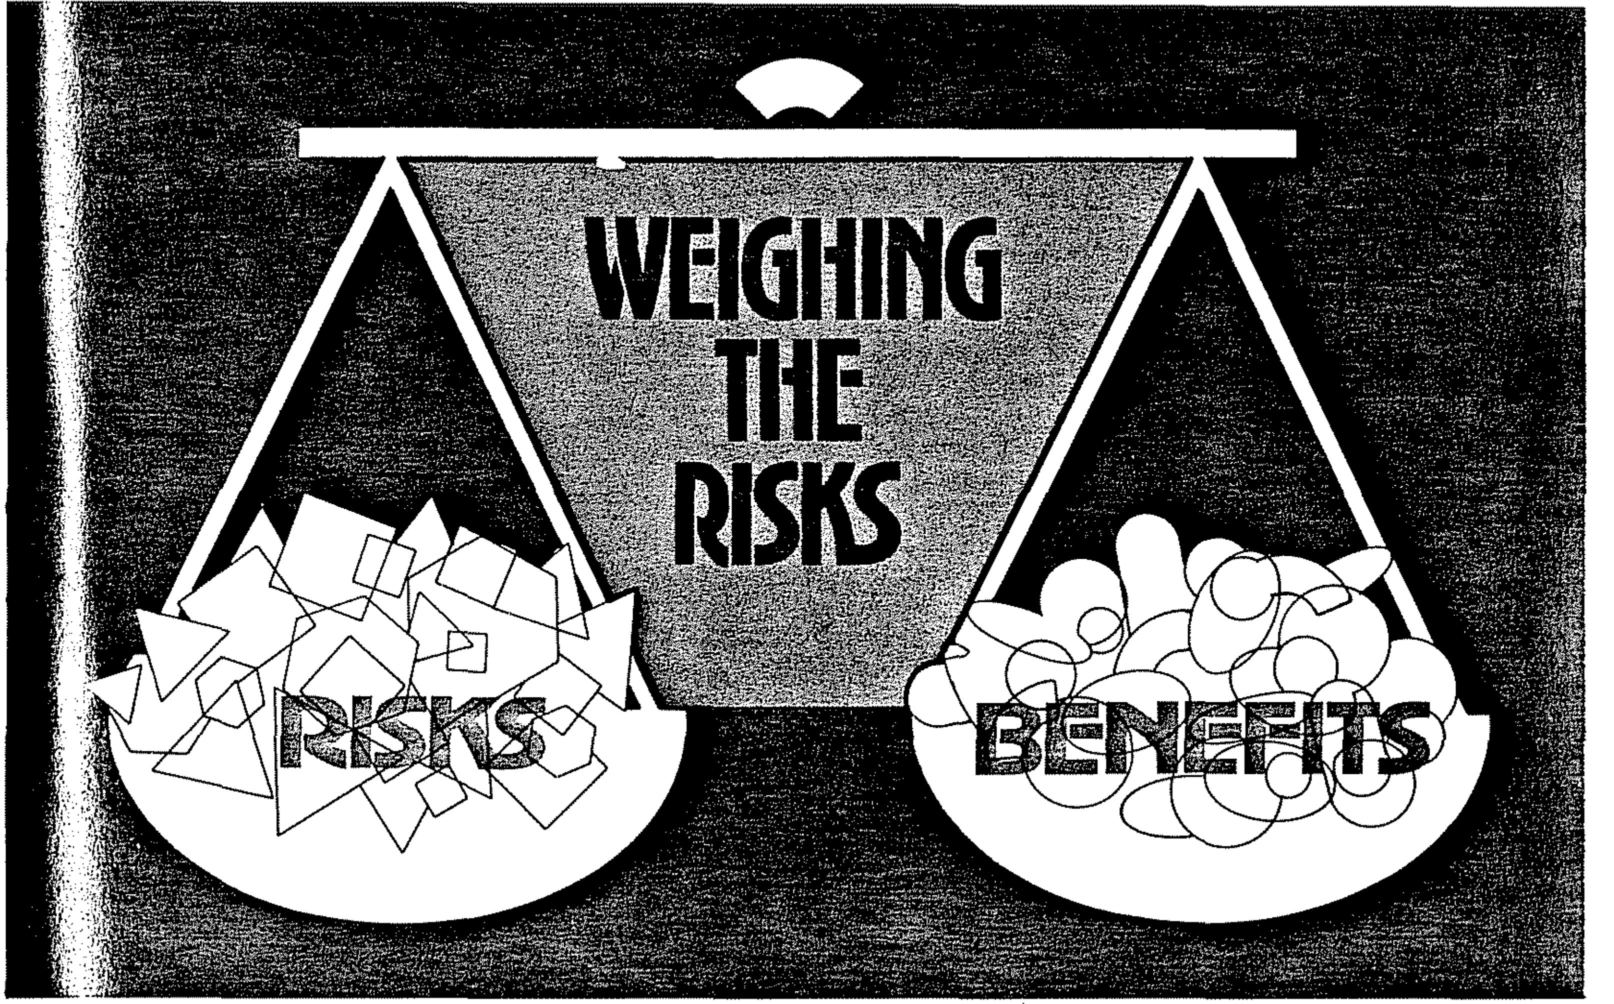

BY BARUCH FISCHHOFF,
PAUL SLOVIC, and
SARAH LICHTENSTEIN

Which Risks Are Acceptable?

BARUCH FISCHHOFF, PAUL SLOVIC, and SARAH LICHTENSTEIN are psychologists specializing in problems of judgment, decision making, and societal risk-taking behavior. They are research associates at Decision Research, a Branch of Perceptronics, Inc., in Eugene, Oregon.

THE BOTTOM LINE IN HAZARD MANAGEMENT is usually some variant of the question, "How safe is safe enough?" It takes such forms as: "Do we need additional containment shells around our nuclear power plants?" "Is the carcinogenicity of saccharin sufficiently low to allow its use?" "Should schools with asbestos ceilings be closed?" Lack of adequate answers to such questions has bedeviled hazard management.

