Working Paper

The Social Mismanagement of Risk?

Risk Aversion and Economic Rationality

Joanne Linnerooth-Bayer

WP-93-71
November 1993

IIASA International Institute for Applied Systems Analysis □ A-2361 Laxenburg □ Austria
Telephone: +43 2236 715210 □ Telex: 079 137 iiasa a □ Telefax: +43 2236 71313

The Social Mismanagement of Risk?

Risk Aversion and Economic Rationality

Joanne Linnerooth-Bayer

WP-93-71
November 1993

International Institute for Applied Systems Analysis □ A-2361 Laxenburg □ Austria
Telephone: +43 2236 715210 □ Telex: 079 137 iiasa a □ Telefax: +43 2236 71313

Contents

The Social Mismanagement of Risk?

Risk Aversion and Economic Rationality

*Joanne Linnerooth-Bayer**

1 Introduction

Is society's system for managing risks to life and limb, as Zeckhauser and Viscusi (1990) put it, deeply flawed? Do our public officials overreact to risks of low probability but high salience, such as those posed by nuclear power, disastrous chemical releases and natural disasters[1], and put too little effort into ameliorating voluntary risks, such as those involving automobiles, diet, and smoking? As worrying as technological and natural disasters are, most of those who die prematurely do not die from any sort of collective disaster, at least in industrialized countries. U.S. life expectancy statistics reveal that individual risks from "private killers" such as heart disease, cigarette smoking, obesity, and even 'being single' rank at least two to three orders of magnitude over such catastrophic events as hurricanes and tornadoes, airline crashes, and accidents from the generation of nuclear energy. Are public expenditures to reduce or avoid low-probability/high consequence (LP/HC) events then flawed if this money could save more lives if spent elsewhere?

Economists usually frame this question as an issue of **risk aversion**, where one asks if the simultaneous death of 100 people is worse than the isolated deaths of 100 otherwise identical people. A close cousin of this risk-aversion concept is a concern not for collective deaths, but for avoiding a "worse-case" eventuality even if the deaths are private and individual.[2] A predilection to avoid the possibility, even if remote, of a large number of "private" deaths has been viewed as a primary motivation for conservative tendencies in risk assessment procedures (Nichols and Zeckhauser, 1986).

Many view such risk-averse preferences as, at best, irrational or uninformed, and, at worst, morally repugnant since the implication is that lives are valued differently whether

*International Institute for Applied Systems Analysis, Laxenburg, Austria. This paper was also partly supported by the Center for Risk Management, Resources for the Future, Washington D.C. I sincerely thank Dominic Golding, Howard Kunreuther, David Ball and Louis Goossens for their suggestions and help. Of course, I take full responsibility for all views expressed in this paper.

[1]The global losses from geophysical hazards (draughts, earthquakes, floods and tropical cyclones) have been estimated to total an annual average of about 250,000 deaths and $15 billion in damage as well as costs of prevention (Burton, et.al. 1978). Not all potential disasters are very low-probability events. Before the 1989 California earthquake, the Federal Emergency Management Agency estimated that a major earthquake in California could be expected with a probability of from 2% to 5% per year. Estimated fatalities were in the order of from 3,000 to 23,000, much higher than what was experienced. FEMA also expressed the view that the nation was essentially unprepared for an earthquake with such devastation (FEMA, 1980).

[2]Kasperson and Pijawka (1985) remind us, for instance, that the accumulated exposure of 8 to 11 million U.S. workers to asbestos is expected to result in as many as 67,000 workers dying prematurely each year. This number exceeds even the well-worn comparison of some 40 to 50 thousand annual fatalities from automobile accidents (National Safety Council, 1987).

lost collectively or individually. The alternative is a **risk neutral** position where the guiding principle for social policy is to save as many lives as possible given the available resources. Reflecting on catastrophe avoidance, Lichtenstein, et.al. (1988) comment, "It seems more compelling to us that it is the moral obligation of our social decision makers to save as many lives as possible; that implies risk neutrality (p. 17). Zeckhauser and Viscusi also view risk neutrality as the remedy for "society's flawed risk-management practices". They advocate a lives-saved standard of value as the most effective means of promoting society's risk-reduction objective (p.248).

While laudable, I show in this paper that this risk-reducing objective is not consistent with the welfare basis of the economist's benefit–cost analysis, unless it can be assumed that the primary concern of people in most risk contexts is maximizing their odds of survival. Zeckhauser and Viscusi assume this to be the case. Behavior that deviates from maximizing these odds illustrates, in their words, "the limits of human rationality" (p. 559).[3] The question they then pose is, "how should we proceed once we admit that individuals do not correctly react to many risks?" (p. 248). Not all economists take this view.[4] Among others, Freeman and Portney (1989) encourage economists to recognize that the public cares about more than simply the statistical magnitude of risks: "Until these concerns are acknowledged and incorporated in economic models, economists may dismiss as irrational responses that make very real sense" (p. 4). The more appropriate question appears then to be: How should we proceed once we admit that many individuals have a broader concept of risk than the probability of survival?

I proceed by setting out a very simplified choice situation that illustrates two different concepts of risk aversion, both of which allow for deviations from a lives-saved objective, but which have different ethical considerations. I then ask how the economist, who follows the precepts of benefit–cost analysis, would evaluate the choice between risk neutrality and risk aversion. Since the economist's reasoning depends on how the affected individuals, themselves, would decide, and since individual choice is often described as irrational, I examine the issue of rationality from different perspectives—decision analysis, cognitive psychology, sociology, and cultural anthropology. I conclude that inconsistent personal risk choices should not disqualify public concerns and citizen choice as a legitimate input to policy decisions involving LP/HC events. At the same time, I question whether the precepts of benefit–cost analysis provide an acceptable normative decision rule for making controversial risk decisions. I illustrate these theoretical issues with policy example involving the transport of radioactive wastes through New York City.

2 Two Versions of Risk Aversion

Ever since Bernoulli's explanation in the 18th century of the famous Petersburg paradox, which centered around the question why most people will pay only a small amount for a gamble of infinite mathematical expectation, the concept of risk aversion has been with us. In the context of expected utility theory (which has been the most influential paradigm of individual decision making), if a person prefers a certain monetary payoff over a gamble with a greater expected payoff, this person is said to be risk averse. The explanation lies simply in the notion of diminishing marginal utility of monetary gains.

[3]Early proponents of this irrationality argument have also included Häfele (1979), Starr and Whipple (1980), Maxey (1979), Comar (1979), and Cohen and Lee (1979). For a discussion, see Schrader-Frechette (1986).

[4]See, for instance, Schultze and Kneese (1980); Lave (1989).

Stripped of social and political context, the more controversial phenomenon of risk aversion—where lives are at stake—is analogous. A risk-averse individual would choose a policy resulting in ten anonymous deaths to a gamble with an expected loss of ten lives, but which offered a small probability of a large number of deaths. In this context, risk aversion is thus a preference for individual "statistical" deaths over more catastrophic events, even if expectations in terms of loss of life are the same. The flip side of risk aversion is a preference for avoiding collective disasters even if more lives could be saved with the same resources were they allocated to saving "statistical" persons.

Closer scrutiny of the risk-aversion concept reveals, however, a number of perplexing issues of definition. What distinguishes, for example, an airline disaster from the daily disaster on our highways? Is the loss of 100 people from a plane crash or from the release of radioactive materials in New York City any more a disaster than 100 people dying individually in their automobiles, or 100 people dying from their exposure to asbestos, or from an earthquake? Is a disaster a disaster because of the proximity of the deaths in time and space, because of the commonality of the cause, or because of other factors such as the collective nature of the deaths?

Multiple Fatalities and Collective Fatalities

I suggest that there are, at least, two dimensions to disasters and thus to risk aversion: an aversion to large numbers of deaths from a common cause and an aversion to collective deaths where the individuals have little influence or control over the outcome. While these two dimensions cannot be wholly separated from the host of psychological concerns found in the risk perception literature, such as control, dread, etc., the point is that risk aversion cannot be, or rather should not be, attributed only to what Green and Brown (1978) have called the "kill size". Concern only with the numbers of persons who are affected by a particular cause leads too readily to the moral objection of seemingly valuing lives more if they happen to be lost in large numbers.

These two versions of risk aversion can be appreciated by considering three risk situations, A, B, and C, shown in Figure 1. RISK A exposes each of 1000 persons to an annual chance of death of 1/1000. In contrast, RISK B takes the form of a group lottery with a one in ten chance of 1000 persons being exposed to a rather high annual chance of death (1/100), and otherwise no additional risk. RISK C takes the form of a group lottery in which there is a 1/1000 chance of a disastrous event taking the lives of all 1000 persons, and a 999/1000 chance of no fatalities at all. If only one of the risks can be eliminated with the available funds, which risk should the public policy maker choose?

Before proceeding, I acknowledge that these choices are highly abstracted versions of life-risking decisions pulled artificially out of any real policy context. The risks shown in Figure 1 show no ambivalence or uncertainty with respect to the probabilities. Full agreement is assumed in the point estimates. The outcome is one dimensional, and the deaths are assumed to occur to present generations. The age of the victims appears to be irrelevant, as is the way in which they die. Moreover, I am framing the risk issue as a matter of probability and magnitude, which may not be essential to public concerns about risky situations (Rayner, 1987). The justification for adopting this formulation is that it is precisely this abstraction that underlies much of the analytical discussion of risk aversion. My purpose is to show that the case for or against risk aversion is complex even when reduced to its simplest, context-free form.

