criticisms, it remains an attractive vehicle for the study of policy decisions on risk when it is compared fairly to the approaches to decision-making with which it competes. The alternative models of policy-making might be grouped roughly into the following categories: intuitive expert judgement, following the polls, or following established procedure.

RBA is indeed vulnerable to profound uncertainty in matters of probability and consequence that are necessary to establish estimates of total risk, but this is because the information is simply not available. Any approach to decision-making will suffer from the same lack of relevant information. The analytic approach has the advantage that it makes the assumptions about uncertain quantities explicit and the impact of the assumptions testable. It also permits explicit representation of the degree of uncertainty in each assumption and in final estimates of total risk. Any other approach would have to involve equally rough assumptions, but would make them implicitly in such a way that they might not be evident, and their impact on predicted outcomes could not be tested. An analytic approach can also assist in defining the crucial information needs of a decision, and can sometimes provide relevant new information by calling attention to significant sources of risk that would otherwise have been overlooked. Kunreuther et al (1982) cite a RBA for a proposed LNG terminal in the Federal Republic of Germany, which showed a significant risk from ship collision that had not been considered, and which thus led to a substantial reduction of the total risk associated with the project through the re-design of a navigation channel.

It is true that RBA conducted at the level of aggregate social risks, costs, and benefits ignores the distributive aspects of decisions. But this is the case as well with all of the alternative approaches, in which distributive issues can only be dealt with implicitly, or on the basis of the relative political power of the various affected groups. Moreover, it is not necessary that analysis of risk decisions only be conducted at the level of social aggregates; the same techniques can be used to estimate total risk to various affected groups, and so can be used to identify possible distributional effects more accurately than any informal procedure can. Disaggregated RBA'' are not frequently performed, because they are long, difficult, and expensive; but if the distributive aspects of a decision are important enough to warrant one, it can be done.

An RBA that seeks a single measure with which to compare alternative decisions does involve comparing the incomparables - weighing economic considerations against community values, scenic beauty, environmental purity, and human life and health. But as was argued in Chapter 3, any decision with effects on all of these factors implies a tradeoff between them. The reality of such tradeoffs can not be avoided in any such decisions. And in fact, it is only in full RBA's that such tradeoffs are made explicitly. It is also possible to use analytic techniques to estimate total risk only, and present total risk along with estimates of costs and benefits to the political process for an informal decision process. It is this kind of calculation that is normally implied by the use of the term Risk Assessment, rather than Risk Analysis. So the comparison of incommesurable values need not be made explicitly in an analytic approach to a risk decision. On the other hand, there may be advantages in doing so. As with the issues of factual uncertainty, adding up risks, costs, and benefits on a single scale requires that the concomitant value decisions be made explicitly, and allows one to investigate the sensitivity of the final recommended decision to changes in the value tradeoffs. An implicit decision process, while it may appear less heart-wrenching, renders the accompanying value judgements inaccessible to scrutiny and makes it impossible to determine their impact effect on the ultimate decision.

The most compelling arguments against the use of RBA concern its capacity for misuse; they suggest the need for certain standards of analysis and presentation to be maintained. The argument is that the flexibility available in every stage of RBA-in the range of choices considered and the basis for their comparison, in the estimates assigned to unknown scientific quantities and the values assigned to unmeasurables-makes it an ideal vehicle for powerful groups in society who will use it to support decisions favourable to their own interests. Furthermore, it is argued that the large systematic biases that can be inserted in an analysis can also be effectively concealed from scrutiny, giving the rigged results a semblance of impartial scientific authority. The critics differ on whether RBA is more likely to favour the interests of those groups who currently enjoy political power or of a new technocratic elite, but they agree that it is anathema to any open, democratic political process.

If RBA is not to be used thus, several issues in its implementation must be responsibly addressed. First is the appropriate breadth of its application. While the decision studied must be important enough to justify the expenditure of resources on an RBA, it must be one of limited breadth. RBA cannot tell what ends society should pursue, because the analysis can only be conducted when the details of approach as well as the goal are specified. It can only be used to estimate the total social benefit of a specified way of seeking a specified goal, or to compare different ways of pursuing a specified, presumably politically chose, goal.

The corollary of this restriction of application is that the significance of the results of RBA must not be over-stated. A favourable RBA does not necessarily mean that society should pursue the goal in question. If an RBA supports the construction of a particular kind of power plant in a particular location, this represents an endorsement only with respect to the alternatives that were considered; it does not indicate the need for such a plant, or indeed for any plant at all.

The possibility of RBA being used for obfuscation imposes a responsibility for clear communication on those performing the analysis. In order for an RBA to be a useful tool in policy debate, it must clearly indicate what problem was studied, what options were compared, and what assumptions were made for each option. All the assumptions made abut unknown scientific quantities, economic values, and value judgements must be clearly stated. When guesses must be made for unknown quantities, the analysts must endeavour to make the guesses as bias-free as possible. Raiffa (1980) has pointed out that this pursuit of freedom from conscious or unconscious biases should include the avoidance of "prudent" or "conservative" assumptions. Prudence, he argues, belongs at the decision-making stages, when the person responsible needs the most accurate possible information, not information that is already biased towards a "prudent" over-estimate of risks. Whenever possible, the impact of the assumptions made should be explained by repeating the analysis with different values and observing the resultant change in total estimated risk. This process, called sensitivity analysis, is particularly important when the quantity in question is disputed or subject to wide uncertainty, and of major impact on the total estimated risk and benefit.

Both Raiffa (1980) and Kunreuther et al (1982) stress the importance of clarity of reporting, possibly even encouraged through establishing a standardized format for reporting the results of RBA'S. This would be particularly valuable if, as Kunreuther et al suggest, major decisions can be studied by two or more groups performing simultaneous analyses. They also recommend that groups performing RBA's should publish their work in order to be subject to critical peer review. The hearing of well-organized opposing views can be encouraged by official funding of objector groups. Ackerman et al (1974) and the Science Council of Canada ( ) have both suggested the establishment of an independent expert body to perform official risk assessments, but Kunreuther et al's (1982) study points out that the problem of finding acceptable experts to sit on the board would be most difficult. These issues of the use of RBA in the decision process are closely related to the process questions, such as Science courts and public inquiries, that are discussed below.

As well as communicating clearly the assumptions that went into the analysis, a responsible analyst should attempt to compare alternatives on as even a basis as possible, making equivalent assumptions about economic and population growth, price of inputs, and rate of technological innovation. It is also important that a sufficiently broad range of options be considered, if possible spanning the entire range of relevant political debate. Paying attention to these caveats can help prevent the suspicion that RBA is being used as its critics charge, to lend apparently authoritative support to decisions made on other grounds. These suggestions can be effectively defused by making the analysis accessible to political debate. Some experiments of this kind are considered in the following section.

Processes of Dispute Resolution

A popular and useful way to study disputes over matters of risk is to split them into disputes over facts and disputes over values. This separation is more than a separation of phases or aspects of any dispute than a viable taxonomy of disputes. It may not be valid; while some disputes will no doubt be dominantly either disputes over facts or disputes over values and others will have both components but be separable, some will no doubt involve both facts and values in a combination that is inseparable. The assumptions behind the separation would be that there exist objective facts, that it is possible to obtain evidence to ascertain what they are, and that every competent and reasonable person will acknowledge the veracity of the facts if presented with the evidence.