Of late, many hazard management decisions are simply not being made—in part because of vague legislative mandates and cumbersome legal proceedings, in part because there are no clear criteria on the basis of which to decide. As a result, the nuclear industry has ground to a halt while utilities wait to see if the building of new plants will ever be feasible,[1] the Consumer Product Safety Commission has invested millions of dollars in producing a few puny standards,[2] observers wonder whether the new Toxic Substances Control Act can be implemented,[3] and the Food and Drug Administration is unable to resolve the competing claims that it is taking undue risks and that it is stifling innovation.

The decisions that are made are often inconsistent. Our legal statutes are less tolerant of carcinogens in the food we eat than of those in the water we drink or in the air we breathe. In the United Kingdom, 2,500 times as much money per life saved is spent on safety measures in the pharmaceutical industry as in agriculture.[4] U.S. society is apparently willing to spend about $140,000 in highway construction to save one life and $5 million to save a person from death due to radiation exposure.[5]

Frustration over this state of affairs has led to a search for clear, implementable rules which will tell us whether or not a given technology is sufficiently safe. Four approaches are most frequently used in attempting to make this assessment. They are cost-benefit analysis, revealed preferences, expressed preferences, and natural standards. Respectively, they would deem a technology to be safe if its benefits outweigh its cost; if its risks are no greater than those of currently tolerated technologies of equivalent benefit; if people say that its risks are acceptable; if its risks are no greater than those accompanying the development of the human species. Each of these approaches has its pros and cons, its uses and its limitations.[6]

Cost-Benefit Analysis

Cost-benefit analysis attempts to answer the question of whether the expected benefits from a proposed activity outweigh its expected costs. The first steps in calculating the expected cost of a project are: to enumerate all the adverse consequences that might result from its implementation; to assess the probability of each such consequence; and to estimate the cost or loss to society whenever the consequence occurs. Next, the expected cost of each possible consequence is calculated by multiplying the cost of the consequence by the probability that it will be incurred. The expected cost of the entire project is computed by summing the expected losses associated with the various possible consequences. An analogous procedure produces an estimate of the expected benefits (see box).[7] The most general form of cost-benefit analysis is *decision analysis*, in which the role of uncertainty, the subjective nature of costs and benefits, and the existence of alternative actions are made explicit.[8]

COST-BENEFIT ANALYSIS

Consider a fictitious new product, Veg-E-Wax, designed to coat fresh fruits and vegetables. Its demonstrated advantages are reducing losses in storage and preserving nutritive value. Aside from the cost of application, its disadvantages are making food look less appetizing and possibly causing cancer to workers who apply it and to consumers who fail to wash fruit. A highly simplified cost-benefit analysis of the decision to apply Veg-E-Wax to a $10 million (market value) shipment of pears bound for storage might appear as follows:

Advantages (benefits)	*$ million*
Guaranteed reduction in storage loss from 30% to 20%	1.0
Improved nutritive value (translating into a 10% increase in market value in the 80% that is not lost in storage).	.8
Total benefits	1.8
Disadvantages (costs)	
Cost of application	.1
Cancer in .1% of 100 workers (@$1 million per case)	.1
Cancer to users (1 million consumers, of whom 10% fail to wash fruit, of whom .0001/ contract cancer as a result, @$1 million per case).	.1
Unappetizing appearance (20% loss in market value of pears not lost in storage).	1.6
Total costs	1.9

In this calculation, the costs slightly outweigh the benefits and the packer should decide not to use Veg-E-Wax. The viability of this conclusion depends upon its capacity to withstand small changes in the figures. If there were only an 18% loss in market value due to the waxy look of the fruit (translating into a cost of $1.44 million), the balance would tip the other way. It might be impossible to predict this loss with the precision needed to take confident action.

Even larger effects may accompany changes in fundamental assumptions. A packer with no social conscience might decide not to worry about the $200,000 in cancer costs, reducing total costs to $1.7 million. Other interested parties, such as consumers interested in maximizing value and minimizing personal risk, might structure the problem entirely differently.

These procedures, and decision analysis in particular, are based on appealing premises and are supported by sophisticated methodology. Furthermore, they permit considerable flexibility; analyses are readily revised to incorporate new options and new information. An important advantage of these methods for decision making in the public sphere is that they are easily scrutinized. Each quantitative input or qualitative assumption is available for all to see and evaluate, as are the explicit computational rules that combine them.

However, decision analysis and its variants have a number of potentially serious limitations, perhaps the most important of which is their unrealistic assumptions about the availability of the data needed to complete the analysis. Performing a full-dress analysis assumes, among other things, that all possible events and all significant consequences can be enumerated in advance; that meaningful probability, cost, and benefit values can be obtained and assigned to them; and that the often disparate costs and benefits can somehow be made comparable to one another.

Unfortunately, it is sometimes impossible to accomplish some of these tasks, while in the case of others, the results are hardly to be trusted. Despite the enormous scientific progress of the last decade or two, we still do not know all or even most of the possible physical, biological, and social consequences of any large-scale energy project.[9] Even when we know what the consequences are, we often do not, or cannot, know their likelihood. For example, although we know that a nuclear reactor core melt-down is unlikely, we will not know quite how unlikely until we accumulate much more on-line experience. Even then, we will be able to utilize that knowledge only if we can assume that the reactor and the attendant circumstances remain the same (e.g., no changes in the incidence of terrorism or the availability of trained personnel). For many situations, even when a danger is known to be present, its extent cannot be known. Whenever low-level radiation or exposure to toxic substances is involved, consequences can be assessed only by tenuous extrapolation from the consequences of high-level exposure to human beings or from observation of exposure in animals.[10]

In all these instances, we must rely upon human judgment to guide or supplement our formal methods. Research into the psychological processes involved in producing such judgments offers reason for concern, since this research demonstrates that people (including experts forced to go beyond the available data and rely on their intuitions) have a great deal of difficulty both in comprehending complex and uncertain information and in making valid inferences from such information.[11] Frequently these problems can be traced to the use of judgmental heuristics—mental strategies whereby people try to reduce difficult tasks to simpler judgments. These strategies may be useful in some situations but in others they lead to errors that are large, persistent, and serious in their implications. Furthermore, individuals are typically unaware of these deficiencies in the judgments.