Figure 1 shows three simplified risk situations, each of which is equivalent with regard to expected outcome, namely one fatality. Yet, the choices differ in form. In the case of

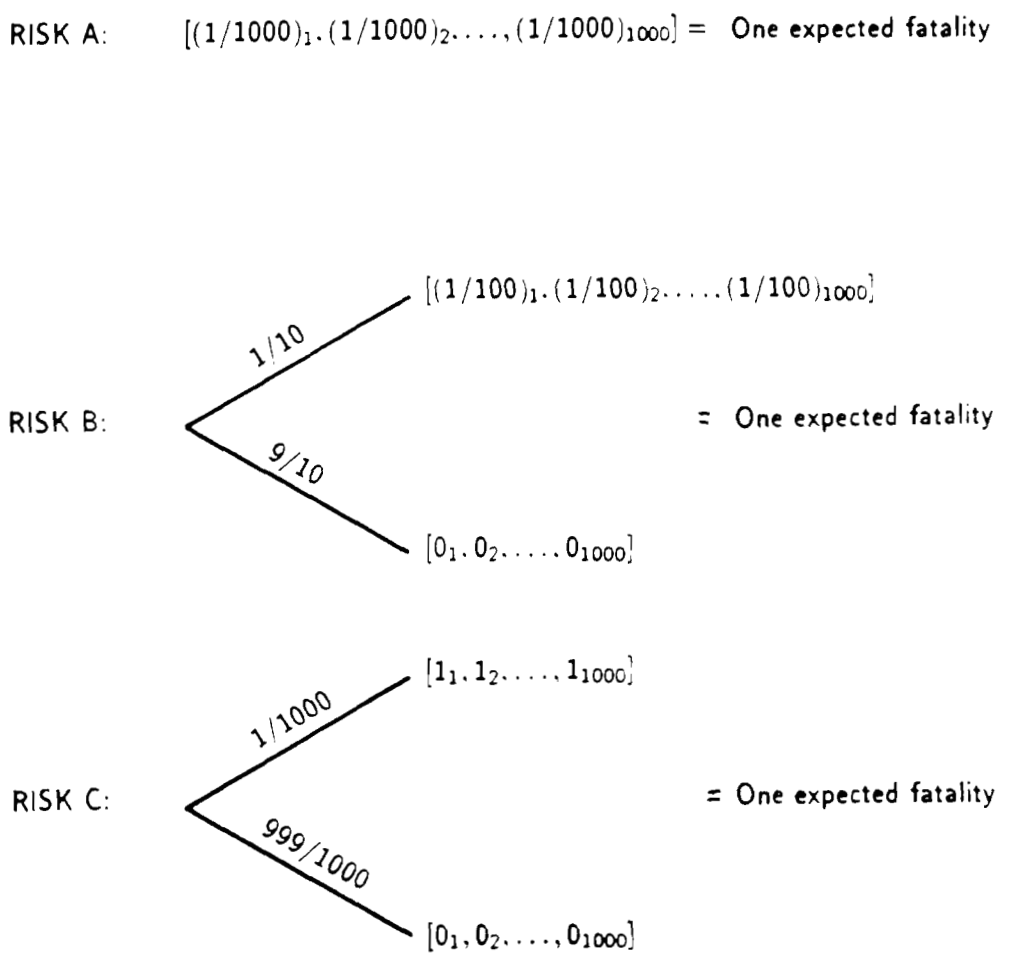

Figure 1: Three Risk Situations Illustrating Risk Aversion.

RISK A, which exposes each of 1000 persons to an annual chance of death of 1/1000, the probabilities are generated independently—each person flips his or her own coin. No one person's fate is conditional on the fate of the others. This could be a group of persons with high cholesterol where the annual chance of death from heart failure for any one individual is 1/1000, or it could be a group of high-risk drivers. The fatalities will likely be anonymous in the sense that little public attention will be focused on them. In contrast, RISK B takes the form of a group lottery with a one in ten chance of 1000 persons being exposed to a rather high annual chance of death (1/100), and otherwise no additional risk. RISK B is characteristic of a range of societal risks where **if** something were to happen (e.g., a release of radioactive materials in New York City, a chemical release such as what occurred at Bhopal or Seveso, or a nuclear power accident such as at Chernobyl), **then** there may be many deaths in the present or in the future.

Nichols and Zeckhauser (1986) offer an intriguing contrast between Risk A and Risk B. Two chemicals are assumed equally effective in treating a medical condition, but which are suspected of being carcinogenic. Chemical #1 poses an annual risk of cancer of 1/1000 for each person. Administering this chemical to a group of 1000 people would result in a risk situation analogous to RISK A in Figure 1. The risk from chemical #2 is uncertain; there is a 90 percent chance that it is perfectly safe (no risk) and a 10 percent chance that it poses a risk of one in a hundred (10 times higher than chemical A). Administering chemical #2 to a group of 1000 people results in a situation analogous to RISK B in Figure 1. Should these two types of risks be evaluated equivalently?

What if instead chemical #2 had a 99.9% chance of being perfectly safe and a 0.1% chance of killing the recipients outright? RISK C thus takes the form of a group lottery in which there is a 1/1000 chance of a disastrous event taking the lives of all 1000 persons, and a 999/1000 chance of no fatalities at all. Like RISK A, each person faces the grim prospect of a 1/1000 chance of death. Also like RISK A, the statistical expectation is one death. Unlike RISK A, however, the probabilities are not independently generated—one coin is flipped determining the fate of the entire group collectively. If we are told that one person is a fatality, we can safely assume that all 1000 are also fatalities. In such cases, the deaths usually (but not always as this example illustrates) occur simultaneously at one place and draw considerable public attention: the collapse of superhighway 19 in the 1989 San Francisco earthquake, for example.

Why might a person choose RISK A over RISK B and C? First, there is an unsettling collective or public dimension to RISKS B and C—a coin is flipped for the entire group determining to varying degrees the risk to each participant. Of course, "who throws the die" is statistically irrelevant when the decision is context free. While this distinction is of little or no importance in the artificial choice setting, it may have significance in a real-world context to the extent that a psychological premium is placed on probabilities that are generated individually: each person "spins his own wheel of fortune", "determines his or her own fate". This 'personal control' is often coupled with additional factors that have proven salient to public concerns, for instance, involuntary and/or uncompensated risks, factors which also tend to characterize RISKS B and C. Our second dimension of concern, the prospect of a disastrous, worst-case sort of outcome with multiple fatalities is also greater for RISK C than for RISK B, and still smaller for RISK A.

The message is that, even in this highly simplified example void of most real-world context, there are two distinct types of risk aversion:

(i) Multiple-fatality risk aversion, which is an aversion to deaths in large numbers. This is sometimes referred to as societal risks, and

(ii) collective-fatality risk aversion, which is a non-neutral valuation of the probability-generating process, where a preference is displayed for independent trials — for "individual", private risks over "collective", public risks.

As I will discuss in the next section, only the first concept is usually considered consistent with "rational" choice, but it is exactly this "rational" dimension of risk aversion that has been challenged as ethically unacceptable. The second concept, the aversion to collective types of fates, may be more acceptable from an ethical standpoint, but is perhaps too often labeled as irrational.

3 The Economist's Perspective

Benefit–cost analyst asks if the benefits of a government policy or project outweigh the sacrifices society must make to have it, a proposition so obvious that many economists regard benefit–cost analysis as nothing more or less than common sense (Freeman and Portney, 1989). Of course, any normative tool for social choice or for determining which goods or services in a society should be provided and how they should be allocated has either an explicit or implicit moral foundation.[5]

A positive benefit–cost ratio implies that the gainers from a policy choice are able to more than compensate the losers, even if this compensation is not actually carried out. This quasi-utilitarian rule, that the gainers gain more than the losers lose, is known as the Kaldor–Hicks criterion and is the main ethical premise of benefit–cost analysis.[6] It is not, however, the only premise since it is also important to decide what counts as gains and what counts as losses. A fundamental notion of economic welfare theory and thus of BCA is that each individual is the best judge of his or her own welfare, and, moreover, that each person pursues his or her own self interest (for a discussion of this assumption, see Sen, 1982). Individual preferences are revealed by choices in the market. Where there is no market for the goods or services being measured, and therefore no revealed preferences, their value can be measured according to what each individual is willing to pay for them or, alternatively, according to how much each individual must be compensated to go without them. These two different measures of value are discussed below.

Benefit-Cost Analysis and Mortality Risk

What do these ethical premises of BCA mean for the public policy maker who is contemplating which risk is most deserving of mitigation from the three risks (A, B, and C) shown in Figure 1? One popular way to proceed is to make another simplifying and very crucial assumption: The persons at risk are concerned only about their probability of survival. They are indifferent as to how they die as well as to the context in which these risks are set. In addition, only those people at risk are willing to pay to reduce the risks—family and friends are willing to contribute nothing. Looking again at Figure 1, risks A, B, and C pose an *ex ante* risk of death of 1/1000 for each of the 1000 people in

[5]Boulding (1969) distinguishes a "moral proposition" from an individual taste in that the former is intended to apply to society more broadly (p.1.).

[6]Schulze and Kneese (1980) point out that benefit–cost analysis is a special case of the utilitarian ethic where individuals' utility functions are linear with identical constant marginal utilities across individuals and where future utilities are identically discounted. Therefore, the Kaldor–Hicks rule does implicitly require interpersonal utility comparisons.

the exposed group. The relevant question is how much each of the 1000 people is willing to pay to reduce this risk to zero. (For a graphical presentation, see the Appendix). The full value of mitigating each risk in Figure 1 would be the aggregate 'willingness to pay' summed over the 1000 individuals.

Framed in this way, the economist's reasoning leads to a risk-neutral position. Since the risk reduction is the same in each case, WPT must also be the same. The risks in Figure 1 are then equally deserving of reduction. Before examining this apparent risk neutrality on the part of the economist more closely, it is important to discuss two aspects of this result which one may find troubling. First, as Freeman (1989) has pointed out, mortality risk is viewed by economists as an *ex ante* exercise in the sense that policy makers should focus on the prospects before the uncertainty is resolved. According to Freeman (p.310), an *ex ante* social welfare function implies that social welfare is a function of changes in expected utility, that is, it reflects a concern with opportunity in the expected value sense, while an *ex post* social welfare function reflects a concern with outcomes. One can appreciate this distinction by examining the risks in Figure 1. The *ex ante* measure is concerned with the value of reducing a 1/1000 risk over 1000 individuals; the *ex post* measure is concerned with the value of saving one life. Since value as measured by the necessary compensation would likely be infinite for the latter, Freeman and others advocate the *ex ante* measure as appropriate for benefit–cost calculations.

This brings up a second difficult aspect of this type of valuation. Clearly a person facing certain (near-term) death would be willing to pay a very large sum to avoid this eventuality, perhaps much more than he or she is able to pay. That the value of risk changes is necessarily bounded by each individual's budget or wealth may be troubling for those who are unhappy with the current distribution of wealth, and, indeed, the BCA practitioner accepts the status quo as a legitimate starting point for the analysis. Moreover, many feel that public expenditures on health and safety should be independent of the wealth of the beneficiaries. The economist rightly points out, however, that society's resources are finite. A benefit–cost assessment is based on the notion of allocating society's scarce resources such that the benefits of risk reduction at the margin are just equal to the social costs of this risk reduction.

While this logic is convincing in the aggregate, it may not be entirely appropriate for individual cases of risk reduction. Even those who advocate the WTP logic for valuing mortality risk may find the budget constraint inappropriate in some cases. In fact, BCA requires that a second criteria also be met before a project is recommended. This second, equally justifiable measure of value asks not what the individual is willing to pay for the risk reduction, but what he or she would have to be compensated if the risk reduction program were not undertaken. This is referred to as a person's willingness to accept (WTA). Most people would not accept even infinite compensation for certain immediate death, whereas they can only pay a finite sum to avoid this rather unpleasant prospect.[7] The higher the risk, the greater the deviation will likely be between a person's WTP and WTA.