It is on the basis of such a set of assumptions that one sees so many calls for better research to resolve the most pressing issues of factual dispute, in order that policy discussion can be focused strictly on value dispute. But the evidence of attempts to resolve specific heated issues of factual dispute relevant to public policy risk issues is not encouraging. A prominent example was the attempt of the Operations Research Society of America to resolve dispute over the need for and viability of the proposed ABM system by an independent evaluation of the research work that supported two conflicting expert viewpoints. The attempted objective evaluation became as mired in accusations of error, distortion and bias as the original pieces of work. Details can be found in Operations Research Society of America (1972) and Doty (1972).

The weaker of the two theoretical objections to this model of dispute resolution, and an effective explanation of the ABM experience, is found in Weinberg's (1972) notion of "trans-scientific questions". If definitive evidence to resolve a factual dispute can in principle be unattainable, then we may be stuck with making decisions on the balances of probabilities rather than facts. A number of suggestions have been made for decision-making structures to balance the evidence available when public policy decisions must be made on the basis of insufficient factual evidence. Principal among them are the Science Court proposal, and the proposal for formal scientific mediation.

The Science Court proposal is associated principally with the name of Arthur Kantrowitz. He has proposed an institution that would apply the judicial format of adversary presentations with cross-examination before a neutral expert judge to the resolution of factual disputes that are relevant to policy making. He suggests that subjecting the conflicting views to detailed examination by scientifically competent peers would allow the best possible judgement to be made about where the truth lies. The court would not be constrained to making categorical judgements as is a law court, but could give judgements reflecting the judge's perception of the balance of probabilities between the two contending cases. Judgements would always be regarded as provisional, because new evidence could appear at any time calling a prior decision into question. Details of the proposal can be found in Kantrowitz (1967, 1977).

Three types of criticisms of the science court proposal have been made. First, it has been argued that the use of the adversary format will not serve the discovery of truth. Lawyers have criticized the proposal saying that in the judicial system the adversary system is not designed necessarily to discover truth, but to resolve conflicts expediently without recourse to violence. In particular, it is alleged that relevant information that supports a view between those of the two adversaries may never be introduced. Since a science court has never been established, there exists some question over the degree to which judicial procedures would have to be adopted; one can imagine novel formats, such as permitting the introduction of evidence by third parties or the court itself, that might obviate some of the problems of the adversary system.

Since the idea of a science court is premised upon the supposed separability of facts and values, those who allege that these are inseparable most often attack the idea on that basis. Clearly, the selection and phrasing of the questions to be addressed by the court would be of grave importance in determining the degree of separability of facts. Neither the advocates nor the opponents of the proposal have often adduced candidate questions to support their cases. The procedural matter of who had the power to invoke the court and who would frame the questions to be considered would obviously be crucial in determining the success of the venture.

Other procedural difficulties concern who would judge the proceedings and who would be chosen to represent each side. It has been suggested that eminent scientists would be reluctant to participate because of possible damage to their reputations, particularly if there were stigma associated with fighting a losing case, and that the amount of time required to prepare of hear a case without significant scholarly rewards would discourage all but the most shallowly partisan from participating. Criticisms of the proposal can be found in Shrader-Frechette (1985), Boffey (1976), Callen (1976), Mazzur (1977), and Green (1977). A jurist's view of the appropriate role of the courts in the resolution of science disputes is set out in Bazelon (1979, 1980).

An alternative proposal is found in the formalized mediation proposal of Abrams and Berry (1977). They suggest that representatives of the two sides of a factual dispute meet with a scientifically competent mediator to attempt to elaborate the nature of their disagreement and attempt to find ways that they can come to agreement. The objective of the venture would be the production of a jointly-authored paper describing the original positions and degree of resolution achieved, if any. They argue that the non-adversarial framework allows a more constructive environment for the pursuit of truth, and would encourage the introduction of a broader range of material relevant to the dispute. Possible problems with this approach include the determination of who would have standing to represent a side, the possibility of co-option of representatives from their constituencies by a skilful mediator, and the question raised above for the science court of how the questions to be addressed would be select and what the official status of result would be.

The stronger objection to this model based on a separation of factual disputes and value disputes is that raised by the anthropological research mentioned above in Chapter 3. Those writers argue that the resolution of factual disputes on contentious public policy debates is impossible in principle. Any evidence or argument adduced in support of one side would simply lead to an elaboration of the belief system of the other side, who hold their views not on the basis of evidence and reason, but as part of a complex socially conditioned world view. An equivalent phrasing of this argument would be that there are no purely factual disputes. These writers would argue that the values of the disputants determine what scientific questions they consider important and what kinds of evidence are believed (Kuhn, 1979), and what sources of risk are worth worry and what responses to them are considered (Douglas and Wildavsky, 1982). Wynne (1982) describes the notions of scientific objectivity and a natural consensus as myths, and cites one of the very few cases in which an attempt was made to resolve fact-based policy disputes by bringing scientific representatives of both sides together for private consultation, away from the glare of public and media attention. The experience is described by Nelkin (1971) in her case study of the controversy over the Cayuga Lake nuclear power plant. Both sides remained firmly unmoved by the evidence and argument produced by their adversaries.

When disputes are acknowledged to be over values, there is little hope in seeking agreement; the issue is simply, who gets to decide and how? In this respect, decisions on risk resemble other policy decisions, except that the fervency with which views are held on risk issues is raising a new level of public interest in old questions of the relative advantages of representative and direct democracy. These broader issues of the appropriate role for direct public participation in political decisions are discussed in Kroeker (1979).

For resolution of risk disputes within the context of representative democracy, much attention has been paid to public inquiries. This may be in part the legacy of the Berger Inquiry, because of the large impact that it had on subsequent policy formation. An extensive survey of public inquiries in Canada has been performed by Salter and Slaco (1981).

Various participants attempt to push inquiries to meeting different purposes, and the net effect is a confused perception of their purpose. Officials often regard them as vehicles for passing expert information from the government to the people, and consequently for enhancing popular support for proposed policies. Intervenors regard them as vehicles for passing populist information from the people to the government, and for revising or influencing the subsequent decision. The inquiry can become a forum for intense and conspicuous competition between these two groups. In fact, the reality of differing values and interests between different groups means that a public inquiry cannot hope to bring harmony to an acrimonious issue; if the anthropologists are correct, I may not even bring about widespread factual agreement.

But there are some ways in which an inquiry can reduce the degree of dissent. The dissemination of clear and responsible information from experts may in some measure reduce the extremity of response. This may be accomplished as much through reversing the traditionally secretive and authoritarian attitude that has been attributed to experts and bureaucrats by the public as through inducing changes of factual opinion. Starr (1985) in particular argues that when there is apparently stubborn and irresponsible resistance to policy initiatives, this may be more a reaction to secretiveness and high-handedness than substantive dispute. He argues that if the process is made more open so that people can see that their concerns and interests are being considered, there will be more social cohesion and less dissent, even when total agreement over the decision outcome is impossible. This possibility of gradually building consensus and acceptance through a process of discussion may be the major political attribute of all these proposals for public participation in decision processes.

Another possibility for improvement of the process of risk decisions is offered by the recent development of powerful and versatile computer models to do analytic calculations related to risk, decision analysis, and economic modelling. Some experiments have been conducted with small groups of conflicting interests, in which the participants come together to build and discuss a model of the situation under dispute. Raiffa (1982), Sebenius (1981), and Holling (1978) have all reported experiences in which the task of building the model helped to create an atmosphere of joint problem-solving among the participants, and led to significant convergence of viewpoints. Raiffa points out that a particular advantage of such a model is that it provides an opportunity for abandoning a prior bargaining position without losing face, by making reference to newer, superior information provided by the model.