Even if all the consequences could be enumerated and their likelihood assessed, placing a price tag on them poses further difficulties. Consider, for example, the problems of placing a value on a human life. Despite our resistance to thinking about life in economic terms, the fact is that, by our actions, we actually do put a finite value on our lives. Decisions about installing safety features, buying life insurance, or accepting a more hazardous job for extra salary all carry implicit judgments about the value we place on a life.

Economists have long debated the question of how best to quantify the value of a life.[12] The traditional economic approach has been to equate the value of a life with the value of a person's expected future earnings. Many problems with this index are readily apparent. For one, it undervalues those in society who are underpaid and places no value at all on people who are not in income-earning positions. In addition, it ignores the interpersonal effects of a death which may make the loss suffered much greater than any measurable financial loss. A second approach, which equates the value of life with court awards, can hardly be considered to be more satisfactory.[13]

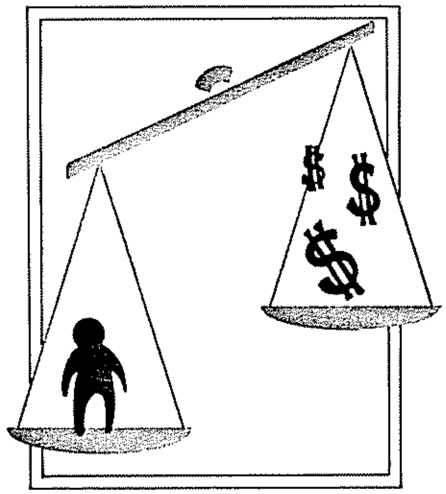

HOW MUCH IS A LIFE WORTH?

Decisions about installing safety devices, buying life insurance, or accepting a more hazardous job for extra salary are all based on implicit judgments about the value one places on a life.

Some have argued that the question, "What is a life worth?" is poorly phrased and what we really want to know is, "What is the value placed on a particular change in survival probability?"[14] One approach to answering this second question is by observing the actual market behavior of people trading risks for economic benefits. For example, one study examined salary as a function of occupational risk and claimed to find that a premium of about $200 per year was required to induce workers in risky occupations (coal mining, for example) to accept an increase of .001 in their annual probability of accidental death.[15] From this finding it was inferred that society should be willing to pay about $200,000 to prevent a death. A replication of this study by Rappoport[16] produced a value of $2,000,000; thus, even if one accepts the assumptions underlying this approach, a definitive value may still elude us.[17]

Decision analysis attempts to accommodate the uncertainties inherent in the assessment of problems and of the values of the variables involved through the judicious use of *sensitivity analysis*. The calculations of expected costs and benefits are repeated using alternative values of one troublesome probability, cost, or benefit. If each reanalysis produces the same relative preponderance of expected costs or benefits, then it is argued that these particular differences do not matter. In the Veg-E-Wax example (see box), changing the estimate of lost market value from 20 percent to 18 percent is a sensitivity analysis. The fact that it tipped the balance from rejection to acceptance of Veg-E-Wax suggests that neither recommendation can be strongly supported.

Unfortunately, however, there are no firm guidelines regarding which of the data might be in error or what range of possible values ought to be tested. A further problem with sensitivity analysis is that it typically tells us little about how the uncertainty from different sources of error is compounded or about what happens when different data are subject to a common bias. The untested assumption is that errors in different inputs will cancel one another, rather than compound in some pernicious way.[18]

In the end, determining the quality of an analysis is a matter of judgment. Someone must use intuition to determine which inputs are of doubtful validity and which alternative values should be incorporated in sensitivity analyses. Essentially, that someone must decide how good her or his own best judgment is. Unfortunately, an extensive body of research suggests that people tend to overestimate the quality of such judgments.[19]

Revealed Preferences

An alternate approach to determining acceptable risks is the method of revealed preferences advocated by Chauncey Starr.[20] This approach is based on the assumption that, by trial and error, society has arrived at an "essentially optimum" balance between the risks and benefits associated with any activity. As a result, it is assumed that economic risk and benefit data from recent years will reveal patterns of acceptable risk-benefit tradeoffs. Acceptable risk for a new technology is defined as that level of safety associated with ongoing activities having similar benefit to society. Starr argued the potential usefulness of revealed preferences by examining the relationship between risk and benefit across a number of common activities.

From this analysis, Starr derived what might be called "laws of acceptable risk":

- ✓ The acceptability of risk is roughly proportional to the third power (cube) of the benefits.
- ✓ The public seems willing to accept risks from voluntary activities, such as skiing, roughly a thousand times greater than it would tolerate from involuntary activities, such as food preservatives, that provide the same level of benefit.
- ✓ The acceptable level of risk is inversely related to the number of persons exposed to that risk.