If the WTP and WTA measures lead to different results in the benefit–cost calculations, Mishan (1988, p.193) advises that the economist forego his or her mandate to recommend a course of action. Both measures must lead to the same result for unambiguous fulfillment of the Kaldor–Hicks criterion. It follows—and this is generally not recognized—that the

[7]Adams (1989) takes issue with the notion that individuals should be indifferent between an amount received or aid for marginal deteriorations or improvements in environmental quality respectively. It should be noted that this is true **only** at the margin, i.e. the slope of the indifference curve. It is **not** true for larger changes.

economist can only justify use of the WTP measure for very small changes in mortality risks (see Linnerooth, 1982).

Is the economist risk neutral?

The risk-neutral outcome of the above analysis is not in any sense intrinsic to the economist's benefit–cost logic, but results solely from the assumption that the people at risk are, themselves, risk neutral. The two key assumptions that drive the risk-neutral result are:

(1) The exposed persons are concerned only with their survival chances, and

(2) only these exposed persons are willing to contribute to reducing the risks.

Economists have long recognized that these assumptions are too restrictive. The way in which people die is important to them, as well as their pain and suffering[8] and the concerns of friends and relatives. Interdependent utilities that take into account the suffering of friends and relatives can be accommodated by welfare theory and BCA; however, a more entractable problem is accommodating what Sen (1982) calls "committment", where the suffering of others does not affect a persons own welfare, but he or she is still prepared to pay to reduce it. Ideally, the economist would have direct information on WTP for each distinct risk situation and for each exposed population, but these valuations can be quite costly (Freeman, 1986).

Lacking estimates on the many different types of risks that policy makers address, a natural way to proceed is to borrow a WTP measure from one risk context for the purpose of applying it to another, e.g., compensation for occupational hazards (Viscusi, 1978; Dardis, 1980; Marin and Psacharopoulos, 1982).[9] Yet, applying WTP measures across the board may be wholly inappropriate. At the least, adjustments to reflect differences in the population and the risk context are needed. Economists have generally focused on the more tangible aspects of the risk context in adjusting these measures, factors such as the age of the affected population, pain and suffering of those concerned, bereavement, and social disruption. Much less attention has been given to the concerns that characterize "collective-fatality" risk aversion, such as the degree to which the risks are assumed voluntarily and the control one feels one has in the risk situation. The extent to which the economist remains loyal to the principle of risk neutrality will depend ultimately on the extent of these adjustments. Bailey (1980), for example, considers them to be minor:

> Although there is no rational case for spending a huge sum to avoid a death from one cause while refusing to spend a relatively small sum to avoid a death from another cause, it can be shown that rational, well-informed citizens do not equalize these incremental sums precisely in private choices. Hence a policy based on such choices will allow some, albeit minor, differences in these sums to remain.

Bailey, like Zeckhauser and Viscusi quoted above, reflect a widely-held view of economists, that large deviations from risk neutrality on an individual level are irrational. In their

[8]Zeckhauser (1976) has reasoned that catastrophic potential may actually reduce WTP in the aggregate since if whole families and communities die together, this reduces pain and suffering!

[9]This information might also be obtained by the "hedonic price" approach which compares, for example, property values of those in a risk area with those in safer areas, or by the "contingent valuation" approach where people are simply asked what they are willing to pay to avoid a risk in a specified context. Both methods have serious drawbacks, but even this imprecise information can be useful.

view, policies based on individual choice are pushing policy makers irrationally into making decisions according to how many people will be frightened rather than how many will be killed (Fremlin; 1984). Framed in this way, the whole issue of risk-neutrality may be seen as resting on the question of individual rationality, which I turn to below.

4 Risk Aversion, Risk Equity, and Rationality

The economic concept of the "rational person" is one who acts in his or her own self interest and maximizes his or her utility. If a person's behavior conforms to certain simple axioms, the existence of an ordinal utility function can be inferred from the individual's choices. Von Neumann and Morgenstern (1947) extended this concept to choices under uncertainty. With a more complex set of axioms,[10] the existence of a utility function with cardinal properties can be inferred from a person's choice among risky alternatives. The so–called weak and strong axioms of revealed preference assure that the "rational person" is an expected utility maximizer. (For a discussion and critique, see Sen (1982b).)

The expected utility model of choice was the basis of an early challenge to risk aversion by Keeney (1980). Without questioning outright the morality of allowing risk-averse preferences to influence public choices, Keeney nonetheless questions the consistency of these preferences with other seemingly laudable goals for public, risk-management decisions; in other words, he questions their rationality. According to Keeney, a policy maker who makes choices in a manner consistent with the expected utility model cannot be both risk averse and at the same time display a preference for distributing risks in an equitable manner. Risk equity in this context means spreading risks as evenly as possible over the population.

As a way of summarizing Keeney's argument, consider again a group of 1000 people, of which only two are exposed to a very high annual risk of death. We will call this RISK D which can be written as follows:

$$\text{RISK D: } ((1/2)_1,\ (1/2)_2,\ 0_3\ ,0\ ,.....0_{1000}) = 1 \text{ Expected Fatality}$$

Not unlike RISK A, B, and C in Figure 1, the statistical expectation of RISK D is one fatality. In sharp contrast, however, RISK D distributes the death warrants rather unevenly, or inequitably. Two of the 1000 individuals face a 50% chance of dying from the risk in question; the remaining 998 face no risk at all. A preference for risk sharing suggests that RISK A would be preferred to RISK D, or

$$(1) \quad A \ > \ D \quad \text{(Risk Equity).}$$

We can recall that risk aversion is consistent with the following preference:

$$\begin{array}{lll} (2) & A > B & \text{(Risk Aversion)} \\ & A > C. & \end{array}$$

Keeney's point is that (1) and (2) are inconsistent for anyone following the maxims of the EU model, according to which the attractiveness or value of a choice is summarized by its expected utility, or the sum of the utilities of each possible outcome weighted by their respective probabilities of occurrence. In other words, if A > D, then $A \not> B$ and $A \not> C$.

[10]These include the axioms of complete–ordering, continuity, independence, unequal probability and complexity.

Without showing the formal proof, Keeney's conclusion can be understood as follows: The more inequitably the risks are spread, or the more they are concentrated among a few persons, the less likely is the chance of a large number of deaths or the chance of exceeding the statistical expectation (in this case, 1 fatality). RISK A, for all its virtues in distributing the risks over many people, still has a very small chance of a large number of deaths—the worst possible case is 1000 fatalities. In contrast, the worst possible case for RISK D is 2 fatalities. It follows that a preference for RISK A over RISK D on equity grounds is at the same time a preference for accepting the remote possibility of many, albeit isolated, deaths. This is consistent only with a risk-prone utility function. RISK A > RISK D implies risk proneness; RISK A > RISK B or C implies risk aversion. A person cannot be both risk averse and risk prone.

For those who feel uncomfortable with this result, who perhaps wish to stick with a preference ordering that has been ruled out as inconsistent or irrational by Keeney, namely, A > B > C > D, some more explanation may be called for. What drives Keeney's result is his full acceptance of the expected utility model. EU theory excludes the possibility of "collective-fatality" risk aversion, that is, it excludes the possibility that people prefer "controlling the odds" or "flipping the coin themselves". Only 'multiple-fatality' risk aversion, or an aversion to accidents claiming many lives at once, is permissible within the EU logic. EU theory dictates that only decision outcomes can have an associated utility. How the probabilities are generated (who flips the coin) cannot be a factor in a person's preference function.[11] The notion of collective–fatality risk aversion, where an extra premium is placed on individual trials, is fully irrational in this context. No distinction can be made between 1000 deaths on the highway and 1000 deaths from a plane crash. It follows that the only variable of concern to Keeney's decision maker (who chooses between RISKS A, B, C, and D) is the probability distribution over fatalities. Keeney's concept of risk aversion translates then into the non-linear utility function over fatalities shown in Figure 3, where DU(2 deaths) > 2 DU(1 death). [12]

This illustrates the core of this discussion. If risk–averse preferences are based only on the notion of non-linear preferences as shown in Figure 2 (which is implicit in most economic analysis), then the whole concept of risk aversion might easily be challenged, as indeed Keeney did, on moral grounds. This interpretation implies that one values lives lost collectively more than those same lives lost individually. But introducing another plausible concept of risk aversion, that people are not indifferent to how the odds are generated or *the manner in which the risks are imposed*, changes the argument. Sen (1982a) argues that this egoistic view of economic rationality is a consequentialist view: judging acts by consequences only (p. 104). Viewing the lotteries shown in Figure 1 only in terms of expected deaths is likewise consequentialistic and does not admit legitimate feelings about how the odds are generated. As Sen notes "Sometimes the lack of personal gain in particular acts is accepted by considering the value of rules of behavior. ...the exclusion of any consideration other than self–interest seems to impose a wholly arbitrary limitation in to notion of rationality" (p.104). The Nichols and Zeckhauser example (Risk B in Figure 1) is interesting from this perspective. The chemical in question poses either a very high risk

[11]This type of preference is ruled out by the continuity axiom and the axiom of complexity.

[12]Collective disasters (RISK C) rank lower than more individual disasters with the same expected deaths (RISKS A&B) for the simple reason that collective disasters have the additional property of skewing the probabilities toward more serious, 'worst case' outcomes (RISK C has a greater chance of 1000 deaths than RISK B and RISK A). Risk A is preferred to Risk C, not because of the independent versus conditional nature of the probabilities, but because Risk A has a lesser chance of a large number of deaths than Risk C.

