

Defining Risk

BARUCH FISCHHOFF

*MRC Applied Psychology Unit, Cambridge, Great Britain and Decision Research,
A Branch of Perceptronics, 1201 Oak Street, Eugene, Oregon 97401, U.S.A.*

STEPHEN R. WATSON

Emmanuel College, Cambridge University, Cambridge, Great Britain

CHRIS HOPE

Cavendish Laboratory, Cambridge University, Cambridge, Great Britain

ABSTRACT

Risk is the focal topic in the management of many activities and technologies. For that management to be successful, an explicit and accepted definition of the term “risk” is essential. Creation of that definition is a political act, expressing the definers’ values regarding the relative importance of different possible adverse consequences for a particular decision. Those values, and with them the definition of risk, can change with changes in the decisionmaker, the technologies considered, or the decision problem. After a review of the sources of controversy in defining risk, a general framework is developed, showing how these value issues can be systematically addressed. As an example, the approach is applied to characterizing the risks of six competing energy technologies, the relative riskiness of which depends upon the particular definition used.

Defining Risk

Managing the risks of technologies has become a major topic in scientific, industrial, and public policy. It has spurred the development of some industries and prompted the demise of others. It has expanded the powers of some agencies and overwhelmed the capacity of others. It has enhanced the growth of some disciplines, distorted the paths of others. It has generated political campaigns and countercampaigns. The focal ingredient in all this has been concern over risk. Yet, the meaning of “risk” has always been fraught with confusion and controversy. Some of this conflict has been overt, as when a professional body argues about the proper measure of “pollution” or “reliability” for

incorporation in a health or safety standard. More often, though, the controversy is unrecognized; "risk" is used in a particular way without extensive deliberations regarding the implications of alternative uses. Typically, that particular way follows custom in the scientific discipline initially concerned with the risk.

However, the definition of "risk," like that of any other key term in policy issues, is inherently controversial. The choice of definition can affect the outcome of policy debates, the allocation of resources among safety measures, and the distribution of political power in society. The present essay begins with an analysis of the key sources of controversy in this definition. It proceeds to advance a highly flexible general approach to defining "risk." Finally, it demonstrates the approach with an analysis of the comparative risks of different energy technologies, showing that the relative "riskiness" of those technologies depends upon the definition used. No definition is advanced as the correct one, because there is no one definition that is suitable for all problems. Rather, the choice of definition is a political one, expressing someone's views regarding the importance of different adverse effects in a particular situation. Such determinations should not be the exclusive province of scientists, who have no special insight into what society should value. As a result, the present approach is designed to offer a way to generate definitions of risk suitable for many problems and value systems.

Dimensions of Controversy

Objectivity. Technical experts often distinguish between "objective" and "subjective" risk. The former refers to the product of scientific research, primarily public health statistics, experimental studies, epidemiological surveys, and probabilistic risk analyses. The latter refers to non-expert perceptions of that research, embellished by whatever other considerations seize the public mind. This distinction is controversial in how it characterizes both the public and the experts.

Although it is tempting (and common) to attribute disagreements between the public and the experts to public ignorance or irrationality, closer examination often suggests a more complicated situation. Conflicts often can be traced to unrecognized disagreements about the topic, including what is meant by "risk." When the public proves misinformed, it is often for good reasons, such as receiving faulty (unclear, unbalanced) information through the news media or from the scientific community (Lichtenstein et al., 1978). In some instances, members of the lay public may even have a better understanding of specific issues (or for the definitiveness of knowledge regarding them) than do the experts (Cotgrove, 1982; Wynne, 1983).

Along with these elements of objectivity in public opinion, there are inevitably elements of subjectivity in expert estimates of risk. Within the philosophy of science, "objective" typically means something akin to "independent of observer." That is, any individual following the same procedure should reach the same conclusion. However meritorious as a goal, this sort of objectivity can rarely be achieved. Particularly in complex, novel areas, such as risk analysis, research requires the exercise of judgment. It

is expert judgment, but judgment nonetheless. Even in those orderly areas for which public health statistics are available, interpretative questions must be answered before current (or even historical) risk levels can be estimated: Is there a secular trend (e.g., are we sitting on a cancer time bomb)? Is the effect of predisposing causes (e.g., poor nutrition) underestimated because deaths are typically attributed to immediate causes (e.g., pneumonia)? Are some deaths deliberately miscategorized (e.g., suicides as accidents when insurance benefits are threatened)? Total agreement on all such issues is a rarity in any active science. Thus, objectivity should always be an aspiration, but can never be an achievement of science. When public and experts disagree, it is a clash between two sets of differently informed opinions. Sciences, scientists, and definitions of risk differ greatly in how explicitly they acknowledge the role of judgment.