Figure 1 depicts the results of Starr's analysis, while Figure 2 shows our own expanded replication of Starr's study, in which we examine 25 activities and technologies, including the 8 he used. In this replication somewhat different methods have been used. Whereas Starr estimated risk in terms of fatality rate per hour of exposure, we have used annual fatalities. This change is motivated in part by the greater availability of data for the latter measure and in part because the definition of exposure to some hazards (for instance, handguns, smoking, antibiotics) is elusive. Whereas Starr measured benefit either by the average amount of money spent on an activity by a single participant

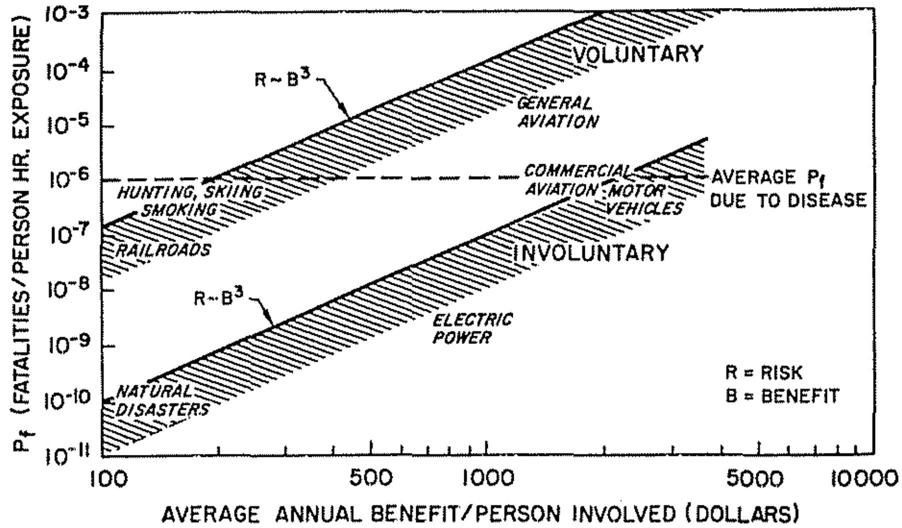

FIGURE 1. A comparison of risk and benefit to U.S. society from various sources. Risk is measured by fatalities per person per hour of exposure. Benefit reflects either the average amount of money spent on an activity by an individual participant or the average contribution an activity makes to a participant's annual income. The best-fitting lines were drawn by eye with error bands to indicate their approximate nature. *Source:* C. Starr. "Benefit-cost studies in sociotechnical systems," in Committee on Public Engineering Policy, *Perspective on Benefit-Risk-Decision Making,* National Academy of Engineering, Washington, D. C., 1972.

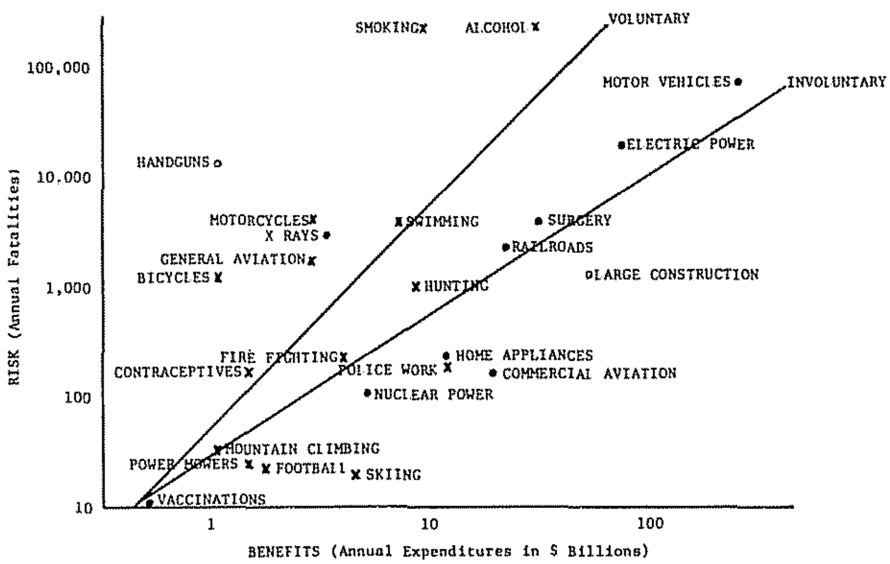

FIGURE 2. One possible assessment of current risks and benefits from 25 activities and technologies. Items are marked with an *X*, if voluntary; with a closed circle, if involuntary. Handguns and large construction cannot be classified as primarily voluntary or involuntary. They are marked here with open circles and are not included in the calculation of the two regression lines shown in the figure.

(continued on p. 32)

Weighing the Risks

(continued from page 20)

or the average contribution the activity made to a participant's annual income, we have used the single measure of total annual consumer expenditure.

Like any other economic measure of benefit, expenditure has its limitations. It includes "bad" as well as "good" expenditures; for example, money spent on the abatement of pollution caused by an industry is weighted as heavily as the value of the product it manufactures. A second problem is that this measure ignores distributive considerations (who pays and who profits). A third problem is that the market price may not be responsive to welfare issues that are critical to social planning. Does the price of cigarettes take account of smokers' higher probability of heart disease or cancer? Does the price of pesticides adequately reflect the increased probability of various deleterious effects on the one hand and the increased yield of foodstuffs on the other?

Expenditures for private goods (whose purchase is the result of the decisions of individual consumers) were obtained from trade and manufacturing associations, while public services, such as police work or fire fighting, were estimated by using government expenditures on payroll and equipment. No attempt was made to calculate the secondary and tertiary economic benefits of a product or service (for example the increase in agricultural yield attributable to the use of pesticides), or the present value of past structural investments (for example, airport terminals, acquisition of wilderness areas), or contributions to distributional equity.

Despite the differences in procedure, our analysis produced results similar to Starr's. Overall, there was a positive relation between benefits and risks (slope = .3, correlation = .55). Furthermore, at any given level of benefit, voluntary activities tended to be riskier than involuntary ones (compare alcohol and surgery or swimming and nuclear power).

To apply these results to Veg-E-Wax (see "Cost-Benefit Analysis" box), consider this technology to be an involuntary activity (imposed upon consumers) with a total economic ' ~nefit to the food industry of $1 billion. Its risks would be tolerable if the expected annual toll were less than 40 lives.