The possibility of wider participation in such populist analytic sessions is enhanced by the newest generation of risk analysis modelling systems, which combine relative ease of use with rapid results, and can perform sophisticated risk assessments with explicit representation of probability distributions. A leading example is the DEMOS system (Henrion et al, 1985) which can blend subjective probability distributions supplied by several people.

Sharing Risks

The essential inequity in how the burden of risk is borne raises the question of how society can distribute the burden more equitably. As there are two kinds of inequity, prior and posterior, there are two possible approaches to risk sharing: compensation in advance for the imposition of elevated risks, and insurance payments to those who suffer harm.

Ex ante compensation is closely related to the questions of how decisions are made to impose elevated risks on one group, and who is selected to take them. If the clear best interest of all requires the imposition of a risk that falls especially heavily on a few, then the rights of society and those of the threatened individuals are in conflict. The resolution of the conflict will depend n where the power lies. If the power to impose risks lies with society and a socially desirable decision entailing risk is to be imposed, then the ensuing competition among “not in my backyard” interests may well ultimately ensure that the risks are imposed on those most ignorant of the risks or on the politically powerless. The question of compensation is in this case not one of immediate necessity, but is still important in terms of ethics and long-term social cohesion.

If the power to refuse risks lies with individuals, then some risks may not be imposed on anyone, and society as a whole may well be worse off. But if the total social benefits really do exceed the costs and risk, it ought to be possible to strike a deal. Compensation is then necessary if the risk is to be imposed at all.

Raiffa (1982, Ch. 21) discusses schemes to determine appropriate compensation to pay to induce a community to accept a hazardous industrial facility. He proposes one scheme in which each candidate community names two prices- the maximum they would pay to another community to induce them to accept the plant rather than not have it built at all, and the minimum that they would have to be paid to accept the plant themselves. The plant is then built in the community with the lowest price to accept it, and the payment is a negotiated quantity between the recipient’s minimum payment and the sum of the willingness to pay figures of all the other communities. The decision of who bears the risk is thus made concurrently with the decision of the amount of compensation, and the risk is borne voluntarily.

Such ex ante payments compensate communities for the loss of property value that results from a risky or unpleasant facility. They thus assume that price changes in real estate markets accurately reflect people's attitudes to the imposed risk. In addition to the questionable validity of this assumption, there are two practical difficulties associated with ex ante compensation. First, the appropriate compensation levels are hard to determine; markets may not do it, and with non-market schemes the affected individuals have incentive to over-state their claims. Second, blanket compensation schemes cannot adequately reflect individual differences in required compensation. Some members or officials of a community may in fact prefer to see a risky facility built for economic or employment reasons, while others perceive the facility as all risks and no benefits. But attempting to find the fair level of compensation certainly seems preferable to letting an imposed risk go uncompensated.

If only some community members are empowered to oppose a decision effectively, there is the danger that they can be bought off with bribes while others who object but are powerless go uncompensated. Kunreuther et al (1982, p. 339) report such a case in German town of Bergkamen, in which a utility company paid members of a citizen's action group for them to cease their opposition. The payment was hotly contested in the media and the courts. If the utility had paid instead to develop parkland or recreational facilities in the town, that would likely have been more widely acceptable.

Insurance for losses suffered in case of accidents may also be a way of redressing inequitable distributions of risks. This can take the form of compulsory liability insurance held by the person or company carrying out the risky activity, or insurance by the government. Historically, it has often happened that insurance was not agreed on in advance, but the government has stepped in with compensation after a disaster. This has been the case both with technological hazards and natural ones.

The main problem with insurance for imposed risks is determining the appropriate level of coverage, particularly since, if losses occur, very high liability payments may be determined by jury verdicts.

It has also been suggested that the expectation that disaster relief will be provided, combined with people's mis-estimation of risks, will lead too many people to choose to put themselves in risky situations , for example by living behind dykes in flood plains. Perhaps one indication that the right size of compensation or insurance package has been provided for the imposition of a man-made hazard would be that there is no net migration into or out of the affected region. Whatever level is chosen, though, equitable distribution of risks might be assisted if the nature and size of compensation or insurance schemes were publicly established at the time that the risk is first imposed.

Regulation can also encourage equity of risk distribution by exploiting different sensitivities to risk. Uranium miners suffer a high risk of lung cancer, but the cancers are very slow in developing (with a latency period of 20 to 30 years), and the risks are vastly greater for smokers. In light of these interactions, perhaps these risks would be borne most fairly if only nonsmokers were allowed to work in uranium mines, or people over

50 years of age, while women of child-bearing age were definitely excluded. (The "charter cases" generated by such practices might turn out to be numerous, though.) Alternatively, better sharing of some risks might be achieved by information programs, disseminating information on the magnitude of risks and differences in sensitivity of different groups.

CHAPTER 5: LIVING WITH RISK

Most debate on risk takes place either at the level of controversy over a particular risk to be accepted, or dispute as to the methods society should employ in making decisions on risk.

But there is another still more general possible level of debate on risk-one in which the broad issues of social attitude to risk are considered, in which the concern is with questions such as: How we can best deal with the inevitably risky nature of the world? To what extent can or should the costs of risk bearing be shared by society? Is it necessarily preferable to achieve the lowest possible level of risk, or can there be advantages in risk bearing? This level of discussion has been insufficiently recognized in public policy debate, and a truly rational approach to risk demands that it be addressed.

It is not possible to be free of risk, or even to know the precise nature and magnitude of the risk you face. This is surely a commonplace observation, but it contradicts much of our behaviour toward risk and misfortune. We seem to be most unwilling to acknowledge the essential uncertainty of events; we reason that anything that happens must have had a prior cause, and its cause must have been either initiated by someone or detectable by someone. There are two beliefs inherent in this reasoning-that misfortunes can be predicted, and that someone is always responsible.

These beliefs can be seen operating in the behaviour of the commissions of inquiry that are frequently appointed in the wake of a tragic accident. The commission seeks the chain of prior events, either necessary or sufficient conditions, that led to the accident. If a set of prior causal conditions can be found, then it may be possible to find a way to prevent that precise path from happening again;. Finding the conditions also allows the identification of the person or persons responsible, in order to assign blame or punishment.

But it may be that the observed chain of events leading to an accident was truly rare and unforeseeable. This possibility calls into question both the efficacy of defending ourselves against what has already happened, and the justice of seeking someone to blame. What do we say, when some surprising misfortune happens-a dam bursts, or a train derails, or a child is gravely injured by a reaction to a vaccine-that its occurrence reveals at least a neglect of responsibility, if not reckless or malicious disregard for people's safety? Or that we have observed a dreadful but uncommon and unforeseeable event? If the former is true, then the responsible person must surely be sought; if the latter, then there may really be no responsible person, even though there may be many who could have deflected the causal sequence of events, had they perceived its import.

If, as seems likely, the degree to which misfortunes are humanly caused or foreseeable varies from 'entirely' to 'not at all', then the degree to which it is appropriate to affix blame must vary accordingly. Following a terrible accident, the tendency surely is too much in the direction of attempting to affix blame. We may demand in retrospect a standard of vigilance and risk avoidance that would not have been possible or practical in advance. Not knowing in advance where or how an accident is to happen, one can not know in what direction to be vigilant, and it is demonstrably not possible to pay full attention to every possibility at all times.