Dimensionality of risk. The risks of a technology are seldom its only consequences. No one would produce it if it did not generate some benefits for someone. No one could produce it without incurring some costs. The difference between these benefits and non-risk costs could be called its net benefit. In addition, risk itself is seldom just a single consequence. A technology may be capable of causing fatalities in several ways (e.g., by explosions and chronic toxicity), as well as inducing various forms of morbidity. It can affect plants and animals as well as humans. An analysis of "risk" needs to specify which of these dimensions will be included. In general, definitions based on a single dimension will favor technologies that do their harm in a variety of ways (as opposed to those that create a lot of one kind of problem). Although it represents particular values (and leads to decisions consonant with those values), the specification of dimensionality (like any other specification) is often the inadvertent product of convention or other forces, such as jurisdictional boundaries (Fischhoff, in press).

Summary statistic. For each dimension selected as relevant, some quantitative summary is needed for expressing how much of that kind of risk is created by a technology. The controversial aspects of that choice can be seen by comparing the practices of different scientists. For some, the unit of choice is the annual death toll (e.g., Zentner, 1979); for others, death per person exposed or per hour of exposure (e.g., Starr, 1969; Wilson, 1979); for others, it is the loss of life expectancy (e.g., Cohen and Lee, 1979; Reissland and Harries, 1979); for still others, lost working days (e.g., Inhaber, 1979). Crouch and Wilson (1982) have shown how the choice of unit can affect the relative riskiness of technologies; for example, today's coal mines are much less risky than those of thirty years ago in terms of accidental deaths per ton of coal, but marginally riskier in terms of accidental deaths per employee. The difference between measures is explained by increased productivity. The choice among measures is a policy question, with Crouch and Wilson suggesting that, "From a national point of view, given that a certain amount of coal has to be obtained, deaths per million tons of coal is the more appropriate measure of risk, whereas from a labor leader's point of view, deaths per thousand persons employed may be more relevant" (p. 13).

Other value questions may be seen in the units themselves. For example, loss of life

expectancy places a premium on early deaths which is absent from measures that treat all deaths equally; using it means ascribing particular worth to the lives of young people. Just counting fatalities expresses indifference to whether they come immediately after mishaps or following a substantial latency period (during which it may not be clear who will die). Whatever individuals are included in a category are treated as equals; these may include beneficiaries and non-beneficiaries of the technology (reflecting an attitude toward that kind of equity), workers and members of the general public (reflecting an attitude toward that kind of voluntariness), or participants and non-participants in setting policy for the technology (reflecting an attitude toward that kind of voluntariness). Using the average of past casualties or the expectation of future fatalities means ignoring the distribution of risk over time; it treats technologies taking a steady annual toll in the same way as those that are typically benign, except for the rare catastrophic accident. When averages are inadequate, a case might be made for using one of the higher moments of the distribution of casualties over time or for incorporating a measure of the uncertainty surrounding estimates (Fischhoff, in press).

Bounding the technology. Willingness to count delayed fatalities means that a technology's effects are not being bounded in time (as they are, for example, in some legal proceedings that consider the time that passes between cause, effect, discovery, and reporting). Other bounds need to be set also, either implicitly or explicitly. One is the proportion of the fuel and materials cycles to be considered: to what extent should the risks be restricted to those directly associated with the enjoyment of benefits or extended to the full range of activities necessary if those benefits are to be obtained? Crouch and Wilson (1982) offer an insightful discussion of some of these issues in the context of imported steel; the U.S. Nuclear Regulatory Commission (1983) has adopted a restrictive definition in setting safety goals for nuclear power (Fischhoff, 1983); much of the acrimony in the debates over the risks of competing energy technologies concerned treatment of the risks of back-up energy sources (Herbert et al., 1979). A second recurrent bounding problem is how far to go in considering higher-order consequences (i.e., when coping with one risk exposes people to another). A third is how to treat a technology's partial contribution to consequences, for example, when it renders people susceptible to other problems or when it accentuates other effects through synergistic processes.

Concern. Events that threaten people's health and safety exact a toll even if they never happen. Concern over accidents, illness, and unemployment occupy people even when they and their loved ones experience long, robust, and salaried lives. Although associated with risks, these consequences are virtual certainties. All those who know about them will respond to them in some way. In some cases, that response benefits the respondent, even if its source is an aversive event. For example, financial worries may prompt people to expand their personal skills or create socially useful innovations. Nonetheless, their resources have been diverted from other, perhaps preferred pursuits. Moreover, the accompanying stress can contribute to a variety of negative health effects, particularly

when it is hard to control the threat (Elliott and Eisdorfer, 1982). Stressors not only precipitate problems of their own, but can complicate other problems and divert the psychological resources needed to cope with them. Thus, concern about a risk may hasten the end of a marriage by giving the couple one more thing to fight about and that much less energy to look for solutions.

Hazardous technologies can evoke such concern even when they are functioning perfectly. Some of the response may be focussed and purposeful, such as attempts to reduce the risk through personal and collective action. However, even that effort should be considered as a cost of the technology because that time and energy might be invested in something else (e.g., leisure, financial planning, improving professional skills) were it not for the technology. When many people are exposed to the risk (or are concerned about the exposure of their fellows), then the costs may be very extensive. Concern may have even greater impact than the actual health and safety effects. Ironically, because the signs of stress are diffuse (e.g., a few more divorces, somewhat aggravated cardiovascular problems), it is quite possible for the size of the effects to be both intolerably large (considering the benefits) and unmeasurable (by current techniques).