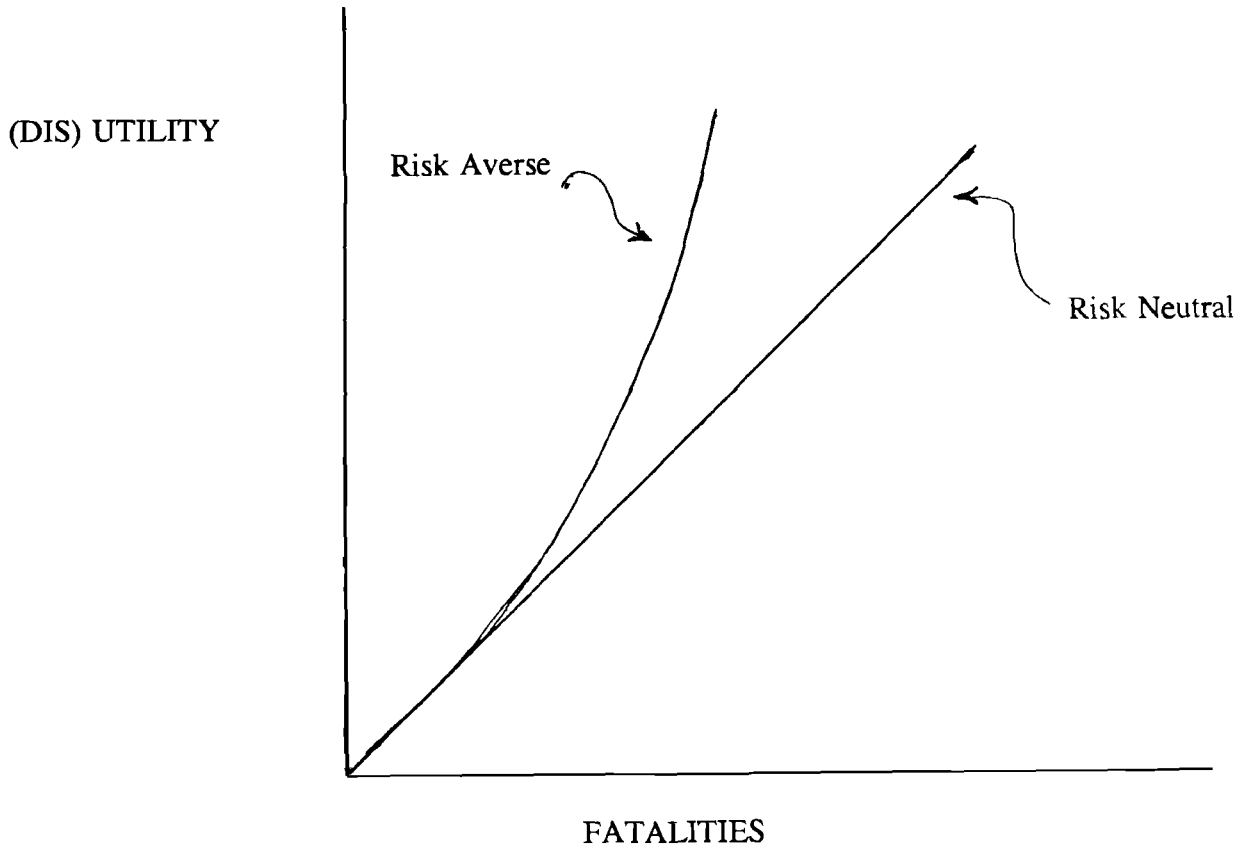

Figure 2: Risk-Averse and Risk-Neutral Utility Functions

(1/100) to each of the 1000 people exposed or no risk at all. The carcinogenity question will be resolved independently of the people involved—a kind of social or group lottery. If positive, then the risk of contacting cancer will likely be viewed as more individually determined, depending on a persons's immune system, lifestyle, age, etc.[13] As such, RISK B is a hybrid of RISK A and C—partially a group lottery and partially independent trials. It may be this "group lottery" aspect of cancer—the exogenous chemical for which little is known and for which the individual has little control—that creates such anxiety about this particular disease.

Yet this "collective–fatality" version of risk aversion is a clear violation of the expected utility model, and, for this reason, might be considered irrational. This concept of irrationality is so closely tied to expected–utility maximization as a normative decision rule that it is important to ask if expected utility is an appropriate yardstick for decision rationality. In a comprehensive review, Schoemaker (1982, p.548) concludes that EU theory fails as both a descriptive and predictive model since it does not take into account various psychological principles of judgement and choice. Very real factors that affect personal choices and are excluded by the EU model include, for instance, regret, ambiguity, and status quo biases.[14] A wealth of variants to the EU model exist that try to capture how people actually make risky choices.

That individuals do not deal well with EU maximization might strengthen the case for more formal decision analysis (and more explicit risk-benefit tradeoffs) as a way of improving the suboptimal nature of intuitive, unaided decision making. In other words, those who prefer RISK A over RISK C in Figure 1 might appreciate being corrected so as not to be in violation of the EU model. Schoemaker, however, points out a number of experiments that lead one to question even the normative appeal of the model. Some of the biases that lead people to behave differently than the model would predict, Schoemaker suggests, may be so basic as to render the normative theory inoperational. In addition, many individuals persist in violating the model even when their biases are made apparent to them. To the extent that this is the case in the choice situation illustrated in Figure 1, Keeney's analysis showing the incompatibility of risk aversion and risk equity (and ultimately the rationality of risk aversion) is based on an inappropriate normative model of choice.

[13]The concept of independent trials is, itself, partly one of perception. For instance, the probability of dying from an earthquake or even a plane crash is not altogether a "group lottery" if each individual has some chance of survival.

[14]EU theory, for example, ignores the fact that people may regret some choices more than others. Regret is magnified by the human tendency to think that an outcome was more predictable before it occurred than it actually was. This hindsight bias (Fischhoff, 1975) is perhaps more acute with low probability events. For example, there is a tendency to argue that the Chernobyl and TMI accidents were both fairly predictable from the evidence that was available at the time.

More then 30 years ago, Allais (1953) and later Kahnemann and Tversky (1979) posited that "certainty" in decision making played a role not anticipated by EU theory. A certain outcome is often preferred to a probabilistic outcome with a higher expected value. A related concept of "ambiguity" or vagueness about probabilities is another attribute ignored by EU theory. Recent research shows that even professional insurers are concerned about the certainty of the probability estimates, which helps explain why the insurance market for many low-probability risks (e.g., nuclear power and environmental pollution) is functioning so poorly in the U.S. and in Europe (Hogarth and Kunreuther, in press). Both "certainty" and "ambiguity" may play a role in risk aversion, e.g., the independent trials of Risk A (Figure 1) result in less ambiguity about the outcome (depending on the number of trials) than the lottery of Risk B.

5 Risk Aversion and Citizen Choice

Advocates of risk-neutral public choices point not only to the equity issues involved, but also to public misperceptions which may lead policy makers away from risk-neutral positions. An important condition for benefit–cost calculations to reflect genuine welfare changes is that the affected public must be reasonably well informed about the issues in question. Early psychological studies shed some doubt on this condition in that the publics' estimates of probabilities in many risk situations, in contrast to the "experts"', differ markedly from actual death rates (Fischhoff, et.al (1978)). The large literature in cognitive psychology shows that people rely on various rules of thumb or "heuristics" in making these estimates. (Tversky and Kahneman, 1973, 1983). One of the most significant of these heuristics is called availability, which refers to the tendency for individuals to think an event is more likely to occur if actual instances of the event come to mind readily. For example, people living in a flood plain tend to underestimate the risk if a flood has not occurred in recent memory; alternatively, widely reported events, such as a plane crash, tend to evoke the opposite response. Nearly everyone considers himself or herself to be better than the average driver (Svenson, 1981), which can also be explained by the availability heuristic in that the daily experience of driving is usually uneventful.

Since judgements about risk are apparently often inconsistent, it is hard not to draw the conclusion that these judgements are inappropriate for guiding public policies. "Given our untrustworthy attitudes", Treuber (1991) questions if a consent-based approach can lead to anything but irrational public policies. This raises our core question: Do misperceptions of probability and inconsistent choices involving risk, even in areas of adequate public information and knowledge, constitute reasonable grounds for rejecting citizen choice (as some aggregate of individual preferences) as a basis for public risk–affecting programs?

There are, at least, three arguments that suggest the contrary:

- We do not demand this perceptional astuteness and consistency in other areas of private and public choice;
- expert judgements are also affected by these biases and heuristics; and,
- the probability estimates often do not reflect the full essence of the risk issue and therefore perceptional accuracy may play only a minor role.

It is noteworthy that in no other area of individual or public choice has so much attention been given to the issue whether personal preferences are correct or rational. We do not question consumer choice on the basis that consumers do not rate the performance of detergents, toothpaste, or automobiles correctly, or that they are subject to 'framing' in the form of packaging of products? We tolerate a large amount of seeming misjudgement or 'irrationality' in private choices because we prize individual autonomy. We try to minimize this 'misjudgement' by policies that encourage honest information and truth in advertising, but we do not scrutinize these choices for consistency and perceptional bias. Benefit-cost analysis purports to build this same individual autonomy into public policies by pricing public goods based on how individuals value them.

Should we single out 'risk' as an area where individual choice should be disqualified? One immediate response is that risk choices are matters of 'life and death', and therefore inconsistency and misperceptions should not be tolerated as a basis for public policy. Proponents of risk neutrality sometimes point out that if the government does not pursue schemes that save the "cheapest" lives first, then it is responsible for the sacrifice in lives which this inefficiency costs (for a discussion of this "responsibility" argument, see

Schraeder-Frechette, 1986). In this spirit, Lichtenstein, et.al. (1988, p. 5), who generally support public input as leading to better social decisions regarding risk, suggest that in certain circumstances social decision makers might in good conscience go against public opinion in order to make a better social decision.

While this argument can be valid, the grounds for singling out risk as a special case because of the life and death nature of the decisions has less merit. The reason is that virtually all private and public expenditures have an opportunity cost in terms of lives and thus can be valued in terms of life extension, an absurd idea that would rule out any program that did not extend life span.[15] The idea of imputing responsibility on the government for not allocating resources in such a way as to maximize life expectancy is therefore misleading by obscuring the fact that we tradeoff life expectancy for other amenities everyday. The extreme importance we place on human life in our "risk" choices may be more of a cloak for questioning the rationality for (the often anti-technology) risk-averse sentiments when viewed in terms of our otherwise reluctance to examine these same irrationalities and opportunity costs in other areas of consumer and public choice.

A second consideration that might deter us from abandoning citizen choice in risk policy making is that expert estimates suffer similar perceptional biases. Challenging the notion that the experts have objective estimates and the public has subjective feelings, Wynne (1989) argues that both the experts and the lay public have perceptions influenced by social biases and assumptions which are usually inadvertent and implicit in technical analyses. The assumptions of the experts may be no better than the lay public's, in fact, they may be worse. In particular, the dominant perception to the risk problem takes no account of the differences which can exist between real world risk systems and the models of such practices which scientists assume in order to structure and analyze a risk problem. (p. 35)

A third reservation about rejecting "incorrect or misperceived" public input is that the relevant probability estimates may have little to do with what people are concerned about in real risk situations. As expressed by Slovic (1987), although the public tends to misjudge probabilities, their conceptualization of risk is often much broader and richer than the experts' reliance on quantitative estimates of lives, injuries and property damage. Research suggests that the primary correlates of public concern are not mortality or morbidity rates, but characteristics such as potentially catastrophic effects, lack of familiarity and understanding, involuntariness, scientific uncertainty, lack of personal control by the individuals exposed, risks to future generations, unclear benefits, inequitable distribution of risks and benefits, and potentially irreversible effects. (Slovic, et.al., 1980; Covello, 1984).

These qualitative risk characteristics are correlated with each other across a wide range of hazards. This means that risk aversion, in its many interpretations, is largely inseparable from other factors of the risk environment. If 'lack of control' (characteristic of collective-fatality risk aversion) or 'catastrophic potential' (characteristic of multiple-fatality risk aversion) are intertwined factors influencing how people judge the seriousness of a risk situation, it is difficult, if not meaningless, to make the adjustments of which Bailey speaks to accommodate different risk contexts.