Given that risk is unavoidable and that it is not possible to make a system or policy perfectly free of error, what then can be done? A more achievable goal would be to make the system tolerant of error. If you cannot ensure that nothing will go wrong, then your next-best choice is to limit the consequences of things going wrong.

This approach might be described as an error-tolerant approach, in contrast to an approach that seeks to prevent errors altogether. The error-tolerant approach is often seen in systems subject to particular kinds of hazards that can be foreseen but not prevented. It is, for example, frequently engineered into control systems of complicated pieces of equipment such as aircraft, automobiles, or computers, to cope with the most likely or most costly kinds of operator error. If your word processing programme makes you ask twice before erasing a large document, it id demonstrating this approach to risk.

But how can such an approach be taken to risks whose origin and size are both unknown? If it is impossible to anticipate every kind of error or hazard that may arise, then it is impossible to take specific measures to reduce the destructiveness of each one. In this case, the error-tolerant approach can only be followed by limiting possible losses through some form of diversification.

An investment fund manager faces a problem of this kind in deciding what investments to buy. Whatever particular security is chosen has a risk of losing value, and there is no way of telling in advance which will gain and which will lose. Managers cope with this risk by dividing their funds among many different investments, expecting that the factors causing losses in some will not affect all, and so losses in some will be more than offset by gains in others. In the absence of risk, the manager would simply put the entire fund into the optimal investment, the one with the highest overall return. But this strategy, which is optimal under certainty, is extremely foolish under risk.

Clark (1978) points out that biological systems such as species and ecosystems must also cope with risk of unknown origin and size, and also manage it with strategies of diversification. In species, this can be most clearly observed in those dwelling in harsh environments. Such species often divide themselves into geographically separated sub-populations. Harsh conditions then may drive one or more of the sub-populations to extinction, but are less likely to extinguish the entire population. When one habitat is emptied by extinction, it is normally re-settled by colonization from the surviving populations, and the diversity is restored.

A similar diversification exists at the level of ecosystems. Ecologists have long been puzzled by the fact that there seem to be many more species of plants and animals than are necessary to do every "job" in an ecosystem. A simple-minded application of the theory of natural selection, without considering uncertainty, would suggest that there exists one species optimally adapted to fulfill each function, and that it eventually should come to dominate and drive its competitors to extinction. Not so. For each function, there are normally several different species, each of which is slightly better adapted under certain conditions. As the environment changes, the advantage periodically passes from one species to another, but conditions are rarely stable for so long that species are driven to extinction. The consequence of this diversity is that if a species is wiped out by some hazard, the entire ecosystem does not collapse; there are other species around to do the essential work. As with investment portfolios, the strategy that would be preferred under certainty is avoided in order to provide resilience against unforeseeable hazards.

The lessons of these analogies to public policy are clear, and provide a cautionary lesson to the disciples of optimization. If some activity is subject to unforeseeable shocks, then it is bet to have several different ways of doing it. We should not keep all our eggs in one basket. And to reply that any fool knows that is to fail to recognize the extent to which demands for efficiency and competitive pressure militate for us to put all our eggs into one basket. After all, one basket is demonstrably cheaper and more convenient than two-as long as there is no risk of dropping it.

There are many particular examples. An economy that is dependent on one source of energy will be vulnerable to price or supply shocks in that source. An economy substantially dependent on a particular skill, process, or labour force, could be held to ransom by them. An agricultural or forestry sector based on a small number of crops cultivated in huge monocultures is vulnerable both to biological hazards and to fluctuations in the markets for the crops. If substances of potential health or environment hazard are used in large quantities throughout the country or the world, then any harm that is realized will be on an enormous scale. And an economy dependent on foreign trade with one other nation is subject to disruption by changes in relations with the trading partner or its economy.

In all of these situations, the market pressures that militate for the most efficient, most profitable way of doing something in the short term-and the static analysis which supports those pressures-may contribute to an unstable system that is vulnerable in the long term. A mix of activities-an industry based on many different fuels, an agriculture producing many different crops, a less-than-fully specialized economy trading with many different partners-may be less than optimal under certainty, but vastly more capable of handling risk and offering some robustness against the shocks and sharks of a dynamic trading environment.

By its very nature, risk will be borne inequitably. This inevitable inequity in the incidence of harm argues for society attempting to share the burden, but the possibility and the appropriate level of sharing is clearly different for different types of risk.

When a risk is taken on voluntarily and knowingly by individuals, the social obligation to share the burden may be limited. For economic risks, one of the bases of the market economy is the ability of individuals to take risks and bear the resultant gains or losses. It is through the discipline of gains and losses to individuals that the economy adapts to changing technology and tastes, and so remains globally competitive. Moreover, total protection against economic losses would give individuals incentive to reckless behaviour, in the expectation that they can receive the benefits of their economic risks if they win but not bear the costs if they lose.

But few would argue that society should abandon even those who make bad economic decisions to starvation; there must be some measure of safety net or protection from deprivation. It is the issue of the appropriate degree of protection, obviously, that is difficult and contentious.

Does the same difficulty apply to physical risk taken voluntarily by individuals? On the one hand, no reasonable person would suggest that hikers who become stranded on mountains that are beyond their ability should be abandoned to their fate; on the other hand, if it is known that society will spare no expense, and even put other people at risk, in order to save individuals who have placed themselves in peril by reckless risk-taking, then one can expect insufficient prudence to be exercised in deciding what mountain to climb or whether to go sailing in bad weather. There is particular tragedy in the stories of a helicopter rescue team dying in the effort to rescue a climber with sprained ankle, or the deaths of eleven people searching for the wreck of a plane with two people aboard.

When risks are imposed on individuals involuntarily, then the obligation of society to protect those who suffer the consequences is presumably greater. The obligation is partly pragmatic, for no group is likely to tolerate being placed at risk unless they believe that their burden will be in some measure shared or insured. But it is also in part ethical, for it is unjust that some should be forced to bear uncompensated losses to advance the welfare of others. This is perhaps particularly true when the probabilistic prior risk burden is inequitable as well as the posterior, after-the-fact, distribution of harm.

But it may not be true that a risk assumed collectively generates a social surplus through which the losers can be compensated. Even if the original decision was correct in terms of expected surplus, there may not be any realized real surplus. If the “thousand-year accident” at a new LNG terminal happens in the first year of operation, then the social surplus will be a long time coming. Alternatively, the social decision may simply turn out to have been incorrect, and so there may not even be an expected surplus. And some risks-earthquakes, floods, war, global environmental hazards and international economic crises-obviously are involuntary even at the collective level, to the extent that these are unmitigable by the policy decisions of a single nation. In any of these cases, the losses of the victims may exceed the total ability of society to insure them.

There are also qualitative limits to society’s ability to compensate losses. Economic or property losses are compensable by payments of money or goods, but losses of life and health are not. No amount of money can truly compensate for the loss of a leg, a disabling handicap or the death of a loved one.

Recognition of the impossibility of true compensation may be behind the present tendency toward extravagant (initial) monetary awards in court cases for personal injury and death. Such awards can be interpreted as attempts to approximate the infinite payments that would be required for true compensation, or alternatively to exact vengeance or to send a message to designers and decision makers. But such awards cannot be justified on any of these grounds. Exacting vengeance for every harm suffered, given the inevitability of some harm, would involve all of society in a blood feud. And although a punitive settlement would be an appropriate means of sending a message that excessive risks were taken, it is not clear that a court has the competence or authority to determine what sizes of risks are acceptable. Even if decisions are made that impose society's ideal level of risk, there will still be victims. How can we face these people, and what can we say to them?