Including concern among the consequences of a risky technology immediately raises two additional controversial issues. One is what constitutes an appropriate level of concern. It could be argued that concern should be proportionate to physical risk. There are, however, a variety of reasons why citizens might reasonably be concerned most about hazards that they themselves acknowledge to be relatively small (e.g., they feel that an important precedent is being set, that things will get worse if not checked, or that the chances for effective action are great). The second issue is whether to hold a technology responsible for the concern evoked by people's perceptions of its risks or for the level of concern that would be evoked were they to share the best available technical knowledge. It is the former that determines actual concern; however, using it would mean penalizing some technologies for evoking unjustified concerns and rewarding others for having escaped the public eye.

The Nature of Risky Decisions

Although a part of all risky decisions, risk is all of very few. Hazard management would be easy if risk were a substance and a technology could be characterized (and managed) effectively in terms of how much of that substance it contained (Watson, 1981). Risky decisions are, however, not about risk alone. Rather, they are choices among options, each of which has a variety of relevant features, including a level of risk. When a technology is adopted, so is its entire package of features. Thus, it is impossible to infer from its adoption that a technology has an acceptable level of risk (Fischhoff et al., 1981; Green, 1980; Otway and von Winterfeldt, 1982). Those adopting it might prefer much less risk, but be unable to obtain it at an acceptable price. In other decisions (or even in that decision should the possibilities change), they might adopt much less risky options.

From this perspective, the most general role for a definition of risk is to provide a

coherent, explicit, consistent expression of one subset of the consequences arising in risky decisions. For deliberative decisionmaking to proceed, it must be complemented by comparable conceptual analyses of the other consequences. With a clear set of concepts, it is possible to begin making the hard tradeoffs between risks and net benefits (which may include any positive value attributed to risk itself due, say, to the thrill or excitement it produces).

There are, however, some reasons for thinking about risks in isolation. One is educating the intuitions. The risks created by many technologies are so diverse that it is hard to think about them all at once. The rem and Sievert, which aggregate diverse radiation doses, attempt to serve this role. A second reason is to summarize the conclusions of policymaking that has considered other factors (Fischhoff, *in press*). Health and safety standards are often expressed in terms of an "acceptable level of risk," even though nonrisk costs and benefits strongly influenced how they were set (otherwise, they would be set at zero risk). That expression may enunciate a political philosophy ("we care about the public to this extent"), or it may provide an operational rule for the technical staff monitoring compliance, or it may be the only legitimate public conclusion of an agency that is mandated to manage risk (but must, in practice, consider risk-benefit tradeoffs). A third role is providing an explicit criterion for guiding and evaluating an agency's actions. A safety measure, such as a mandatory seat belt law, might have quite a different effect on "risk" if that is defined as deaths, serious injuries, or all injuries. Evaluating it fairly requires knowing what it was intended to accomplish.

Aspects of Risk

The first step in defining "risk" is determining which consequences it should include. Because that determination depends upon the particular problem, some context must be specified in order to produce even a hypothetical example. The context adopted here is evaluating the risks of competing energy technologies, as a component of setting national energy policy. Like any other choice of context, this one renders consequences that none of the competing options create – "unimportant for present purposes" – whatever their overall importance to society. This particular choice means that the selection of consequences (like other aspects of the definition process) should reflect "society's values," rather than those of any single interest. If one wished to revise or criticize this example, that effort, too, should begin with its selection of consequences.

Figure 1(B) shows that selection. Three kinds of risky consequence are included: fatalities, concern, and morbidity. Each is meant to include consequences whose magnitude is known, even though the identity of the casualties is not. For example, fatal accidents are a risk to those exposed to motor vehicles, even though the annual death toll is quite predictable. Each is meant to exclude anything but threats to human health and safety (e.g., accompanying property and financial risks). Actually choosing among energy technologies would require consideration of the broader set of consequences appearing in Fig. 1(A). The general form of Fig. 1(B) takes a position on one of the five




Fig. 1. Possible dimensions of consequence: (A) for decisionmaking, (B) for risk index.

sources of controversy regarding risk, its dimensionality. The specific contents take a position on a second, whether to consider concern a consequence of risk, despite its not having an obvious physical or physiological measure. Positions on the others are taken below in the course of developing a procedure for expressing the risks of the energy technologies.

Figure 1(B) makes two further distinctions. One is between mortal risks to the general public and mortal risks to workers in the technology. Such a distinction is made in many industries, with safety standards for workers being (very) roughly one-tenth as stringent as those for nonworkers (Derr et al., 1983). Making the distinction here allows one to decide whether this common practice should be accepted and enshrined in public policy by assigning a different weight to public and worker deaths. As discussed above, distinctions might also be made on the basis of whether those who die also benefit from the technology, have consented to exposure, or lose their lives in catastrophic accidents.