One concern of particular relevance to this discussion is different forms of control over

[15]Moreover, this life-saving numeraire is not at all straightforward when one considers Douglas and Wildavsky's (1982) argument that "richer is safer". This means that the opportunity cost of public expenditures, if lives are the numeraire, should also take into account the loss of life expectancy resulting from the net reduction in personal income from directing income to public uses. Keeney (1989) has calculated this opportunity cost to be in the order of one life per $7 million expenditure.

risks. Collective-fatality risk aversion may quite rationally be based on an apprehension of risks which individuals feel are out of their control. Wynne (1988) ties this to public trust in risk-controlling institutions, "judgements of the performance, attitudes, openness and overall 'social demeanor' of the relevant industries and regulatory bodies; ....in defining these factors out of the risk analysis frame, experts, consciously or not, are incorporating assumptions that these institutions will behave in a perfectly trustworthy manner" (p. 35).

It bears emphasis that this research shows that perceptions of the riskiness of situations that differ from the probability estimates do not always reflect an ignorance of the statistics. Risk perception research, as one of its pioneers notes (Otway, 1985), originated in the belief that the factors leading to 'misperceptions' could be identified and corrected. In spite of research findings to the contrary, conflicts about risk are still attributed to public misperceptions of risks and associated misunderstandings of science. The U.K. Royal Society (1985) attributes this misunderstanding as the root cause of the objective-perceived risk controversy between experts and the public.

The psychological research constitutes the most acknowledged challenge to the risk-neutrality prescription. A policy maker who is sensitive to public concerns cannot justifiably base his or her policies on mortality risk irrespective of its context. Moreover, conflicts over technologies which have the potential for large-scale effects often appear to arise from conflicting evidence on the probabilities. This has led many to frame the 'public-acceptance' issue in terms of individual, differential perceptions of the probabilities or 'hazardness', a notion that is supported by the richness of the contextual concerns brought to light by the risk perception studies.

Cultural theory (see Douglas, 1978, 1982; Thompson, 1983 a, b; Gross and Rayner, 1985; Wildavsky and Dake, 1991) challenges this thesis that differential risk perceptions are fundamental to societal risk conflicts. The aggregation of risk perception data, it is argued, has concealed underlying group phenomena that override individualistic notions of perception. Viewing individuals as the active organizers of their own perceptions, cultural theorists suggest that individuals choose what to fear in order to support their way of life (Douglas, 1978). Selective attention is given to risks according to cultural biases, that is, to worldviews or ideologies that entail deeply held values and beliefs.

These views, the anthropologists suggest, are rooted more in social contexts than individual perceptions of probability. People who feel that society should take more technological risks, even those of the LP/HC sort, can be described as patient, forbearing and orderly; this pro-risk personality of an obedient citizen deferential to authority fits well with the political culture of hierarchy. By contrast, those who perceive greater risk with respect to technologies and the environment consider autonomy and change important. This risk-averse personality holds for those who endorse egalitarianism. According to research by Wildavsky and Dake (1991), these cultural biases can predict risk-taking preferences better than measures of knowledge and at least as well as political orientation.

Accordingly, risk-taking preferences are rooted more in social contexts than in individual perceptions of probability and consequence. Each actor is perfectly rational, given his or her convictions as to how the world is. The situation is therefore one of plural rationalities (Schwarz and Thompson, p. 6). This theory rejects self interest, distorted by misperceptions of probability, as the motivating force behind policy debates. In Sen's (1982) terminology, individuals are **committed** to a world view or social ideal that goes beyond the utilitarian notion of self interest. According to Schwarz and Thompson, "the conceptual advance of cultural theory lies not in the rejection of the idea of competing interests, but in making it contingent upon the culturally induced biases in perception

of policy actors who are operating within a social arena that they themselves collectively shape and maintain" (p.49).

If individual interests are contingent upon cultural affiliation, what implications does this have for the welfare economist? Since the benefit–cost practitioner is little concerned about **why** people are willing to pay for risk reduction and can easily accept the notion that these preferences are socially determined and contingent upon cultural bias, cultural interpretations do not change the analysis. There may be some pragmatic implications, however, insofar as the practitioner must assert more care in obtaining a representative sample for the WTP measures, and extrapolations across different risks and different groups will be inappropriate. Viewed through a different lens, however, cultural theory presents a more formidable challenge to welfare economics. If preferences (and WTP measures) are determined more by individual adherence to different world views than by the actual benefits or costs of a policy in question, the BCA model may be less useful. The reason is that the conflicting groups may not accept a welfare concept based on the potential for a Pareto improvement when they feel strongly and morally committed to a position far beyond their personal risks or disbenefits. The welfare criteria may not be an acceptable policy criteria when debates are so polarized that conflicting groups view each other's preferences as unreasonable.

6 A Policy Example

The transport of very hazardous materials through densely populated areas is a topical and contentious policy area, and one characterized by LP/HC risks. As early as 1976, the City of New York banned the road transport of large-quantity shipments of radioactive materials through the city. Five years later, the US Department of Transportation (DOT) published a Final Rule (HM 164) which allowed the shipment by road throughout the nation of all types of radioactive materials. One avowed purpose of this rule was to override local prohibitions against such shipments.

New York City challenged HM-164 on the basis that DOT misjudged the significance of the unlikely possibility of a catastrophic accident. Relying mainly on estimates from the Sandia report,[16] According to one study (SANDIA, 1980), the most lethal credible accident was one involving plutonium shipments, which could cause an estimated 5 early fatalities, 1800 latent cancer fatalities, and 290 early morbidities, but only with the very low probability of around one in a million. This estimate, however, excluded the possibility of sabotage, which another report crudely estimated could result in as many as a million latent cancer fatalities. DOT claimed that HM-164 would have no significant impact on the human environment. This excluded the Department from preparing an Environmental Impact Statement (EIS) as required by the National Environmental Policy Act (NEPA) and from the 1978 "Research and Worst Case" regulation of the Council on Environmental Quality[17] (CEQ), which required an agency confronted with unavoidable uncertainty to include a worst case analysis and an indication of the probability or improbability of its occurrence.[18]

The District Court upheld New York City's challenge on the basis of the large uncertainties in the probability estimates, especially insofar as the SANDIA analysis did not adequately take into account human error and sabotage, and the gravity of the potential

[16]The Nuclear Regulatory Agency also published a risk assessment with similar estimates (NRC, 1977).

[17]The Council on Environmental Quality assigns binding procedures for compliance with NEPA.

[18]40 C.F.R., Section 1502.22.

consequences. Judge Sofaer acknowledged that "significance depends on the estimation of the credible consequences discounted by their improbability", yet, he went on to state that in some circumstances "significance will depend heavily on the gravity of the potential consequences, for some consequences are so grave that, however unlikely, their mere "credibility" makes the impact of an agency action significant. (Sofaer, 1982, p. 53)

If "mere credibility" of a disastrous consequence is significant, then this statement is a clear rejection of an expected-value rule or of risk neutrality.[19] Judge Sofaer justifies this position on the grounds that adhering to an expected-value rule bypasses any consideration of public concern about catastrophic events. In his opinion, DOT unjustifiably dismissed public concern as unwarranted.

Judge Sofaer's decision was reversed by the Appeals Court.[20] In the words of the Appeals Court judge, "Disquieting as it may be even to contemplate such matters", DOT's decision that a remote possibility even of a serious accident does not pose a significant risk for the human environment "cannot be said to be an abuse of discretion" (Newman, 1982, p. 41). According to a dissenting judge, however, "Given the 'unique (and) unknown risks' associated with and the 'highly controversial' nature of nuclear waste transport", the DOT decision was "arbitrary, capricious, and not in accordance with law", [21] (Oakes, 1982, p.5).

This case encapsulates the controversy and judicial insecurity regarding the treatment of LP/HC events. In fact, several earlier court cases had set opposing legal precedents regarding the interpretation of NEPA and the CEQ regulation,[22] In one notable case,[23] it was ruled that **no** limitations should be recognized on the remoteness of the probability or on the speculative nature of the data associated with catastrophic events. This was clearly moving away from a risk-neutral position. This oscillation in judicial interpretations of the National Environmental Policy Act and the drift of some courts away from a 'rule of reason', i.e., risk neutrality, led the CEQ to amend its regulation in 1986 such that only known or reasonably foreseeable consequences are considered in an Environmental Impact Statement.[24] By this revision, the CEQ clearly intended to rescind the prior duty of agencies to consider remote and conjectural consequences of major projects thus moving away from risk aversion toward more risk-neutral reasoning (Weiss, 1988).

Although closely intertwined with other issues, especially questions of what scientific

[19]It is perhaps of interest that this 'expected value rule' was not new to the courts. It was introduced into judicial interpretations of tort law as early as the 1920's by Judge Learned Hand who gave the standard for negligent conduct a simple mathematical interpretation: risk is expressed as the product of probability times consequences. Provided the product remains constant, an accident with low probability and high consequences is weighted equally in assessing liability with an accident with high probability and low consequences (see Yellin (1977) p. 982).

[20]Based on an earlier Supreme Court decision (Baltimore Gas and Electric Co. vs. Natural Resources Defense Council (Vermont Yankee IV), 462 U.S. 87, 101 (1983), which emphasized that courts should be deferential where an agency is making predictions within its area of special expertise (at the frontiers of science), the only role for the court was to ensure that the agency, in this case the DOT, had taken a 'hard look' at the environmental consequences.

[21]Commenting on this case, Henry (1986) argues that by allowing the DOT to avoid a worst case analysis required by the CEQ regulations, the Appeals Court "created a hole in NEPA large enough to drive a truckload of radioactive waste through" (p.565).

[22]See, for example, Sierra Club v. Sigler, 695 F.29 957 (5th Cir, 1983).

[23]Southern Oregon Citizens Against Toxic Sprays, Inc. v. Clark, 720 F.2d 1475 (9th Cir. 1984).

[24]The worst case analysis requirement was replaced by a 'threshold of reasonably foreseeable impacts based on credible scientific evidence' While worse case analysis focuses first on exposing the potential types of consequences of a federal action and then on determining their likelihood, the "credible scientific evidence" threshold concentrates on the uncertain probability of a known or reasonably foreseeable consequence. (Weiss, 1988. p.817).

evidence is permissible regarding very low probability events,[25] the issue of risk neutrality was central to the NYC case and more generally to U.S. legislative history dealing with LP/HC events. Given estimates of low probabilities and grave consequences, the underlying issue is whether an expected-value rule should be adhered to in view of the public's concerns about catastrophic accidents. DOT was clearly against weighing public concern or public risk aversion in its considerations. The public, DOT stated, is incapable of rationally appraising the consequences of accidents; it tends to be disproportionately influenced by potentially large consequences (Federal Register, 1980). The Agency expressed its explicit view that, where it finds public concern unjustified, it will refuse to regulate with a view to relieving this concern.[26] Moreover, DOT blamed irrational public anxiety as responsible for actions such as New York City's transport ban which tended to show local governments to be unfit to legislate in this area.