Although based upon an intuitively compelling logic, the method of revealed preferences has several drawbacks. It assumes that past behavior is a valid predictor of present preferences, perhaps a dubious assumption in a world where values can change quite rapidly. It is politically conservative in that it

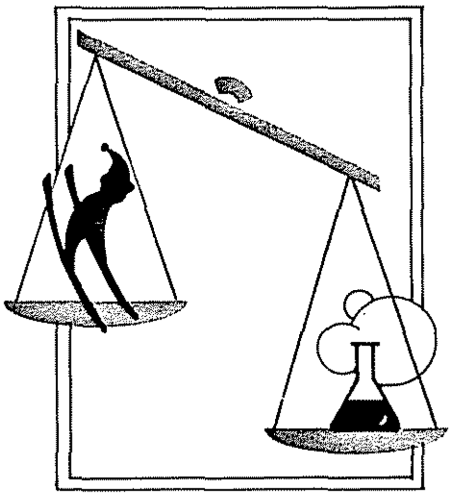

VOLUNTARY VS. INVOLUNTARY RISKS

One of Starr's "laws of acceptable risk" says the public seems willing to accept risks from voluntary activities (such as skiing) roughly 1,000 times greater than it tolerates from equally beneficial involuntary risks (such as food preservatives).

enshrines current economic and social arrangements. It ignores distributional questions (who assumes what risks and who gets what benefits). It may underweigh risks to which the market responds sluggishly, such as those involving a long time lag between exposure and consequences (as in the case of carcinogens).

It makes strong (and not always supported) assumptions about the rationality of people's decision making in the marketplace and about the freedom of choice that the marketplace provides. Consider the automobile for example. Unless the public really knows what safety is possible from a design standpoint and unless the industry provides the public with a set of alternatives from which to choose, market behavior may not indicate what a reflective individual would decide after thoughtful and intensive inquiry.

A revealed preference approach assumes not only that people have full information but also that they can use that information optimally, an assumption which seems quite doubtful in the light of much research on the psychology of decision making. Finally, from a technical standpoint, it is no simple matter to develop the measures of risks and benefits needed for the implementation of this approach.

Expressed Preferences

Both cost-benefit analysis and revealed preference analysis must infer public values indirectly, using procedures that may be both theoretically and politically untenable. The expressed preference approach tries to circumvent this problem by asking people directly what levels of safety they deem acceptable.

The appeal of this approach is obvious. It elicits current preferences; thus it is responsive to changing values. It also allows for widespread citizen involvement in decision making and thus should be politically acceptable. It allows consideration of all aspects of risks and benefits, including those not readily converted into dollars and body counts. Some ways of obtaining expressed preferences are through referenda, opinion surveys, detailed questioning of selected groups of citizens, interviewing "public interest advocates," and hearings.

Recently, we conducted a series of expressed preference studies paralleling Starr's revealed preference study.[21] We asked people to judge the total risk and benefit for each of thirty activities and technologies, including those used by Starr. Contrary to Starr's presumption, our respondents did not believe that society had managed these activities and technologies so as to allow higher risk only when higher benefit is obtained (see Figure 3a). In their view, society presently tolerates a number of activities with very low benefits and very high risks (alcoholic beverages, handguns, motor-

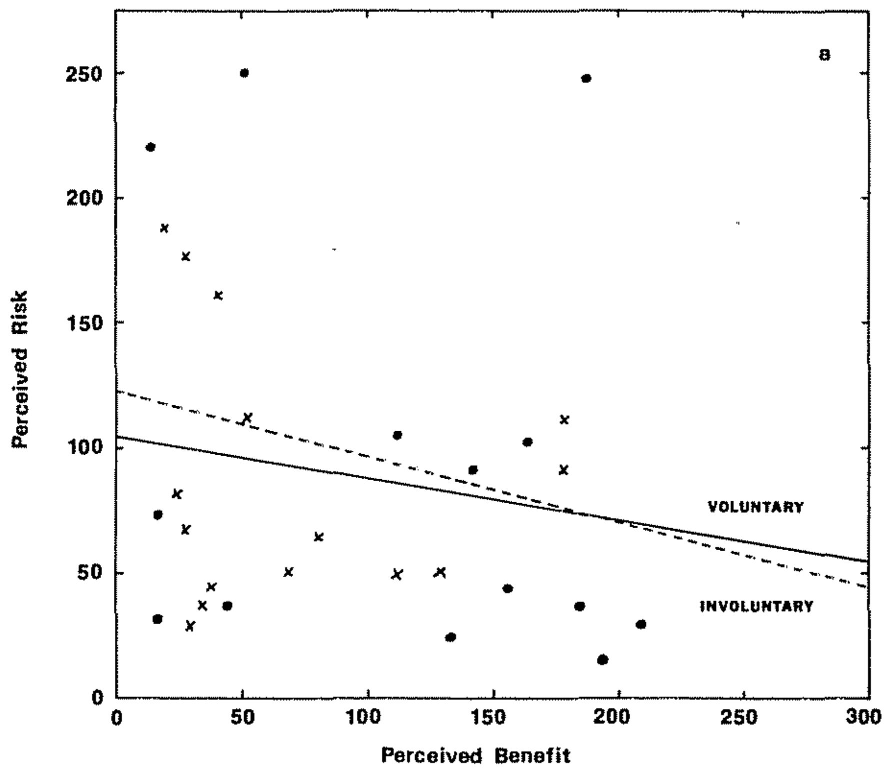

FIGURE 3. (a) A comparison of judged current risks and judged current benefits for 30 activities and technologies. The points are the average judgments of 76 members of the Eugene, Oregon, League of Women Voters. These individuals saw no systematic relationship between current risks and current benefits for these enterprises, a sharp contrast to the relationships shown in Figures 1 and 2. Similar patterns emerged from judgments by students and members of a community service club. As the two best-fit lines indicate, these individuals did not believe that society had to date achieved a lower level of risk for involuntary activities (indicated by *X*). (b) A comparison between judgments of acceptable risk levels and current benefits. These individuals (members of the League of Women Voters) believed that if risks were adjusted to socially acceptable levels, greater risks should be tolerated for more beneficial activities. Furthermore, for any given level of benefit, greater risks should be tolerated for voluntary (circled points) than for involuntary activities. Summarizing the contrasts between Figures 1, 2, and 3, the individuals we questioned believed that Starr's hypothesized relationships should be obtained in a society in which risk levels are adequately regulated. However, they also thought that their current world did not achieve that ideal. *Source:* B. Fischhoff, P. Slovic, S. Lichtenstein, S. Read and B. Combs (1978), Footnote 21.