The second distinction is between two kinds of concern. Studies of lay risk perceptions (Fischhoff et al., 1978; Slovic et al., in press; Vlek and stallen, 1981) have shown that concern about technologies' risks can be predicted quite well by two "subjective dimensions of risk." These dimensions summarize a large number of individual determinants of perceived risk. They may be described as reflecting (a) the degree to which the risk is *unknown* and (b) the degree to which the risk evokes a feeling of *dread*. The former expresses aversion to uncertainty, and thus represents cognitive (or intellectual) aspects of concern, whereas the latter captures a risk's ability to evoke a visceral response. This usage takes these dimensions from the domain of prediction to the domain of prescription, from antici-

pating how people will respond to technologies to guiding how those technologies will be shaped.

The morbidity aspect is intended for all nonfatal injury and illness. It would also include genetic damage, whether expressed in birth defects or latent in the population. (The production of spontaneous abortions might appear here or in the previous category.) It should also include the unpleasantness of the period preceding death, which should be negligible for some accidents and considerable for some lingering illnesses.

Constructing Risk Indices

Choosing a set of attributes to describe a risk prospect creates a vector, each element of which expresses one dimension of consequence. The definition process could be terminated at that point, leaving users to integrate the elements intuitively. Alternatively, the elements can be combined into an aggregate measure of risk in order to eliminate the costs, errors, and vagueness that come with intuitive integration. The essence of aggregation is determining the relative importance of the elements. The following is a generalized scheme that can be adapted to the needs of many problems and value systems. It is drawn from multiattribute utility theory, fuller expositions of which can be found elsewhere, for both the theory itself (Keeney and Raiffa, 1976) and its application to risk problems (Ahmed and Husseiny, 1978; Keeney, 1980; Lathrop and Watson, 1982). The present application differs from its predecessors in focussing on risk, emphasizing the full range of options, and offering no opinion as to the correct solution. The logic of such procedures is as follows:

An aggregation procedure for risks should characterize a technology by a single value such that activities having a higher value will be more risky (in the eyes of those whose values the procedure represents). After the components have been selected, the next step is their operationalization. For example, if the consequences are called x_i , then x_1 might be the number of (additional) deaths among members of the general public; x_2 , the number of years of incapacitating illness caused; x_3 , a measure of public concern, etc. These consequences are often called "attributes"; we will use the terms interchangeably. This operationalization creates a vector of measurements $(x_1, \dots, x_n)$.

This vector expresses the measure, but not the worth, of those consequences. "Utility" is commonly used as a generalized unit of worth. Utility theory offers a wide variety of procedures for converting each such vector into a single number, representing its overall (un)desirability. These procedures can incorporate highly complex value systems and varying degrees of uncertainty regarding which consequences will, in fact, be experienced. In practice, though, various simplifying assumptions are adopted to render the analysis more tractable. For example, instead of explicitly modeling uncertainty (complete, say, with the elicitation of probability distributions for the values of different parameters), the analyst might treat all results as certainties. However, at the end of the analysis, each parameter will be varied through a range of plausible values to see whether such sensitivity analyses affect the previously reached conclusions.

In order to simplify the exposition, the present example makes two potentially controversial assumptions regarding the value structure. The first of these is *risk neutrality*, meaning, for example, that the certain loss of one life is just as bad as one chance in 10^5 of 10^5 deaths. Although it has been argued that people are particularly averse to losing many lives at once (as opposed to losing as many lives in separate accidents), this tendency appears to be due primarily to the great uncertainty that surrounds technologies capable of producing large accidents. A death is a death, whether it comes alone or with many others (Slovic et al., in press). In this definition, the uncertainty surrounding those deaths can be incorporated in the sensitivity analyses and the first “concern” attribute. The second simplifying assumption asserts that the underlying value structure is not overly complex. It is formalized as the property of *mutual preference independence*. Roughly speaking, two attributes are preference independent of all others if tradeoffs between them do not depend upon the levels of the other attributes.

If either of these assumptions seems wanting, then it is straightforward (if cumbersome) to repeat the analysis with alternative assumptions. If they seem adequate, then it is possible to express the index as

$$R = \sum_{j=1}^n w_j y_j \quad (1)$$

where y_j is the expected utility for attribute j and w_j is a weighting factor, expressing its relative importance. “Expected utility” is the product of a consequence’s utility and the probability of it being incurred if a technology is pursued. For example, if x_i were the number of public deaths, then y_i would be the expected utility for public deaths, which would consider not only the probability for different losses, but also any changes in the significance of marginal deaths as a function of total deaths.

Risks of Electricity Generation: Different Definitions

Problem Description

Electricity generation is an interesting case for two reasons. One is the evidence that disagreement about the definition of key terms (including “risk”) has contributed to the bitterness of many energy debates. The second is that important issues tend to generate research, producing data upon which risk estimates may be more soundly based.