Not all U.S. agencies take such an explicit risk-neutral position. Recognizing how essential public acceptance is to its future, the nuclear power industry has explicitly addressed the quantification of public risk aversion. Exponential weighing functions on consequences with multiple, simultaneous fatalities are most often discussed (Farmer, 1967; Wilson, 1975; Ferreira and Slesin, 1976; Griesmeyer and Okrent, 1980; Higson, 1986). Another approach is to set absolute levels on the acceptability of individual and societal risks, regardless of the costs of avoiding them.[27] Many propose that these acceptability limits be more stringent for risks with catastrophic consequences (Versteeg, 1986; Levine, 1980). This form of risk aversion has been adopted in the safety targets of the U.S. Nuclear Regulatory Agency[28] and proposed by the International Atomic Energy Agency (Niehaus, 1987). In the Netherlands, risk aversion has been formalized in legislation dealing with safety targets for industrial facilities.

By explicitly stating that it would refuse to regulate with a view to relieving public concern if it found this concern unwarranted, DOT disregarded public risk–averse preferences in favor of a more risk–neutral position. This position violates the principles underlying benefit–cost analysis, most importantly that public preferences should guide public policies. DOT's paternalistic position raises an important question: At the most fundamental level, might there be some areas where the public would reject the principles underlying BCA in favor of a more paternalistic model of social choice? Might there be concerns and preferences which an individual holds on a personal level that he or she would not wish to affect public policy? As Lichtenstein, et.al. (1989) query, are there circumstances where

[25]For a good discussion of how agencies should deal with uncertainties in risk estimates, see Finkel (1990).

[26]DOT, Summary of comments, DOT Motion, EX. E, pt.H.4.at 4.

[27]The U.K. Health and Safety Executive suggests that acceptability for major societal risks can be estimated from risks that appear to be accepted in society. For instance, the industrial installations at Canvey Island were calculated to pose a risk of about 1 in 5000 per annum of a major accident causing more than 1500 casualties. When the Thames Barrier was built, the design specification insisted that the chances of its being overtopped by a freak tide should be less than 1 in 1000 per annum, which is the predicted approximate annual chance of an aircraft crash somewhere in the UK killing 500 or more people. According to the HSE, this suggests that society requires the risk of major accidents to be less than 1 in 1000 and if possible less than 1 in 5000 (U.K. Health and Safety Executive, 1987, p.24).

[28]The NRC safety goals set quantitative safety objectives only for the risk of prompt fatality to the average individual within a mile of a plant site boundary, which was equivalent to an average individual risk of death of 5 $times 10^{-7}$. After Chernobyl, which surprised the nuclear community with its few immediate fatalities, there was concern that this individual risk figure would permit an accident frequency that was too high. This risk aversion led the NRC to establish an additional safety guideline of 1 in 1 million per year frequency for large releases of radioactive materials from nuclear power plants (see Whipple and Starr, 1989).

social decision makers might, in good conscience, go against public opinion (as motivated by self interest) in order to make a better social decision?

These questions raise a number of issues that have been addressed extensively by economists and philosophers. Libertarians, for instance, consider the economic welfare paradigm as inappropriate for social choices on a number of grounds, but most basically because the Kaldor–Hicks criterion permits uncompensated losers (Kneese, 1991). Indeed, libertarians liken the governments' mandate to redistribute between winners and losers in environmental policy making with socialism and centralized planning. Sagoff (1993) recommends instead that policy be directed by a more explicit respect of property rights. In the New York City case, this would suggest that radioactive shipments are unacceptable unless those exposed to the risks are compensated adequately. Libertarians are not the only ones opposed to uncompensated risk imposition, however. In an early paper on this subject, Schulze and Kneese (1980) examined three ethical systems in addition to the libertarian and concluded that even in those cases where the benefits clearly outweigh the costs, each of the ethical systems (in contrast to that of BCA) reject uncompensated risks at least some of the time.

In advocating the protection of property rights for motivating environmental policy, Sagoff rejects the notion that preferences should play the dominant role in public policy. If government protects property rights or the status quo, and does not respond to the preferences and tastes of the public, this implies a more paternalistic view of public policy making which characterizes most European policy processes. Kopp (1993) rightly notes that the issue can only be decided as an individual value issue: "Are we more willing to live with our mistakes than we are willing to give up our freedom to make mistakes?" (p.12)

A more long–standing critique of the utilitarian basis of BCA is based on the rejection of self–interested preferences for determining public policy. Consider, for example, a person with an aversion to spiders, who is observed as killing each and every spider he encounters. From this behavior one can infer that spiders rank rather low in this person's utility. Yet, he may recognize the ecological importance of spiders and not wish that his personal preference for killing all spiders influences public policy; in fact, he may be opposed to a public policy that eradicates spiders. In the words of Sen (1982), this person has a commitment to spiders beyond his personal preference, and this commitment will not be reflected in his actual personal choices (and thus not in his derived utility function).[29] Whether or not these social preferences factor into the benefit–cost ledger will depend on whether an attempt is made to solicit them directly, e.g. willingness to pay to eradicate spiders, or whether the social preferences are a summation of preferences derived from personal behavior.

The other side of this issue is not the inclusion of altruistic, social concerns, but the exclusion of preferences that are deemed asocial or immoral, for example, based on jealousy or racial prejudices. Sen also addresses this issue, concluding that noone would suggest that all preferences are morally acceptable. Judgements therefore must be made about what should and what should not count in the ledger, or as Mishan (p.165) rather chauvinistically puts it, "what men of good will regard as reasonable". With a little reflection, it seems apparent that benefit–cost practitioners cannot operate fully on the premise that every effect on individual utility should be included in the analysis.

The dilemma is that where issues are strongly polarized and where underlying values on

[29]In an early paper, Harsanyi (1955) raised the possible distinction between a persons subjective preferences, which are motivated by self interest only, and his ethical preferences, which express impersonal social considerations.

the direction of society are at stake, there may be little social agreement or consensus on what views are reasonable or moral. Cultural theory proposes that "reasonableness" or "rationality" is very closely tied to social relationships and corresponding adherence to different myths of nature. Underlying any judgement on the reasonableness of preferences for benefit–cost purposes, therefore, is a judgement of what preferences are sanctioned by the organization doing the calculations. Thus, we find analyses which adjust for the risk context by adding factors for social disruption, for the pain and suffering caused to others and the age of the victims. Some studies find it 'reasonable' to adjust (usually arbitrarily) for the catastrophic potential. But these adjustments would be dismissed as trivial by egalitarians who might consider this same catastrophic risk as vastly more than social disruption, perhaps as one event in a chain of possible calamities that signal the breakdown of social/ecological systems—a contingency that must be avoided at all costs.

Where such polarization exists, what can be the role of benefit–cost analysis? Recognizing that the analyst cannot avoid making judgements about what to include in the analysis, the authority and effectiveness of BCA will depend ultimately on the trustworthiness and reputation of the political institution doing the assessing. This trust is waning. As Jasanoff (1991) observes, there is a growing reluctance for powerful social actors to accept expert knowledge as the basis for their political assessments of risk. Because of scientific pluralism, and also because of the relatively non-hierarchical organization of science as well as politics in the U.S., different constructions of risk often enter the debates. As a case in point, when a dissenting judge in the NYC case was presented with seemingly cut-and-dried evidence for risk neutrality in the form of comparisons of days of life lost from, eg., remaining unmarried, smoking a pack of cigarettes a day and the transportation of radioactive materials, he responded that "quantifications such as these are absurd" (Oakes, p.4). This comment reflects the polarized risk constructions apparent in the public debate.

Despite Judge Sofaer's misgivings, DOT's risk-neutral stance cannot be faulted on procedural or legal grounds; however it may prove to be short sighted. As Wynne (1983) has argued, the authority and effectiveness of government policies for regulating and managing hazardous activities (and the effectiveness of formal analyses) often founder upon disagreements about risks. These disagreements have led to a disturbing loss of public confidence in regulatory institutions, which means that their viability may depend as much on building public trust as in managing risks efficiently. This will at a minimum require a meaningful dialogue with public groups and a sincere effort to accommodate public concerns.

Where trust and legitimacy are threatened and analytical expertise challenged, a BCA can probably contribute little to building a stable social consensus. To quote Mishan one last time: "It follows that there will be many issues in the formulation of economic policies which cannot call for guidance on a welfare economics that raises itself only on those ethical premises sanctioned by (benefit–cost analysis) ...it should be clear that while some issues of economic policy are fit subjects for welfare economics, it has little to contribute to other issues" (p.166). Commenting more generally on social choice theories (including the economists' welfare theory), Sen (1986) also stresses the importance of recognizing the heterogeneity of different types of social issues and choices and the folly of applying any given set of complete ideas to all types (p. 238).

This does not mean that benefit–cost analysis does not have a role in controversial LP/HC cases. Agencies can rely on it to support and justify their policy decisions by explicitly bringing in public concerns. But the analyses will necessarily include agency judgements, from framing the issues to screening the preferences for 'reasonableness',

and these judgements will naturally reflect the goals and interests of the agencies doing the analyses. The advantage of the benefit–cost methodology is that it makes these judgements more transparent and open to public scrutiny. Its usefulness then rests on its contribution to the political debate and process, but not on it serving as a by-pass to this process.[30]

7 Concluding Remarks

In this paper, I have addressed the priority our social decision makers should place on preventing or reducing the consequences of potentially disastrous events that will occur with very low frequency. The question I pose is whether risk neutrality, where the social objective is to reduce the loss of life as efficiently as possible given competing demands on resources, should be a guiding principle even if individuals often deviate from this objective in their personal choices? At the core of this issue lies a fundamental question: To what extent, under which circumstances, and how should individual preferences count as legitimate input for public policy?

While there is always room for improvement in the public understanding of technical risks, and admittedly most people are poor estimators of probabilities, I have argued that this should not disqualify informed public preferences as influencing "life and limb" policy decisions. Even contemplating the loss of life that deviations from risk neutrality implies, this does not warrant exceptional treatment of such decisions since, in fact, all allocative decisions have an opportunity cost in terms of human lives. Another argument for not disqualifying public concern in favor of "expert objectivity" is that the expert estimates of low probability events are also subject to heuristics and analytical biases. Finally, the probability estimates may not constitute the essence of the hazard situation for many people, and therefore public decisions should be based on more than the death count. In the words of Judge Sofaer, "Public reaction is a manifestation of collective wisdom based on human experience. It should not be lightly dismissed as unscientific" (Sofaer, 1982, p.76). That DOT labeled public concerns as "irrational", or "as disqualifying local governments to legislate in these areas", may ultimately reduce DOT's authority in future routing or other risk issues. This may lead to further polarization, and, in so doing, distort and impoverish the policy-making process.