One thing that can be done is to recognize that while some losses cannot be compensated, some can. Medical care, therapy and replacement of lost income are matters for which compensation to victims can be made; pain and suffering, disability, and loss of loved ones are not.

In attempting to compensate victims, and in attempting to determine responsibility, there is a "moral hazard" problem that must be recognized. Just as post hoc inquiries assigning blame for non-predictable failure influence people's behaviour to make them excessively shun risks, so can the expectation that losses will be fully insured or compensated influence people's behaviour towards excessive risks. This is clearly the case when individuals are led to expect that they can retain the benefits of success from economic risk taking without having to bear the costs of failure. It has occasionally been suggested that large damage settlements against property-owners or municipalities for injuries sustained on their property may encourage too much risk-taking by individuals, but this seems unlikely; the prospect of monetary settlement for disability can surely not make the prospect of disability seem any less horrible if considered in advance. On the other hand, these settlements do promote too much risk avoidance on the part of the people held liable; it surely can not be socially desirable that cities close their playgrounds out of fear of responsibility in massive damage suits.

Finally, it must be recognized that when we take risks the uncertain consequences of our actions might be benefits rather than losses. By shunning uncertainty as much as possible, one counts oneself out of the gains as well as the losses. It is clear that individuals know this, for they take risks economically all the time for the prospect of resultant gains; this is the essence of innovation. They also take risks in order to learn; this is the essence of experiment. People can also be observed to accept risks to their life and health in return for financial or other gains; for many people, the benefits of sky-diving and motorcycle riding are clearly sufficient to offset the substantial associated risks.

If economic gains, learning and pleasure can all justify physical risk for individuals, can not the same argument be made for collective decisions? To pursue the minimum possible risk is to expend resources in the pursuit of a phantom, to reduce adaptability, and to miss opportunities for learning. Successfully minimizing risk in a certain policy area would require devoting all possible attention to analyzing the current state of the world and fine-tuning the current optimal solution. But the world is always changing, and such a strategy leaves nothing available with which to anticipate or respond to these changes.

This may sound like a fine philosophy, but it cannot be employed if any error is catastrophic or is punished severely. If any risk carries a significant probability of disaster, then one would be a fool to do other than attempt to minimize risks. But still new information and learning are only gained through taking risks.

The challenge then for public policy on risk is to find a socially-appropriate level of risk taking, somewhere between rigid caution and outright recklessness, and in particular to ensure that as much as possible, risks are taken in an environment of error tolerance. To a great extent, our economic and political institutions are established upon an assumption that some errors can be tolerated, in the investment of decision-making power in non-experts through juries and elections, for example.

Our collective decisions, intentional and unintentional, are at present a confusing mixture of too much risk and too little. While certain developments on a global scale, such as the nuclear arms race and the possibility of CO2-induced climate changes have us counting unimaginable risks, the reluctance to acknowledge the inevitability of risk, , and the excessive post hoc assignment of blame, lead many individuals and public servants to excessively shun risks. In some areas we even forbid individuals to assume risks that they might wish to take on in full knowledge, as in the lengthy testing and licensing process for new prescription drugs and the restriction of health care practice by non-physicians. In other areas we give incentives for reckless risk-taking by insuring individuals against economic losses, as when the federal government insures large investors in banks offering high rates of return in order to support high-risk lending policies.

The pursuit of error tolerance argues for diversification =in all realms of activity- in industrial strategy, agriculture, foreign trade, and environmental intervention. When, for example, a new activity or substance of small but unknown risk is considered, it might be treated in a number of ways. If the risk can be effectively confined to those who choose it voluntarily, as with the introduction of new prescription drugs, then a less paternal approach that focuses on informing potential consumers about the current state of knowledge regarding its potential risks and benefits is likely preferable to restricting its use. There would no doubt then be a large individual variation in chosen degrees of risk-taking, with benefits both to individual welfare and social knowledge and adaptability. Perhaps this is one answer to the ethical problem of experiment on humans; only experiment on those who choose to take the risk when fully informed of the current state of knowledge (although that in itself may be quite incomplete and subject to debate). The controversy over the prohibition of “premature” release of potentially efficacious drugs to AIDS victims thought to be terminally ill is a relevant example here.

The appropriate degree to which society can and ought to take the burden of risks off individuals, and the determination of the degree of social responsibility for individual suffering are matters that cut to the heart of the ambiguous nature of social organization. The conflict between protecting individuals and benefiting the collective arises in all aspects of public policy, but is particularly poignant on matters of risk. Neither extreme position is tenable. Because individuals matter, a position of social darwinism that cavalierly dismisses individual suffering in order to maximize the welfare of the surviving members of society is clearly unacceptable. But neither is it possible for society to completely insure individuals against all losses suffered, both because true compensation for physical suffering is impossible, and because the attempt to compensate all victims may cripple the adaptive capacity of society as a whole.

BIBLIOGRAPHY

Abrams, N.E. and Berry, R.S. (1977) "Mediation: A Better Alternative to Science Courts." Bulletin of the Atomic Scientists 133:4 April 1977, pp.50-53.

Ackerman, et al. (l974) referenced IIASA LNG pp. 313 **FIND**

Ackoff, R. (1974) Redesigning the Future: A Systems Approach to Societal Problems. New York: Wiley.

Allman, W.F. (1985) "Staying Alive in the 20th Century." Science 85, *p.31*.

Arrow, K. (1971) Essays in the Theory of Risk Bearing. Chicago: Markham.

Arrow, K.J. (l982) "Risk Perception in Psychology and Economics." Economic Inquiry 20:1, January, pp.1-9.

Bazelon, D.L. (1979) "Risk and Responsibility." Science 205, 20 July 1979, pp.277-280.

Bazelon, D.L. (1980) "Science, Technology, and the Court." Science 208, 16 May 1980.

Boffey, P.M. (1976) "Experiment Planned to Test Feasibility of Science Court." Science 193, 9 July, 1976, pp.129.

Borkan, B., and Kunreuther, H. (1979) "Contingency Planning Against Low Probability Events." Best's Review 80:106-I 12.

Burton, I.; Fowle, C.D.; and McCullough, R.S. (1982) Living with Risk:

Environmental Risk Management in Canada. Environmental Monograph Number 3, Institute of Environmental Studies, University of Toronto.

Burton, I.; Kates, R.W.; and White, G.F. (1976) The Environment as Hazard. Oxford University Press.

Callen, E. (1976) "The Science Court." Letter to Science 193, 10 Sept 1976.

Camerer, C. (1986) "Judgment Bias in Experimental Asset Markets: The Role of Experience." Unpublished manuscript presented at TIMS/ORSA meeting, Los Angeles, April 1986.

Clark, W.C. (1978) Managing the Unknown: An Ecological View of Risk Assessment. Working Paper W-26, Institute of Resource Ecology, University of British Columbia.

Clark, W.C. (1978) "Witches, Floods, and Wonder Drugs: Historical Perspectives on Risk Management." In Societal Risk Assessment: How Safe is Safe Enough?, eds. Schwing and Albers. New York: Plenum Press.

Coase, R.H. (1980) "The Problem of Social Cost." Journal of Law and Economics *3:1-44*.