In this analysis, six energy technologies are considered. Five of these, coal, hydropower, large-scale windpower, small-scale windpower, and nuclear power, can increase the supply of electricity. The sixth, energy conservation, can reduce the demand for electricity, thereby freeing existing supplies for use elsewhere.

Attribute Definition

The five attributes of these technologies are those shown in Fig. 1(B). They are opera-

tionalized as follows: Both kinds of death are measured in terms of the expected number of deaths per Gigawatt year (GWyr) of electricity generated or saved. Choosing this summary statistic means taking positions on two additional dimensions of controversy regarding the definition of risk: Broad bounds are set on the technologies, so as to attribute to them all casualties incurred in conjunction with generating electricity. Deaths are just tallied, without regard for the number of years taken off each, the extent of each victim's exposure, the distribution of deaths over time, or any of the other features discussed earlier. Morbidity will be measured by expected person-days of incapacity per GWyr of electricity.

The two attributes associated with concern will be specified in terms of the technologies' ratings on the two comparable factors in psychometric studies of perceived risk (e.g., Slovic et al., in press). These studies have produced sufficiently robust results to make reliance on them conceivable; perceptions of risk have proven sufficiently good predictors of attitudes and actions for them to serve as reasonable indicators of level of concern. What is most arguable about such reliance is treating the expression of concern as evidence of adverse consequences. As discussed above, one ground for that claim is that concern itself is an adverse consequence, which should not be imposed upon people without compensating benefit; a second ground is that concern is associated with stress which is, in turn, associated with various physiological effects that are so difficult to measure that it is reasonable to use concern as a surrogate for them.

Evaluating Consequences

Having defined the attributes, the next step is to evaluate each possible outcome on each (e.g., how bad is it to incur 10 or 100 worker deaths). In technical terms, this means defining a utility function for each attribute. A convenient way of doing so, given the assumptions made here, is to use a 100-point scale for each attribute, where 0 represents the least extreme possible consequence and 100 the most extreme possible consequence. (If both good and bad consequences were being considered, then a distinction between positive and negative scores would be necessary. Here, 100 is the worst possible outcome.) Intermediate values are defined appropriately. Although linear scaling is possible, it is not necessary. For example, for most people winning \$100 will not be 10 times as satisfying as winning \$10. Setting the end points of each scale requires a factual (or scientific) judgment regarding what consequences are possible. Setting the midpoints requires a value judgment regarding how those intermediate consequences are regarded.

A natural zero point for a casualty scale is zero casualties. It will be used here, recognizing that no deaths to workers, no deaths to the public, and no person-days lost are practically unachievable with any energy technology. On the basis of worst-case analyses, scores of 100 on attributes 1 and 2 are defined, respectively, as 10 public deaths and 10 occupational deaths per GWyr of electricity generated or saved. Similarly, 60,000 person-days of incapacity per GWyr would merit a score of 100 on attribute 3. Intermediate scores are assigned linearly (e.g., on attributes 1 and 2, one death receives 10, two

TABLE 1
The Components of Attributes 4 and 5

Attribute	Score of 0 implies risk has these properties	Score of 100 implies risk has these properties
4. Unknown risk	Observable Known to exposed Effect immediate Old Known to science	Not observable Unknown to exposed Effect delayed New Unknown to science
5. Dread risk	Controllable Not dread Not global catastrophic Consequences not fatal Equitable Individual Low future risk Easily reduced Decreasing Voluntary Doesn't affect me	Uncontrollable Dread Global catastrophic Consequences fatal Not equitable Catastrophic High future risk Not easily reduced Increasing Involuntary Affects me

Source: Slovic, Fischhoff and Lichtenstein, 1984.

deaths receive 20, etc.), reflecting a desire to assign an equal value to each casualty (as distinct, perhaps, from the decreasing sensitivity to additional casualties that people might actually experience).

Table 1 shows the characteristics that would give a technology scores of 0 and 100 on attributes 4 (unknown risk) and 5 (dread risk). Research has shown that although no technology quite reaches either extreme, mountain climbing and handguns score close to zero on attribute 4 (at least in the U.S.A., as do home appliances and high school football on attribute 5. At the other extreme, DNA research is rated as sufficiently unknown to receive a score in the 90s on attribute 4, while nuclear weapons do likewise on attribute 5.

Making Tradeoffs

The final step in specifying an evaluation scheme is to assign weights reflecting the relative importance of the different attributes. As these weights reflect value judgments, disagreements are legitimate; in the present context, they are to be expected. Table 2 presents four sets of weights, each reflecting a different set of values.