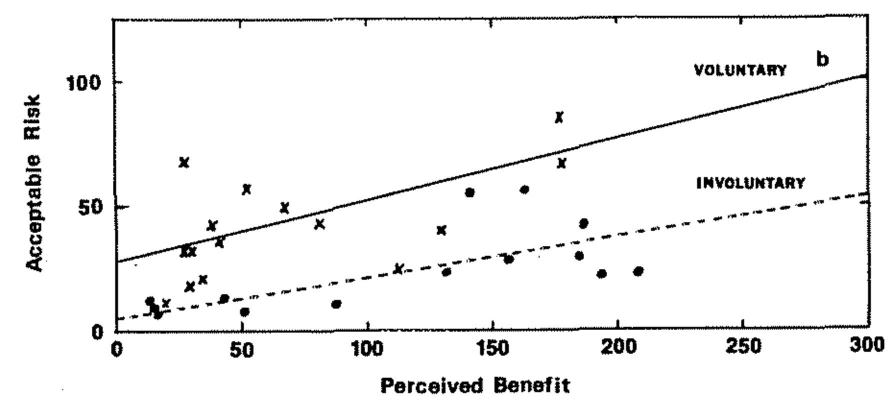

cycles, and smoking). Some very safe activities were judged to have very great benefits (antibiotics, railroads, vaccinations).

When we asked people what level of safety would be acceptable for each of the thirty activities and technologies, they responded that current levels were too safe 10 percent of the time, about right 40 percent of the time and too risky about 50 percent of the time ("too risky" was defined as "indicating the need for serious societal action"). Thus for these individuals, the historical record used by the revealed preferences approach apparently would not be an acceptable guide to future action.

When acceptable levels of safety were compared with perceived benefits, a relationship emerged much like the one obtained by Starr. Participants believed that greater risk should be tolerated for more beneficial activities and that a double standard is appropriate for voluntary and involuntary activities (Figure 3b).[22]

Similar studies were conducted with students, members of the (generally liberal) League of Women Voters, and members of a (generally conservative) community service club. Although the groups disagreed on the evaluation of particular items, their judgments showed the same general pattern of results as shown in Figure 3.

One frequent criticism of the expressed preferences approach is that safety issues are too complicated for ordinary citizens to understand. However, the results just cited suggest that, in some situations at least, motivated lay people can produce orderly, interpretable responses to complex questions.

A related criticism is that, when it comes to new and complex issues, people do not have well-articulated preferences. In some fundamental sense their values may be incoherent—not thought through.[23] In thinking about acceptable risks, people may be unfamiliar with the terms in which the issues are formulated (social discount rates, miniscule probabilities, megadeaths). They may have contradictory values (a strong aversion to catastrophic losses of life and a realization that they are not more moved by a plane crash with 500 fatalities than one with 300). They may occupy different roles in life (parents, workers, children) which produce clear-cut but inconsistent values. They may vacillate between incompatible, but strongly held, positions (freedom of speech is inviolate but it should be denied to authoritarian movements). They may not even know how to begin thinking about some issues (how to compare the opportunity to dye one's hair with a vague, minute increase in the probability of cancer twenty years from now). Their views may change over time (say, as the hour of decision or the consequence itself draws near), and they may not know which view should form the basis of a decision.

In such situations, where people do not know what they want, the values they express may be highly unstable. Subtle changes in how issues are presented—how questions are phrased and responses are elicited—can have marked effects on their expressed preferences. The particular question posed may evoke a central concern or a peripheral one; it may help clarify the respondent's opinion or irreversibly shape it; it may even create an opinion where none existed before.

Three features of these shifting judgments are important. First, people are typically unaware of the extent of such shifts in their perspective. Second, they often have no guidelines as to which perspective is the appropriate one. Finally, even when there are guidelines, people may not want to give up their own inconsistency, creating an impasse.

Natural Standards

A shared flaw of the approaches described above is that all of them are subject to the existing limitations of society and its citizens. It might be desirable to have a standard of safety independent of a particular society, especially for risks whose effects are collective, cumulative, or irreversible. One such alternative is to look to "biological wisdom" to insure the physical well-being of the species.[24] Rather than examining (recent) historical time for guidelines, one might look to geological time, assuming that the optimal level of exposure to pollutants is that characteristic of the conditions in which the species evolved.

Specific proposals derived from this approach might be to set allowable radiation levels from the nuclear fuel cycle according to natural background radiation and to set allowable levels of chemical wastes according to the levels found in archaeological remains.[25] These standards would not constitute outright bans, as some level of radiation-induced mutation is apparently good for the species and traces of many chemicals are needed for survival. Since exposure has varied from epoch to epoch and from place to place, one could establish ranges of tolerable exposure.

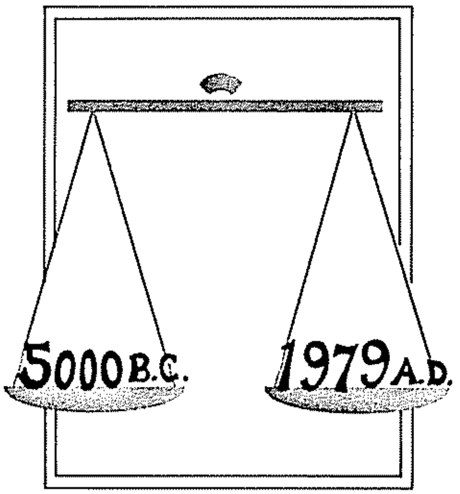

NATURAL STANDARDS

A commitment to "natural standards" assumes that the optimal level of exposure to pollutants is that characteristic of the conditions in which the species evolved.