Comar, C.L., and Sagain, L.A. (1976) "Health Effects of Energy Production and Conversion." Annual Review of Energy 1:558-99.

Combs, B., and Slovic, P. (1979) "Newspaper Coverage of Causes of Death." Journalism Quarterly 56:837-43.

Conrad, J. (1980) Society, Technology and Risk Assessment. New York: Academic

Press.

Covello, V.T., and Mumpower, J. (1985) "Risk Analysis and Risk Management: An Historical Perspective." Risk Analysis 5:2, p.103.

Crump, K.S. (1984) "Mechanisms Leading to Dose-Response. In Principles of Health Risk Assessment, ed. P.F. Ricci. Englewood Cliffs, NJ: Prentice-Hall.

Crouch, E.A.C., and Wilson, R. (1982) Risk-Benefit Analysis. Cambridge, MA:Ballinger.

Dasgupta, A.K., and Pearch, D.W. (1972) Cost-Benefit Analysis: Theory and Practice. London:Macmillan.

Doern, G.B. (1982) "The Politics of Risk: The Identification of Toxic and Other Hazardous Substances in Canada." Special Report, The Royal Commission on Asbestos January, Study No.4.

Doty, P. (1972) "Can Investigations Improve Scientific Evidence: The Case of the ABM." Minerva 10, 2 (Apr 72), pp.280-294.

Douglas, M., and Wildavsky, A. (1982) Risk and Culture. University of California Press, 1982.

Farber, E. (1981) "Chemical Carcinogens." New England Journal of Medicine 305: 1379.

Ferguson, R.A.D. (1981) Comparative Risks of Electricity Generating Fuel Systems in the United Kingdom. Stevenage, UK: Peter Peregrinus.

Fischoff, B., et al. (1981) Acceptable Risk. Cambridge University Press.

Gold, L.S., et al. (1984) "A Carcinogenic Potency Data Base of the Standardized Results of Animal Bioassays." Environmental Health Perspectives 58:9.

Graham, J.D., and Vaupel, J.W. (1981) "Value of a Life: What Differences Does it Make?" Risk Analysis 1:89.

Green, H.P. (1975) "The Risk-Benefit Calculus in Safety Determinations." George Washington Law Review, Washington DC, April 1976.

Hapgood, F. (1979) "Risk-Benefit Analysis: Putting a Price on Life." The Atlantic 243: 1, pp.35-38.

Hausch, D.B.; Ziemba, W.T.; and Rubinstein, M. (1981) "Efficiency of the Market for Racetrack Betting." Management Science 27:12, p.1435.

Henrion, M., et al. (1985) "Quantitative Risk Assessment of Automobile Restraint Systems." Risk Analysis 5:2.

Hertz, D.B., and Thomas, H. (1983) Risk Analysis and its Applications. New York: John Wiley and Sons.

Hohenemser, K.H. (1975) "The Failsafe Risk." Environment 17:6.

Hohenemser, C., et al. (1977) "The Distrust of Nuclear Power," Science, I 96:25-34.

Holdren, J.P.; Smith, K.R.; and Morris, G. (1979) letter, Science 208, pp.1242-1244.

Holling, C.S., ed. (1978) Adaptive Environmental Assessment and Management. Chichester:John Wiley and Sons.

Holling, C.S., and Clark, W.C. (1975) "Notes Towards a Science of Ecological Management." In Unifying Concepts in Ecology, eds. W.H. van Dobben and R.H. Lowe-McConnell. The Hague:W. Junk Publ., pp.247-251.

Howard, R.A.; Matheson, J.E.; and Miller, K.L. (1979) Readings in Decision Analysis. Decision Analysis Group, Stanford University.

Hynes, M., and Vanmarcke, E. (1976) "Reliability of Embankment Performance Predictions." Proceedings of the ASCE Engineering Mechanics Conference. University of Waterloo Press.

IAEA (1984) Risk and Benefits of Energy Systems. Vienna: Int. Atomic Energy Agency.

Inhaber, H. (1982) Energy Risk Assessment. New York: Gordon and Breach.

Int. Agency for Research on Cancer (1980) Long-Term and Short-Term Screening Assays for Carcinogens: A Critical Appraisal. Lyon, France: IARC.

Jackson, J., and Kunreuther, H. (1981) "Low Probability Events and Determining Acceptable Risk: The Case of Nuclear Regulation." Professional Paper, PP-81-87, May, Laxenburg, Austria: IIASA.

Jasanoff, S., and Nelkin, D. (1981) "Science, Technology, and the Limits of Judicial Competence." Science 214: 121-1215.

Kahneman, D., and Tversky, A., (1979) "Prospect Theory: An Analysis of Decisions Under Uncertainty." Econometrica 47:2, pp.263.

Kantrowitz, A., (1967) "Proposal for an Institute of Scientific Judgment." Science 156, pp.763-764.

Kantrowitz, A. (1977) "The Science Court: Criticisms and Responses." Bulletin of the Atomic Scientists 133:4, pp.43-49.

Kaplan, S., and Garrick, B.J. (1981) "On the Quantitative Definition of Risk." Risk Analysis 1:11.

Kleindorfer, P. (1982) "Group Decision Making Methods for Evaluating Social and Technological Risk." In Controversy: An Institutional Perspective, eds.

Kunreuther, H. and E. Ley. New York: Springer-Verlag (forthcoming).

Kleindorfer, P., and Kunreuther, H. (1982) "Misinformation and Equilibrium in Insurance Markets." In Economic Analysis of Regulation Market, ed. J. Finsinger. London: MacMillan (in press).

Kluge, E.H.W. (1986) "What is a Human Life Worth?" Canadian Public Administration (forthcoming).

Kroeker, H.V., ed. (1979) Sovereign People or Sovereign Governments. Montreal:Institute for Research on Public Policy.

Kuhn, T.S. (1962, 1979) The Structure of Scientific Revolutions. University of Chicago Press.

Kunreuther, H. p1981) Risk: A Seminar Series. Laxenburg:International Institute for Applied Systems Analysis.

Kunreuther H.; Linerrooth J., et al. (1982) Risk Analysis and Decision Processes. Laxenburg, Austria: IIASA.

Kunreuther, H., et al. (1978) Disaster Insurance Protection. New York: Wiley.

Kunreuther, H., and Lathrop, J. (1981) "Siting Hazardous Facilities: The Case of LNG Terminals." Risk Analysis 1:4, p.289.

Kunreuther H., and Linnerooth, J. (1984) "Low Probability Accidents." Risk Analysis 4:2, p.143.

LaBreque, M. (1981) "On Making Sounder Judgments." Psychology Today, June, pp.33.

Lamarsh, J. (1975) Introduction to Nuclear Engineering. Reading, MA:Addison-Wesley.

Lathrop, J., and Linnerooth, J. (1982) "The Role of Risk Assessment in a Political Decision Process." In Analysing and Aiding Decision Processes, ed. P. Humphreys and A. Van. Amsterdam,North Holland:IIASA.

Lave, L. (1982) Quantitative Risk Assessment in Regulation. Washington:Brookings Institution.

Lave, L. (1986) "The Problem with Risk Analysis and Management." Risk Newsletter, March, p.2.

Layard, R., ed. (1972) Cost-Benefit Analysis: Selected Readings. Penguin Books.

Lesourne, J. (1975) Cost-Benefit Analysis and Economic Theory. New York: North Holland/Elsevier.