Brief descriptions might help explicate the perspectives that could motivate each set's adoption. The first rejects anything but readily measured physiological effects; treats a death as a death, whether it befalls a worker or a member of the public; views a life as equal to 6000 person-days of incapacity. Set B reflects a belief that concern is a legitimate consequence, that public deaths are twice as important as worker deaths, and that a

TABLE 2
Four Possible Sets of Weights for Five Risk Attributes

Attributes	A	B	C	D
1. Public deaths	0.33	0.40	0.20	0.08
2. Occupational deaths	0.33	0.20	0.05	0.04
3. Morbidity	0.33	0.20	0.05	0.40
4. Unknown risk	0	0.10	0.30	0.24
5. Dread risk	0	0.10	0.40	0.24
Sum of weights	1	1	1	1

worker death should be treated as equivalent to the loss of 6000 person-days. As Dunster (1980) argues, "it is not easy to weigh the benefits of reducing anxiety against those of saving life, but our society certainly does not require the saving of life to be given complete priority over the reduction of anxiety" (p. 127). Set C increases the importance ratio for public to occupational deaths and assigns major significance to concern. The specific weights imply a willingness to tradeoff 10 public deaths per GWyr to move from a technology causing extreme dread to one that is about average, perhaps feeling that the toll from concern-generated stress is large or that even minor accidents in a dread technology can cause enormously costly social disruption. The D weights represent a paramount concern with the suffering of the living, whether through injury or anxiety, rather than with the number of deaths.

Whatever one's value system, the weights assigned should be very sensitive to the range of outcomes considered on each attribute. If, for example, 100 on attribute 1 meant 50 public deaths per year (rather than 10), then Set A would have to assign a larger value to attribute 1 to achieve the same effect of weighting a public and a worker death equally.

Scoring Technologies

In order to apply this scheme to technologies, it is necessary to assess the magnitude of the consequences that each produces on each attribute. This is a scientific, not a value question. It should be informed by the best available technical knowledge. However, applying that knowledge in the present case requires the exercise of judgment, to choose, weigh, and extrapolate from existing studies. Despite having a commitment to objectivity, we cannot escape some subjectivity in attempting to derive this sort of policy-oriented advice.

Table 3 provides point estimates roughly summarizing the research reported in the following sources: Baecher et al. (1980), Birkhofer (1980), Bliss et al. (1979), Budnitz and Holdren (1976), Comar and Sagan (1976), Department of Energy (1979), Dunster (1980), Greenhalgh (1980), Hamilton (1980), Okrent (1980), Rogers and Templin (1980), and Slovic et al. (1980, in press). This literature reveals both substantial differences of opinion and substantial areas of ignorance. As two examples: The extreme values for

TABLE 3

The Scores of Six Technologies of Five Risk Attributes by One Expert

Attribute	Coal	Hydro	Large scale wind	Small scale wind	Nuclear	Conser- vation
1. Public deaths	80	10	20	5	10	5
2. Occupational deaths	30	20	10	30	5	10
3. Morbidity	20	20	40	50	10	40
4. Unknown risk	70	60	90	50	80	40
5. Dread risk	50	50	40	20	90	10

expected occupational deaths from coal were 0.7 and 8 deaths per GWyr of electricity generated. Very few risk data were available for either small-scale wind power or conservation; these scores were liberally adapted from knowledge of other technologies. Where available, the concern scores required the least exercise of judgment. Technologies that have been rated have proven to have rather robust scores on these dimensions, regardless of who does the rating, how the rating is carried out, and what other technologies are in the rating set (Slovic et al., in press). However, several energy technologies have yet to be evaluated in this way. Their scores were derived by conjecture. For example, the scores for conservation on attributes 4 and 5 were averages of those for home appliances and bicycles.

Given the unreliability of these estimates, any attempt to establish the risks of energy technologies would have to address the uncertainty surrounding them, with either sensitivity analyses or explicit assessment of probabilities. Given the illustrative nature of the present example, that exercise will be foregone as misplaced imprecision.

Computing Risk

Using these values and Set A's weights, Expression (1) shows the risk from coal to be $0.33(80) + 0.33(30) + 0.33(20) = 42.9$. Other scores are computed similarly and displayed in Fig. 2. Because the scores are standardized to range from 0 to 100 and the weights to sum to 1.0, it is possible to compare scores across technologies and across weighting schemes. That comparison shows that the riskiness of coal, small-scale windpower, and conservation vary little across these four sets of weights, whilst those for hydro, large-scale windpower, and particularly nuclear power vary greatly. Thus, if one accepts the consequence estimates of Table 3, then the riskiness of these last three technologies depends upon the importance assigned to the different consequences.

Table 4 shows how this sensitivity expresses itself in terms of the relative riskiness of the six technologies. Coal, for example, ranks consistently low, whereas nuclear may be best or worst depending upon the definition used. These enormous variations occur despite complete agreement regarding the magnitude of the consequences. Thus, arguments over relative risk may reflect only disagreements about values.




Fig. 2. The risk indices of six technologies on four sets of weights.