As for catastrophe avoidance, there is ample evidence that "catastrophic potential" is so closely linked to other psychological dimensions of concern such as feelings of personal control and social equity, that risk aversion cannot be attributed solely to an aversion against fatalities in large numbers. My argument against risk neutrality, or for public risk-averse preferences serving as a legitimate input to benefit–cost calculations, begs however the question whether benefit–cost analysis is an appropriate criterion on which to base public policy on controversial risk issues. If, as cultural theory suggests, individual preferences for risk-mitigating policies are determined more by social affiliation and cultural biases than by survival interests, then it is doubtful that the Kaldor–Hicks principle underlying benefit–cost calculations can be a socially-acceptable, normative criterion for valuing public policies, at least to the extent that the disparate preferences are viewed as "unreasonable" or "irrational". In these cases, Mishan appears correct in his assessment

[30]Hardly any major environmental statutes require explicit estimates of benefits and costs. Moreover, although the role of BCA has been strengthened by executive orders mandating that benefit–cost analysis accompany any new proposed or final regulations, regulatory agencies have often resisted (Freeman and Portney, 1989).

that if we base our social choices purely on utility, they will run into ethical objections and lose credibility.

This means simply that benefit–cost reasoning cannot and should not be the final word in mitigating or reducing LP/HC risks, but it can be an important part of the debate. If the analytical judgements are made explicit, a benefit–cost analysis has the advantage of openly addressing public concerns and public risk aversion as well as exposing the opportunity costs of risk-averse policies. As I argued with regards to the controversial issue of transporting radioactive materials through New York City, analyses which explicitly address public concerns may be more appropriate than uncritically accepting a position of risk neutrality. Society's system for managing risks to life and limb is not then flawed by deviations from risk neutrality, but this system is flawed if citizens' concerns are ignored or deemed irrational.

APPENDIX I

The amount a person is willing to pay to reduce his or her mortality risk can be illustrated with the simple indifference curve shown in Figure 2. In this figure, a person's wealth (W) is shown on the vertical axis and his or her probability of dying from the cause in question (PD) is shown on the horizontal axis. The indifference curve (I) represents combinations of W and PD that have the same utility. The person represented in Figure 2 would be willing to pay a maximum amount of $W_1 - W_2$ to reduce the risk from 1/1000 to zero, and thus suffer no loss in utility. This "willingness to pay" corresponds to the concept of "compensating variation" (CV), defined as the money transfer necessary, following some economic change, to maintain the individual's welfare at its original level.

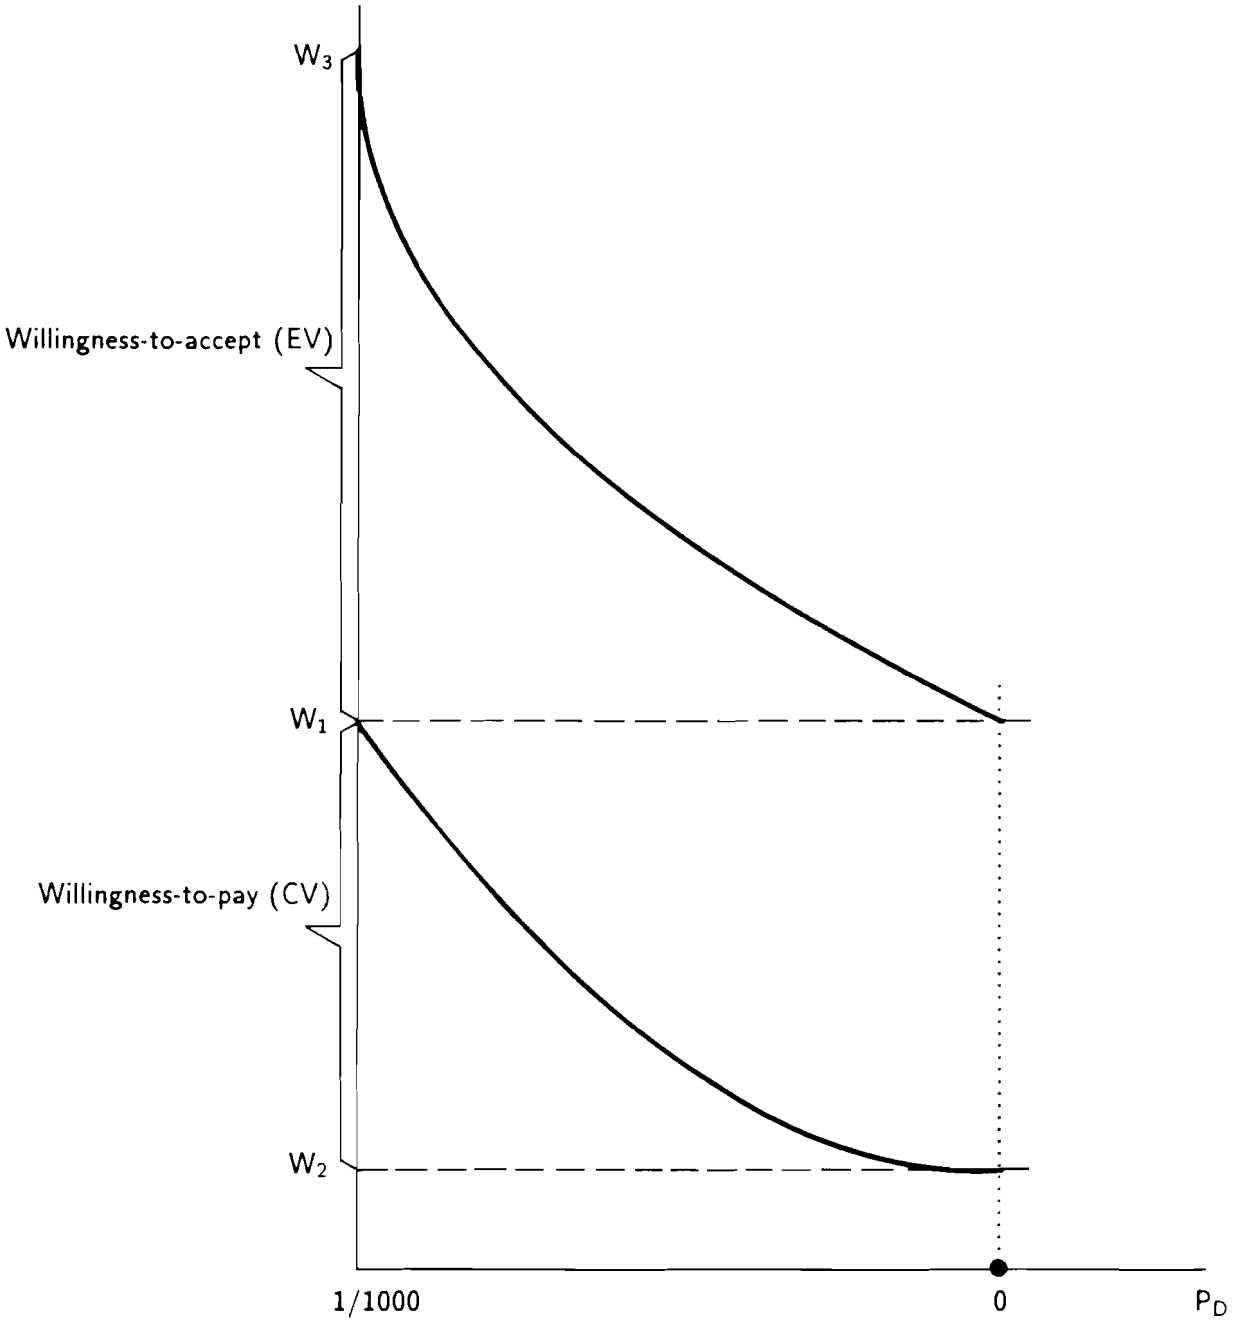

Figure 3: An Indifference Curve for Wealth and Fatality Risk.

References

Adams, J. (1989). Unsustainable Economics. A review for *International Environmental Affairs*, University College, London.

Allais, M. (1953). Le Comportement de L'Homme Rationel Devant le Risque, Critique des Postulates et Axiomes de Lécole Americaine. *Ecometrica* 21:503-546.

Bailey, M.J. (1980). *Reducing Risks to Life: Measurement of the Benefits*, American Enterprise Institute, Washington, D.C.

Boulding, K. (1969). Economics as a Moral Science, *American Economic Review.* 59:1-12.

Cohen, B. and I. Lee (1979). A Catalog of Risks, *Health Physics* 36:707.

Comar, C. (1979). Risk: A Pragmatic De Minimus Approach, *Science* 203:319.

Covello, V. (1984). Social and Behavioral Research on Risk: Uses in Risk Management Decisionmaking, *Environmental International* 4 (Dec. 1984).

Dardis, R. (1980). The Value of Life: New Evidence from the Labor Market. *American Economic Review 50:1077-1082.*

Douglas, M. (1978). *Cultural Bias.* Occasional Paper No. 35, London: Royal Anthropological Institute of Great Britain and Ireland.

Douglas, M. (1982). *Essays in the Sociology of Perception.* Routeledge and Kegan Paul, London.

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

Farmer, F.R. (1967). Reactor Safety and Siting: A Proposed Risk Criterion. *Nuclear Safety*, 8:539-548.

Ferreira J. and L. Slesin (1976). Observations on the Social Impact of Large Accidents. Massachusetts Institute of Technology, Cambridge, Mass.

Finkel, A. (1989). Is Risk Assessment Really too Conservative? Revising the Revisionists. *Columbia Journal of Environmental Law.* 14.2.

Fischhoff, B. (1975). Hindsight Does Not Equal Foresight: The Effect of Outcome Knowledge on Judgement under Uncertainty. *Journal of Experimental Psychology: Human Perception and Performance* 1:288-299.

Fischhoff, B., P. Slovic, S. Lichtenstein, M. Layman and B. Combs. (1978). Judged Frequency of Lethal Events. *Journal of Experimental Psychology: Human Learning and Memory* 4:551-578.

Freeman, A.M. (1989). Ex Ante and Ex Post Values for Changes in Risks, *Risk Analysis* 9: 309-317.

Freeman, A.M. (1989). Benefit-Cost Analysis of Policies toward Risk. Discussion Paper QE 89-13. Resources for the Future, Washington, D.C.

Freeman, A.M. and P.R. Portney (1989). Economics and the Rational Management of Risk, Center for Risk Management, Discussion Paper CRM 89-05, Resources for the Future, Washington. D.C.

Freeman, A.M. and P.R. Portney (1989). Economics Clarifies Choices about Managing Risk, *Resources* 95:1-4.