Lichtenstein, S., et al. (1978) "Judged Frequency of Lethal Events." Journal of Experimental Psychology: Human Learning and Memory 4:551-78.

Lind, R., et al. (1982) Discounting for Time and Risk in Energy Policy. Baltimore, MD: Johns Hopkins University Press for Resources for the Future.

Linnerooth, J. (1982) "Murdering Statistical Lives?" IIASA Working Paper, May.

Lipfert, F.W. (1985) "Mortality and Air Pollution: Is There a Meaningful Connection?" Environmental Science and Technology 19, p.764.

Lovins, A.B. (1977) "Cost-Risk-Benefit Assessments in Energy Policy." George Washington Law Review 45:5 (Aug 1977) pp.911.

Lowrance, William W. (1976) Of Acceptable Risk. Los Altos, California: William Kaufman, Inc.

Manne, A.S., and Richels, R.G. (1980) "Probability Assessment and Decision Analysis of Alternative Nuclear Fuel Cycles." In National Energy Issues: How do we Decide?, ed. Sachs. Ballinger.

Marschak, J (1950) "Rational Behaviour, Uncertain Prospects, and Measurable Utility." Econometrica April.

Mazur, A. (1977) "Science Courts." Minerva 15:1, pp. 1-14, Spring.

McClory, R. (1986) "Murder on the Shop Floor." Across the Board, June.

McCormick, N.J. (1981) Reliability and Risk Analysis. New York: Academic Press.

Meisner, J. (1976) "Processing and Evaluating Expert Opinion." Paper presented at the Second European Congress on Operations Research, Stockholm.

Miller, A.B. (1982) "Epidemiology in the Assessment of Risk." Proceedings of the Symposium on Risk in Canada, Royal Society of Canada, Oct 18-19, 1982.

Mishan, E.J. (1971) "Evaluation of Life and Limb: A Theoretical Approach." Journal of Political Economy 79: 687-705.

Mishan, E.J. (1975) Cost-Benefit Analysis (2nd ed.). London: George Allen and Unwin.

Molton, L.S., and Ricci, P.F. (____) "Increased Risk: Science and Law", Cambridge, MA:Arthur D. Little Inc.

Morgan, M.G.; Ameral, D.A.L.; Herrion, M.; and Morris, S.C. (1982) Technological Uncertainty in Policy Analysis. Pittsburgh, PA: Dept. Engineering and Public Policy, Carnegie-Mellon University.

National Research Council (U.S.) Committee on Risk and Decision-Making (1982) Risk and Decision-Making: Perspectives and Research. Washington: National Academy Press.

NAS/NRC (1982) Risk and Decision Making: Perspectives and Research. Washington, DC: National Academy Press.

NAS/NRC (1983) Risk Assessment in the Federal Government: Managing the Process. Washington, DC: National Academy Press.

Nisbett, R., and Ross, L. (1980) Human Inference: Strategies and Shortcomings of Social Judgement. Englewood Cliffs, N.J.: Prentice-Hall.

Novegno (1985) Risks from Energy and Other Complex Industrial Systems. IAEAS/UNEP/WHO, International Project.

OECD (1983) Environmental Effects of Energy Systems. Paris.

Operations Research Society of America, "Guidelines for the Practice of Operations Research." Operations Research 19:5 (Sept 1971), pp.1123-1158.

Otway, H.J., and von Winterfeldt, D. (1981) "Risk Management and Acceptable Risk Criteria." In Kunreuther (1981).

Raiffa, H. (1968) Decision Analysis. Reading, MA: Addison-Wesley.

Raiffa, H. (1980) "Science and Policy: Their Separation and Integration in Risk Analysis." Address to 6th Symposium on Statistics and the Environment, Oct 6, 1980.

Raiffa (1982) The Art and Science of Negotiations. Cambridge, MA: Belknap/Harvard U. Press.

Ramsey, W. (1979) Unpaid Costs of Electrical Energy: Health and Environmental Impacts from Coal and Nuclear. Baltimore: Johns Hopkins University Press for

Resources for the Future.

Rappoport, E., (1981) Doctoral Dissertation, Department of Economics, UCLA.

Rasmussen, N.C. (1981) "The Application of Probabilistic Risk Assessment Techniques to Energy Technologies." Annual Review of Energy 1981, pp.123-138.

Rescher, N. (1983) Risk: A Philosophical Introduction to the Theory of Risk Evaluation and Management. Washington, DC: University Press of America.

Rethans, A. (1979) "An Investigation of Consumer Perception of Product Hazards." Doctoral Dissertation, University of Oregon.

Ricci, P.F., and Molton, L.S. (1981) "Risk and Benefit in Environmental Law." Science 214:1096.

Ricci, P.F., and Molton, L.S. (1984) "Risk Analysis in United States Law: Cost Versus Risk." In Low Probability/High Consequence Risk Analysis, Plenum Press.

Ricci, P.F., and Molton, L.S. (1985) "The Regulation of Cancer Risks." Environmental Science and Technology 19:473.

Ricci, P.F., and Molton, L.S. (1986) "Health Risk Assessment: Science, Economics, and Law." Annual Review of Energy 11, p.77.

The Royal Society (1983) Risk Assessment, A Study Group Report. London: The Royal Society.

Ruckelshaus, W.D. (1983) "Science, Risk, and Public Policy." Science 220:1026.

Salter, L., and Slaco, D. (1981) Public Inquiries in Canada. Background Study Number 47, Science Council of Canada.

Sass, R. (1986) "The Need to Broaden the Concept of Risk in Workplace Health and Safety." Canadian Public Policy 12:2, p.286.

Savage, L.J. (1972) The Foundations of Statistics (2nd ed.). New York: Dover.

Schelling, T. (1968) "The Life you Save May be Your Own." In Problems in Public Expenditure Analysis, ed. Chase. Washington: Brookings Institution.

Schneiderman, M.A. (1980) "The Uncertain Risks we Run." Societal Risk Assessment, How Safe is Safe Enough?, ed. R.C. Schwing, W.A. Albers. New York: Plenum.

Science Council of Canada (1977) Policies and Poisons. Report No.28.

Sebenius, J.K. (1981) "The Computer as Mediator: Law of the Sea and Beyond." Journal of Policy Analysis and Management 1:1, Fall 1981, pp.77-95.

Sheena, M.F. (1982) "Importance of the Burden of Proof in Environmental Litigation." The Environmental Professional 4:75-83.

Shrade-Frechette, K.S. (1985) Risk Analysis and Scientific Method. Dordrecht/Boston: D. Reidel Publishing Company.

Shrader-Frechette, K.S. (1985) Science Policy, Ethics, and Economic Methodology. Boston: D.Reidel.

Shrecker, T.F. (1984) Political Economy of Environmental Hazards. Law Reform Commission of Canada Study Paper.

Slovic, P., et al. (1977) "Preferences for Insuring Against Probable Small Losses: Insurance Implications." Journal of Risk and Insurance 44:237-58.

Slovic, P.; Fischhoff, B.; and Lichtenstein, S. (1979) "Rating the Risks." Environment 21:3, pp. 4-20, 36-39.

Slovic, P.; Fischhoff, B.; and Lichtenstein, S. (1980) "Risky Assumptions." Psychology Today June 1980, pp.44.

Slovic, P.; Fischhoff, B.; and Lichtenstein, S. (1982) "Facts Versus Fears: Understanding Perceived Risk." In Judgement Under Certainty: Heuristics and Biases, eds. D. Kahnemann, P. Slovic, C.A. Tversky. Cambridge:Cambridge University Press.