Conclusion

An effective decisionmaking process, whether conducted by individuals or societies, requires agreement on basic terms. Without such conceptual clarity, miscommunication and confusion are likely. Definitional ambiguity regarding the term “risk,” in particular,

TABLE 4

The Risk-To-Human-Health Rankings of Six Technologies Given Four Sets of Weights on One Set of Five Attributes and the Scores of One Expert

Rank		Set of Weights			
		A	B	C	D
Best	1.	Nuclear	Conservation	Conservation	Conservation
	2.	Hydro	Hydro	Small wind	Hydro
	3.	Conservation	Nuclear	Hydro	Small wind
	4.	Large wind	Small wind	Large wind	Coal
	5.	Small wind	Large wind	Coal	Nuclear
Worst	6.	Coal	Coal	Nuclear	Large wind

has spawned needless (and irresolvable) conflict over the relative riskiness of different technologies. At the same time, it has obscured the need for debate over the value issues involved in specifying what "risk" means.

The present analysis presents a framework for defining risk which directly faces those inherent conflicts. Indeed, it forces one to adopt an explicit position on each aspect of the controversy before a workable definition can be created. As a result, the specific indices of risk developed here are controversial by design. However, they are also expendable by design. The general framework is highly flexible, capable of fitting many problems and many value systems. Its use in a particular problem makes possible a diagnosis of the extent to which conflicts reflect disagreements about facts or disagreements about values. In the former case, one can hope that consensus about risk will evolve as scientific research progresses. In the latter case, consensus will only emerge if there is effective public debate about what society should value. That debate can be informed (and spurred) by ethical and policy analyses, but it cannot be resolved by them.

Applying this potentially rich procedure required a series of simplifying assumptions. These included taking only five risky consequences from the vector of possibilities, asserting risk neutrality and mutual preference independence, and representing effect magnitude by point estimates. Despite these restrictions, this illustrative analysis showed that the relative riskiness of different energy technologies is quite sensitive to how risk is defined.

The emphasis here has been on the logic of the analysis, rather than on its content. Making a definitive statement regarding the risks of competing energy technologies would require definitive estimates of both the magnitude and the importance of those consequences for a particular society. Neither was attempted here. An additional caution is that even a definitive statement about risk would have no necessary implications for most policymaking. People do not accept risks, but technologies, one of whose significant features may be their risks. Developing an index of risk allows systematic treatment of one aspect of those decisions, but only one aspect. Analogous treatments of other consequence domains would be needed to complete the picture.

Developing a definition of risk requires a variety of explicit value judgments. Choosing to express risk in a numerical index may itself make a statement of values. The present exposition emphasized the possibilities that an index offers for including different people's values in policymaking. However, it may also be used to exclude the people themselves from the policymaking process, with policy experts serving as self-appointed spokespeople for what the public wants. Even if careful research is conducted to identify the public values that are to be incorporated in society's risk index, such technical recognition need not substitute for active, personal participation. If it is used that way, then the index may be blamed for the faults of a political process that can tolerate public opinion, but not the public.

Acknowledgments

This research was supported by the National Science Foundation under Grant PRA-8116925 to Perceptronics, Inc. Any opinions, findings and conclusions expressed in this article are those of the authors and do not necessarily reflect the views of the National Science Foundation. We give our thanks to Vincent Covello, Jack Dowie, Ken Hammond, Paul Slovic, and Ola Svenson for perceptive comments on previous drafts.

Notes

Chris Hope is now at the Department of Fuel Sciences, University of Leeds, Leeds, Great Britain. Requests for reprints should be sent to Baruch Fischhoff, 1201 Oak Street, Eugene, Oregon 97401, U.S.A.