Fremlin, J.H. (1984). *Power Production: What are the Risks?* Adam Hilger, Bristol and Boston.

Green, C. and R. Brown (1978). Life Safety: What is it and how much is it worth? Building Research Establishment Current Paper, Hertfordshire.

Griesmeyer, J. and D. Okrent (1980). Risk Management and Decision Rules for Light Water Reactors, *An Approach to Quantitative Safety Goals for Nuclear Power Plants*, NUREG-0739, U.S. Nuclear Regulatory Commission, Washington D.C.

Gross, J. and S. Rayner. (1985). *Measuring Culture*, New York: Columbia University Press.

Harsanyi, J. (1955). Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility, *Journal of Political Ecoomy*, 63:315.

Häfele, W. (1979). Energy, Science, Technology, and the Human Prospect (C. Starr and P. Retterbush, eds.) New York: Pergamon.

Henry, E. (1986). The Council on Environmental Quality's Research and Worst Case Regulation: The Recent Litigation, *Oregon Law Review* 64:547-570.

Higson, D.J. (1986). Some Questions on the Development of Probabilistic Risk Criteris. Paper presented at the Technical Committee Meeting, "Status, Experience, and Future Prospects for the Development of Probabilistic Safety Criteria, Vienna, 27-31 January 1986.

Hogarth, R. and H. Kunreuther (in press). Risk, Ambiguity, and Insurance, *Journal of Risk and Uncertainty.*

Jasanoff, S. (1991). American Exceptionalism and the Political Acknowledgment of Risk, *Daedalus* 119:61-82.

Kahnemann, D. and A. Tversky (1979). Prospect Theory, *Econometrica* 47:263-292.

Kasperson, R. and K. Pijawka (1985). Societal Response to Hazards and Major Hazard Events: Comparing National and Technological Hazards. *Public Administration Review*: Special Issue, 7-18.

Keeney, R.L. (1980). Utility Functions for Equity and Public Risk, *Management Science*, 26: 345-353.

Keeney, R.L. (1990). Mortality Risks Induced by Economic Expenditures, *Risk Analysis* 10: 74-81.

Kneese, A. (1992). Benefit-Cost Analysis: The Historical and Ethical Context, Draft Paper. Resources for the Future, Washington D.C.

Kopp, R. (1993). Environmental Economics: Not Dead But Thriving. *Resources* 111:7-13.

Kunreuther, H., R. Ginsberg, L. Miller, P. Sagi, P. Slovic, B. Borkin and N. Katz (1978). *Disaster Insurance Protection: Public Policy Lessons*, New York: Wiley.

Lave, L. (1989). Risk Assessment and Regulatory Priorities, *Columbia Journal of Environmental Law* 14:307-314.

Levine, S. (1980). Various Applications of Probabilistic Risk Assessment Techniques Related to Nuclear Power Plants. Paper presented at the Annual Meeting of the National Safety Council, Chicago, Oct. 1980.

Lichtenstein, S., R. Gregory, and P. Slovic (1988). When Lives Are In Your Hands: Dilemmas of the Societal Decision Maker. Unpublished draft paper.

Linnerooth, J. (1982). Murdering Statistical Lives...? *The Value of Life and Safety*, Jones-Lee, ed. North Holland, Amsterdam.

Marin, A. and G. Psacharopoulos (1982). The Reward for Risk in the Labor Market: Evidence from the United Kingdom and Reconciliation with Other Studies, *Journal of Political Economy* 90:827-853.

Maxey, M. (1979). Managing Low-Level Radioactive Waste, *Low-Level Radioactive Waste Management*, (Watson, J. ed.). Williamsburg, Virginia, Health Physics Society.

Mendeloff, J. and R. Kaplan (1989). Are Large Differences in "Lifesaving" Costs Justified? A Psychometric Study of the Relative Value Placed on Preventing Deaths, *Risk Analysis* 9:349-363.

Mishan, E.J. (1988). *Cost-Benefit Analysis*, Unwin Hyman Press, London.

National Safety Council (1987). Accident Facts, National Safety Council, Chicago.

Neumann, J. von, and O. Morgenstern (1947). *Theory of Games and Economic Behavior.* Princeton, N.J.

Newman (1983). Appeal from the May 6, 1982, Judgment of the District Court for the Southern District of New York. 539 F. Supp. 1237, Dissenting Statement, US. Court of Appeals for the Second Circuit, Nos. 415, 451. August Term.

Niehaus, F. (1986). Status, Experience, and Future Prospects for the Development of Probabilistic Safety Criteria, Draft Report, International Atomic Energy Agency, Vienna.

Nichols, A. and R. Zeckhauser (1986). *Regulation* 10:13-19.

Oakes (1983). Appeal from the May 6, 1982, Judgment of the District Court for the Southern District of New York. 539 F. Supp. 1237, Dissenting Statement, US. Court of Appeals for the Second Circuit, Nos. 415, 451. August Term.

Organization for Economic Development (OECD) (1989). Risk Assessment and Risk Management for Accidents Connected with Industrial Activities, *Environment Monographs* 19:39 Paris.

Otway, H. (1985) Introduction, *Regulating Industrial Risks*, (Otway, H. and H. Peltus, eds.) Butterworths, London.

Pijawka, K., B. Cuthbertson, and R. Olson (1988). Coping with Extreme Hazard Events: Emerging Themes in Natural and Technological Disaster Research, *Omega* 18:281-297.

Rayner, S. and R. Cantor (1987). How Fair is Safe Enough? The Cultural Approach to Societal Technology Choice, *Risk Analysis* 7:3-9.

Royal Society of London (1985). *The Public Understanding of Science*, London, Royal Society.

Sagoff, M. (1993). Environmental Economics: An Epitath. *Resources* 111:2-7.

Sandia National Laboratories (1980). *Transportation of Radionuclides in Urban Environs: Draft Environmental Assessment*, Report prepared for the U.S. Nrc, July.

Schrader-Frechette, K. (1986). Risk-Cost-Benefit Methodology and Equal Protection, *Risk Evaluation and Management* (Covello, V. , J. Menkes, and J. Mumpower, eds.), 275-296, Plenum, New York.

Schulze, W.D. and A. V. Kneese (1980). Risk in Benefit-Cost Analysis, *Evaluation Studies Review Annual*, 263-270.

Schwarz, M. and M. Thompson (1990). *Divided We Stand: Redefining Politics, Technology and Social Choice*, University of Pennsylvania Press, Phila.

Sen, A. (1982a). Rational Fools: A Critique of the Behavioural Foundations of Economic Theory. In *Choice, Welfare and Measurement*, Blackwell Press, 84-106.

Sen, A. (1982b). Choice Functions and Revealed Preference, *Choice, Welfare and Measurement.* Blackwell, 41-54.

Sen, A. (1986). Foundations of Social Choice Theory: An Epilogue. In *Foundations of Social Choice Theory.* Jan Elster and Aanund Hylland, eds., 213-248. Cambridge University Press.

Shoemaker, P. (1982). The Expected Utility Model: Its Variants, Purposes, Evidence and Limitations, *Journal of Economic Literature* 20:529-563.

Slovic, P., S. Lichtenstein, and B. Fischhoff (1980). Facts and Fears: Understanding Perceived Risk, in R. Schwing and W. Albers (eds), *Societal Risk Assessment: How Safe is Safe Enough?* (Plenum Press, New York, 1980), 181-216.

Slovic, P. (1987). Perception of Risk, *Science* 236:280-285.

Sofaer, A. D. (1982). Opinion, United States District Court, Southern District of New York, 81 Civ. 1778 (ADS).

Starr, C. and C. Whipple (1980). Risks of Risk Decisions, *Science* **208**.

Svenson, O. (1981). Are We All Less Risky and More Skillful Than Our Fellow Drivers Are? *Acta Psychologica*, 47:143-148.

Thompson, M. (1983a). Postscript: A Cultural Basis for Comparison, in H.C. Kunreuther, J. Linnerooth, et. al., *Risk Analysis and Decision Processes*, Springer Verlag, Berlin.

Thompson, M. (1983b). The Aesthetics of Risk: Culture or Context? In R. Schwing and W. Albers (eds.), *Societal Risk Assessment*, New York: Plenum.

Treuber, A. (1991). Justifying Risk, Daedalus 119:235-254.

Tversky, A and D. Kahneman (1973). Availability: A Heuristic for Judging Frequency and Probability, *Cognitive Psychology*, 5:207-232.

Tversky, A. and D. Kahneman (1983). Exceptional vs. Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment, *Psychological Review* 90:293-315.

U.K. Health and Safety Executive (1987). *The Tolerability of Risk from Nuclear Power Stations*, London: Her Majesty's Stationery Office.

U.K. Royal Society (1992). Risk: Analysis, Perception and Management. London: The Royal Society.

U.S. Nuclear Regulatory Commission (1977). *Final Environmental Statement on the Transportation of Radioactive Material by Air and Other Modes*, Nureg-0170, December.

U.S. Federal Register (1980). 45, 7141.

Versteeg, M.V. (1986). External Safety Policy in the Netherlands: An Approach to Risk Management. Paper presented at the Technical Committee Meeting, "Status, Experience, and Future Prospects for the Development of Probabilistic Safety Criteria", Vienna, 27-31 Jan.

Viscusi, K. (1978). Labor Market Valuations of Life and Limb: Empirical Evidence and Policy Implications, *Public Policy* 26, 339-386.

Weiss, C. (1988). Federal Agency Treatment of Uncertainty in Environmental Impact Statements Under the CEQ's Amended NEPA Regulation Sec. 1502.22: Worst Case Analysis or Risk Threshold? *Michigan Law Review* 86, 777-820.

Whipple, C. and C. Starr (1989). Nuclear Power Safety Goals in Light of the Chernobyl Accident, *Risk Assessment in Setting National Priorities*, (J. Bonin and D. Stevenson, eds.) Plenum, New York.

Wildavsky, A. and K. Dake (1991). Theories of Risk Perception: Who Fears What and Why? *Daedulus* 119: 41-61.

Wilson, R. (1975). The Costs of Safety, *New Scientist* 68:274-275.

Wynne, B. (1983). Redefining the Issues of Risk and Acceptance: The Social Viability of Technology, *Futures* 15:13-32.

Wynne, B. (1989). Frameworks of Rationality in Risk Management: Towards the Testing of Naive Sociology, *Environmental Threats: Social Science Approaches to Public Risk Perceptions* (J. Brown, ed.), Belhaven, London.

Yellin, J. (1977). Judicial Review and Nuclear Power: Assessing the Risks of Environmental Catastrophe, *The George Washington Law Review* 45, 969-992.

Zeckhauser, R. (1976). Valuing Lives. In *Technological Risk Assessment*, NATO ASI Series, p. 133.

Zeckhauser, R. and W. K. Viscusi (1990). Risk Within Reason, *Science* 248: 559-564.