Perhaps the best-known criteria for risk acceptability based on natural standards are those for ionizing radiation set by the International Commission on Radiological Protection. The standards set by this small, voluntary, international group are subscribed to by most countries in the world. Their underlying assumptions include the following:

The maximum permissible dose levels should be set in such a way that, in the light of present knowledge:

(a) they carry a negligible *probability of severe somatic or genetic injuries; for example, leukemia or genetic malformations that result from exposure to individuals at the maximum permissible dose would be limited to an* exceedingly small fraction *of the exposed group; and*

(b) the effects ensuing more frequently are those of a minor nature that would not be considered unacceptable by the exposed individual and by the society of which he is a part. Such frequently occurring effects might be, for example, modifications in the formed elements of the blood or changes in bone density. Such effects could be detected only by very extensive studies of the exposed individual. Effects such as shortening of life span, which may be proportional to the accumulated dose, would be so small that they would be hidden by normal biological variations and perhaps could be detected only by extensive epidemiological studies.[26]

Figure 4 shows how U.S. Atomic Energy Commission standards compared with natural background levels of radiation in 1976. It also compares current levels of SO_2 and NO_2 with background levels, indicating the implications of invoking natural standards in these contexts.

Natural standards have a variety of attractive features. They avoid converting risks into a common monetary unit (like dollars per life lost). They present issues in a way that is probably quite compatible with people's natural thought processes. Among other things, this approach can avoid any direct numerical reference to very small probabilities, for which people have little or no intuitive feeling.[27] Use of natural standards should produce consistent practices when managing the same emission appearing in different sources of hazards.

As a guide to policy, natural standards are flawed by the fact that our natural exposure to many hazards has not diminished. Thus, whatever new exposure is allowed is an addition to what we are already subjected to by nature and thereby constitutes excess "unnatural" exposure (although conceivably within the range of toleration).

A second problem is that most hazards increase some exposures and reduce others. Trading off different exposures brings one back to the realm of cost-benefit analysis.

Another problem arises when one considers completely new substances for which there is no historical tolerance (saccharin, for example). In such cases, a policy based on natural standards would tolerate none of the substance at all, unless it involved no risk. The Delaney Amendment, which outlaws the addition of any known carcinogen to food, is consistent with this approach.

The technical difficulties of performing this type of analysis are formidable. Indeed, while there may be some hope of assessing natural exposure to chemicals and radiation that leave traces in bone or rock, appraising the natural incidence of accidents and infectious disease is probably impossible. Furthermore, should such an analysis be completed, it would quickly become apparent that the ecology of hazard in which humans live has changed drastically over the eons—mostly for the better, as in the case of the reduced incidence of infectious disease.[28] The biological wisdom (or importance) of restoring one component of the mix to its prehistoric values would demand careful examination.

In addition to whatever difficulties there may be with their internal logic and implementation, natural standards are likely to fail as a sole guide to policy because they ignore the benefits that accompany hazards and the costs of complying with the standards.

Multiple Hazards

Our discussion so far has focused on the acceptable risk associated with individual hazards. What additional problems are created by considering many hazards at once? There are some 60,000 chemicals and 50,000 consumer products in common use in the United States.[29] If even a small fraction of these presented the legal and technical complexities engendered by saccharin or flammable sleepwear (not to mention nuclear power), it would take legions of analysts, lawyers, toxicologists, and regulators to handle the situation. If hazards are dealt with one at a time, many must be neglected. The instinctive response to this problem is to deal with problems in order of importance. Unfortunately, the information needed to establish priorities is not available; the collection of such data might itself swamp the system.

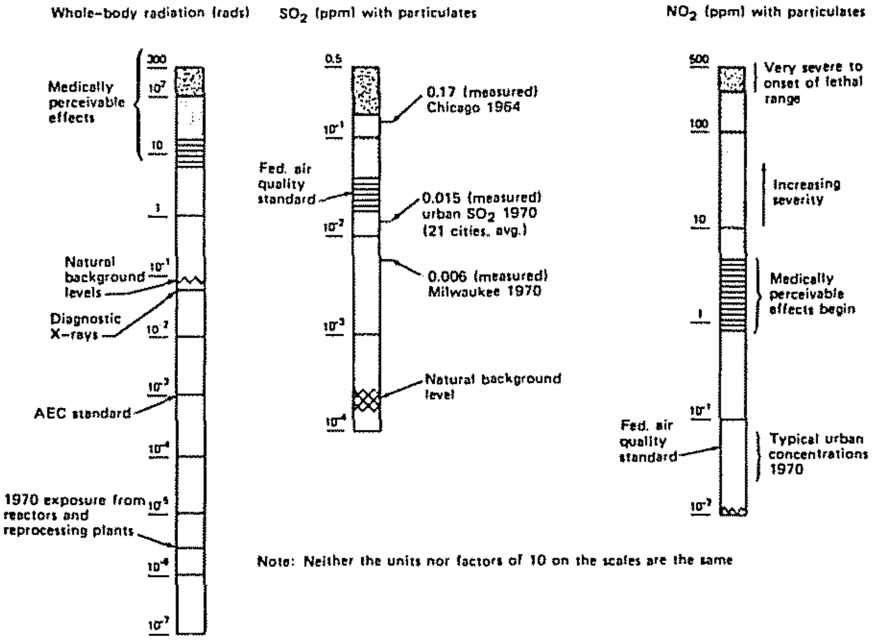

FIGURE 4. Comparison of pollutant standards, background levels, man-made exposures, and health effects for radiation, SO_2 and NO_2. *Source*: WASH-1224, U.S. Atomic Energy Commission.