References

- Ahmed, S. and Husseiny, A. A. (1978). "A multivariate utility approach for selection of energy sources," *Energy* 3: 669-700.
- Baecher, G. B., Pate, M. E. and de Neufville, R. (1980). "Risk of dam failure in benefit-cost analysis," *Water Resources Research* 16: 449-456.
- Birkhofer, A. (1980). "The German risk study for nuclear power plants," *IAEA Bulletin* 22 (5/6): 23-33.
- Bliss, C., Clifford, P., Goldgraben, G., Graf-Webster, E., Krickenberg, K., Maher, H. and Zimmerman, N. (1979). *Accidents and Unscheduled Events Associated with Non-Nuclear Energy Resources and Technology*. Washington, D.C.: MITRE Corporation for Environmental Protection Agency.
- Budnitz, R. J. and Holdren, J. P. (1976). "Social and environmental costs of energy systems," *Annual Review of Energy* 1: 553-580.
- Cohen, B. and Lee, I. (1979). "A catalog of risks," *Health Physics* 36: 707-722.
- Comar, C. L. and Sagan, L. A. (1976). "Health effects of energy production and conversion," *Annual Review of Energy* 1: 581-660.
- Cotgrove, A. (1982). *Catastrophe or Cornucopia? The Environment, Politics and the Future*. New York: Wiley.
- Crouch, E. A. C. and Wilson, R. (1982). *Risk/Benefit Analysis*. Cambridge, Mass.: Ballinger.
- Department of Energy (1979). *Energy Technologies for the United Kingdom, Vol. I and II*. London: HMSO.
- Derr, P., Goble, R., Kasperson, R. E. and Kates, R. W. (1983). "Responding to the double standard of worker/public protection," *Environment* 25 (6): 6-11, 35-36.
- Dunster, J. H. (1980). "The approach of a regulatory authority to the concept of risk," *IAEA Bulletin* 22 (5/6): 123-128.
- Elliott, G. R. and Eisdorfer, C., (Eds.) (1982). *Stress and Human Health*. New York: Springer Verlag.
- Fischhoff, B. (1983). "Acceptable risk: The case of nuclear power," *Journal of Policy Analysis and Management* 2: 559-575.
- Fischhoff, B. (1984). "Standard setting standards," *Management Science* in press.
- Fischhoff, B., Lichtenstein, S., Slovic, P., Derby, S. L. and Keeney, R. L. (1981). *Acceptable Risk*. New York: Cambridge University Press.
- Fischhoff, B., Slovic, P. and Lichtenstein, S. (1983). "The public vs. The experts: Perceived vs. actual disagreements about the risks of nuclear power," in V. Covello, G. Flamm, J. Rodericks, and R. Tardiff (Eds.), *Analysis of Actual vs. Perceived Risks*. New York: Plenum.
- Fischhoff, B., Slovic, P., Lichtenstein, S., Read, S. and Combs, B. (1978). "How safe is safe enough? A psychometric study of attitudes towards technological risks and benefits," *Policy Sciences* 8: 127-152.
- Green, C. H. (1980). "Risk: Attitudes and beliefs," in D. V. Canter (Ed.), *Behaviour in Fires*. Chichester: Wiley.
- Greenhalgh, G. (1980). *The Necessity for Nuclear Power*. London: Graham & Trotman.
- Hamilton, L. D. (1980). "Comparative risks from different energy systems: Evolutions of the methods of studies," *IAEA Bulletin* 22 (5/6): 35-71.

- Herbert, J. H., Swanson, L. and Reddy, P. (1979). "A risky business," *Environment* 21 (6): 28-33.
- Inhaber, H. (1979). "Risk with energy from conventional and nonconventional sources," *Science* 203: 718-723.
- Keeney, R. L. (1980). "Evaluating alternatives involving potential fatalities," *Operations Research* 28: 188-205.
- Keeney, R. L. and Raiffa, H. (1976). *Decisions with Multiple Objectives. Preferences and Value Trade-offs*. New York: Wiley.
- Lathrop, J. W. and Watson, S. R. (1982). "Decision analysis for the evaluation of risk in nuclear waste management," *Journal of the Operational Research Society* 3: 407-418.
- Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M. and Combs, B. (1978). "Judged frequency of lethal events," *Journal of Experimental Psychology: Human Learning and Memory* 4: 551-578.
- Okrent, D. (1980). "Comment on social risk," *Science* 208: 372-375.
- Otway, H. and von Winterfeldt, D. (1982). "Beyond acceptable risk: On the social acceptability of technologies," *Policy Sciences* 14: 247-256.
- Reissland, J. and Harries, V. (1979). "A scale for measuring risks," *New Scientist* 83: 809-811.
- Rogers, D. W. O. and Templin, R. J. (1980). "Errors in a risk assessment of renewable resources," *Energy* 5: 101-103.
- Slovic, P., Fischhoff, B. and Lichtenstein, S. (1980). "Facts and fears: Understanding perceived risk," in R. Schwing and W. A. Albers Jr. (Eds.), *Societal Risk Assessment: How Safe is Safe Enough?* New York: Plenum.
- Slovic, P., Fischhoff, B. and Lichtenstein, S. (1984). "Characterizing perceived risk," in R. W. Kates and C. Hohenemser (Eds.), *Technological Hazard Management*. Cambridge, Mass.: Oelgeschlager, Gunn & Hain, in press.
- Slovic, P., Lichtenstein, L. and Fischhoff, B. (1984). "Modeling the societal impact of fatal accidents," *Management Science* in press.
- Starr, C. (1969). "Social benefit versus technological risk," *Science* 165: 1232-1238.
- U.S. Nuclear Regulatory Commission (1982). *Safety Goals for Nuclear Power Plants: A Discussion Paper*. NUREG-0880. Washington, D.C.: The Commission.
- Vlek, C. A. J. and Stallen, P. J. M. (1981). "Risk perception in the small and in the large," *Organizational Behavior and Human Performance* 28: 235-271.
- Watson, S. R. (1981). "On risks and acceptability," *Journal of the Radiological Protection Society* 1: 21-25.
- Wilson, R. (1979). "Analyzing the daily risks of life," *Technology Review* 81 (4): 40-46.
- Wynne, B. (1983). "Institutional mythologies and dual societies in the management of risk," in H. C. Kunreuther and E. V. Ley (Eds.), *The Risk Analysis Controversy*. New York: Springer Verlag.
- Zentner, R. D. (1979). "Hazards in the chemical industry," *Chemical and Engineering News* 57 (45): 25-27, 30-34.