Smith, L.G. (1982) "Mechanisms for Public Participation at a Normative Planning Level in Canada." Canadian Public Policy 8:4, 561-572.

Somers, E., and Krewski, D. (1982) "Risks from Environmental Chemicals." Proceedings of the Symposium on Risk in Canada. Royal Society of Canada, Oct. 18-19, 1982.

Starr, C. (1969) "Social Benefits Versus Technological Risk." Science 165:1232.

Starr, C., (1985) "Risk Management, Assessment, and Acceptability." Risk Analysis 5:2, pp.97-102.

Starr, C.; Rudman, R.; and Whipple, C. (1976) "Philosophical Basis for Risk Analysis. Annual Review of Energy 1976, pp.629-662.

Starr, C., and Whipple, C. (1980) "The Risks of Risk Decisions." Science 208, pp.1114-1119.

Stokey, E., and Zeckhauser, R. (1978) A Primer for Policy Analysis. New York: Norton.

Svenson, 0. (1978) "Are we all Less Risky and More Skillful than our Fellow Drivers?" Acta Psychologica 47, pp.143-148.

Thaler, R., and Rosen, S. (1976) "The Value of Saving a Life: Evidence from the Labour Market." In Household Production and Consumption, ed. Terlecky. Columbia University Press.

Thompson, M (1981) "Beyond Self-Interest: A Cultural Analysis of a Risk Debate." IIASA working paper, Feb 1981.

Thompson, M. (1981) "To Hell with Turkeys." Postscript to Kunreuther, et al. 1981.

Tuohy, C.J., and Trebilcock, M. (1982) "Policy Options in the Regulation of Asbestos-Related Health Hazards." Study Series, Royal Commission on matters of health and safety arising from the use of asbestos in Ontario.

Tversky, A., and Kahneman, D. (1974) "Judgment under Uncertainty: Heuristics and Biases." Science 185, pp.1124-1131.

Tversky, A., and Kahneman, D. (1981) "The Framing of Decisions and the Psychology of Choice." Science 211, pp.453-458.

Tversky, A. (1986) "Rational Choice and the Framing of Decisions." Address to TIMS/ORSA conference, April 16, 1968, Los Angeles.

US Atomic Energy Commission (1974) Reactor Safety Study: An Assessment of Accident Risks in US Commercial Nuclear Power Plants, WASH-1400.

US Nuclear Regul. Commission (1981) Fault Tree Handbook, NUREG/CR-0492.

US Office Techn. Assessment (1985) Chemical Carcinogens: A Review of the Science and its Associated Principles, Washington, DC.

University of Waterloo, Institute for Risk Research (1984) "Risk Management in the Transportation, Storage, and Disposal of Hazardous Materials"." Draft report for the Canadian Committee for IIASA.

Von Neumann, J., and Morgenstern, 0. (1947) Theory of Games and Economic Behaviour. Princeton University Press.

Weinberg, A.M. (1972) "Science and Trans-Science." Minerva 10, pp.209-222.

Wynne, B. (1981) "Institutional Mythologies and Dual Societies in the Management of Risk." IIASA working paper, June.

Zeckhauser, R. (1975) "Procedures for Valuing Lives."

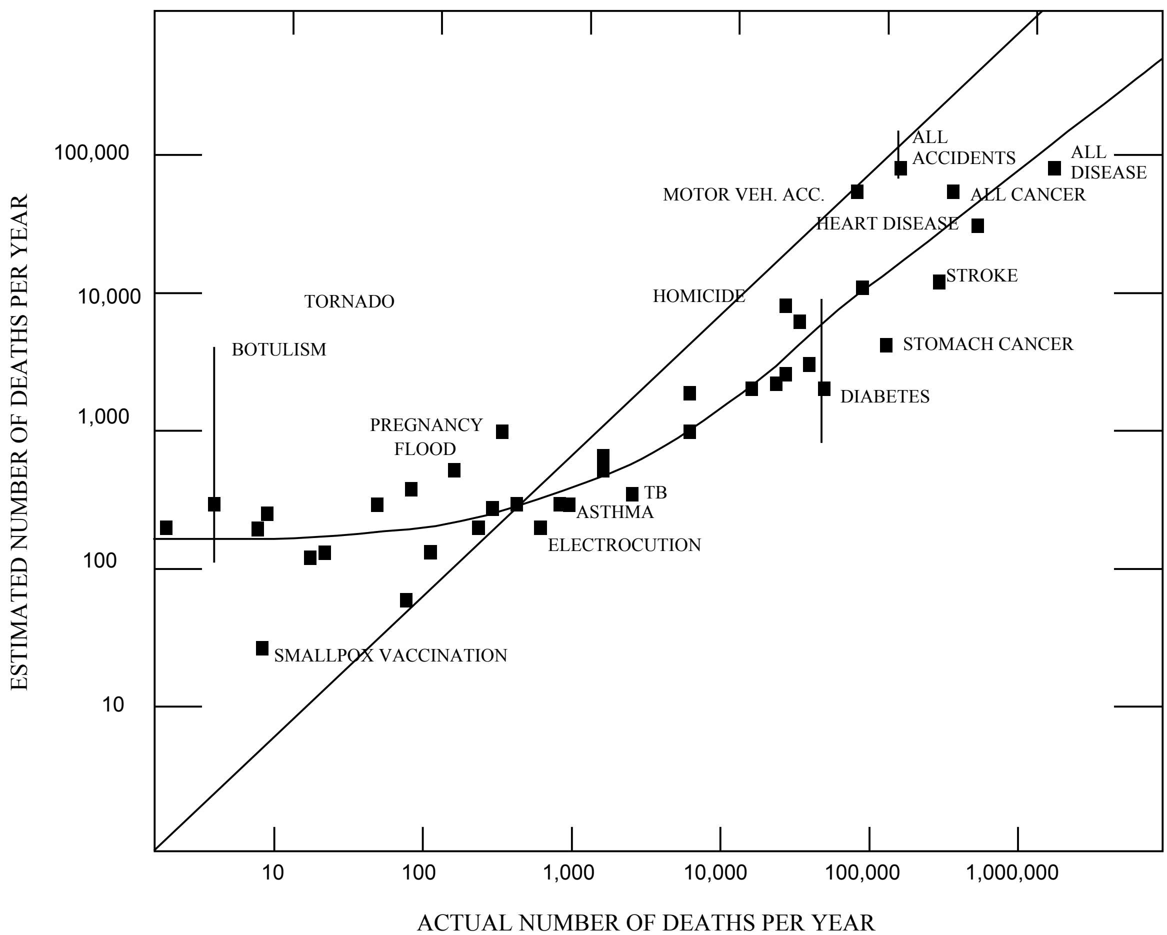
ESTIMATED NUMBER OF DEATHS PER YEAR
100,000
10,000
1,000
100
10
ACTUAL NUMBER OF DEATHS PER YEAR
10
100
1,000
10,000
100,000
1,000,000
ALL ACCIDENTS
ALL DISEASE
MOTOR VEH. ACC.
ALL CANCER
HEART DISEASE
STROKE
TORNADO
HOMICIDE
BOTULISM
STOMACH CANCER
DIABETES
PREGNANCY
FLOOD
TB
ASTHMA
ELECTROCUTION
SMALLPOX VACCINATION

Figure 2