Even if legions of hazard managers were available, the wisdom of tackling problems one at a time is questionable. Responsible management must ask not only which dangers are the worst but which are the most amenable to treatment. A safety measure that is reasonable in a *cost-benefit* sense may not seem reasonable in a *cost-effectiveness* sense. That is, if our safety dollars are limited, finding that the benefits of a particular safety measure outweigh its costs does not preclude the possibility that even greater benefits could be reaped with a like expenditure elsewhere. The hazard-by-hazard approach may cause misallocation of resources across activities (for instance, giving greater protection to nuclear plant operators than to coal miners) or even within activities (protecting crop dusters but not those in the fields below).[30]

The cumulative danger from a problem that appears in many guises may be hidden from a society that tackles hazards one by one. The current cancer crisis seems to reflect an abrupt realization of the cumulative impact of a risk distributed in relatively small doses over a very large number of sources. The nuclear industry has only recently been alerted to the possibility that temporary workers who receive their legal limit of radiation exposure in one facility frequently move on unnoticed to another and another.[31]

Proponents of new products or systems can often argue persuasively that the stringent risk standards imposed upon them by the public constitute an irrational resistance to progress. After all, many currently tolerated products have much greater risks with appreciably less benefit. The public may, however, be responding to its overall risk burden, a problem outside these proponents' purview. From that perspective, one of the obvious ways to reduce a currently intolerable risk level is to forbid even relatively safe new hazards unless they reduce our dependence on more harmful existing hazards.

Treating hazards individually may obscure solutions as well as problems. Hazard managers must worry not only about how to trade lives and health for dollars but also about how to do so in an equitable fashion. Resolving equity issues in the context of an individual hazard often demands either heroic theoretical assumptions or considerable political muscle. Looking at the whole portfolio of hazards faced by a society may offer some hope of circumventing these problems. No one escapes either the risks or the benefits of all aspects of a society. Indeed, they are often implicitly traded between individuals. I live below the dam that provides you with hydroelectric power in the summer while you live near the nuclear power plant that provides me with electricity in the winter. In this example, the participants might view the trade as equitable, without recourse to complex distributional formulas. While such simple dyads may be rare, looking at the total distribution of risks and benefits in a society may possibly produce clearer, sounder guidelines for resolving equity issues than would solutions generated for individual hazards.

Responsible management must ask not only which dangers are the worst but which are most amenable to treatment. Even if a safety measure is cost-beneficial, greater benefits may be reaped with a like expenditure elsewhere.

Facing Political Realities

Models that do not capture the critical facts about a hazard will not pass muster before the scientific community. Approaches that fail to represent the political realities of a situation will be rejected by those interests that are underrepresented. No one method can serve the needs of all the environmentalists, industrialists, regulators, lawyers, and politicians involved with a particular hazard. These people appropriately view each specific decision as an arena in which broader political struggles are waged.

In theory, any of the approaches described here should find some support among "public interest" advocates and some resistance among technology proponents since all of them make the decision process more open and explicit than it was in the dark ages of hazard management when matters were decided behind closed doors. However, the enchantment of the public wanes some when closed doors are replaced by opaque analyses that effectively transfer power to the minute technical elite who perform them.[32] In such cases, "public interest" advocates may resist formal analysis, feeling that avoiding disenfranchisement is more important than determining acceptable levels of risk. The battle brewing in the United States over the use of cost-benefit analysis to regulate toxic substances and other hazards may largely hinge on these concerns.[33]

For other members of the public, the openness itself is a sham, since each of these approaches makes the political-ideological assumption that society is sufficiently cohesive and common-goaled that problems can be resolved by reason and without confrontation. Sitting down to discuss a decision analysis would, in this view, itself constitute the surrender of important principles. Cooperation may even be seen as a scheme to submerge the opposition in paper work and abrogate its right to fight the outcome of an analysis not to its liking.[34] Such suspicions are most easily justified when the workings of the decision-making process are poorly understood. It is not hard to imagine the observers of a decision analysis accepting its premises but balking at its conclusions when the results of the analysis are complex or counter-intuitive. At the extreme, this would mean that people will only believe analyses confirming their prior opinions.

Proponents of a technology would probably prefer to have the determination of risk acceptability left to their own corporate consciences. Barring that (or the equivalent captive regulatory system), proponents may find it easier to live with adversity than with uncertainty. As a result, one would expect industry increasingly to advocate routinized approaches with rigorous deadlines for making decisions. From this perspective, the zenith of the influence of the Toxic Substances Control Act may have been reached immediately after its enactment. At that moment, industry practice could respond only by making all products as safe as possible, not knowing which substances would actually be dealt with nor how stringently. Cynically speaking, the sooner and more precisely the rules are laid down, the more efficacious the search for loopholes can be.

One could draw similar caricatures of the hidden agendas of other (would-be) participants in hazard management. The point of such an assessment is not to argue that reasonable management is impossible but that all approaches must be seen in their political contexts. Such a broadened perspective may help us to understand the motives of the various participants and the legitimacy that should be assigned to their maneuvers.

In so doing, a crucial issue will be deciding whether society should have higher goals than maximizing the safety of particular technologies. Such goals might include developing an informed citizenry and preserving democratic institutions. In this case, the process could be more important than the product, and it would be important for society to provide the resources needed to make meaningful public participation possible.[35] Such participation would require new tools for communicating with the public—both for presenting technical issues to lay people and for eliciting their values.[36] It might also require new social and legal forms, such as hiring representative citizens to participate in the analytic process, thereby enabling them to acquire the expertise needed by the governed to give their informed consent to whatever decision is eventually reached. Such a procedure might be considered a science court with a lay jury. It would consider any or all of the analytic techniques described here as possible inputs to its proceedings. It might also place the logic of jurisprudence above the logic of analysis, acknowledging that there is no single way to determine what risks are acceptable.

The forums in which safety issues are currently argued were not designed to deal with such problems. H. R. Piehler has, in fact, argued that the legal system could hardly have been designed more poorly for airing and clarifying the technical considerations which arise in product liability suits.[37] Much public opinion about hazards derives from the testimony of experts. Often this testimony is offered in rancorous debates between experts trying to cast doubt on the probity of their opponents.[38] In addition to creating negative attitudes toward scientists, such spectacles tend to destroy public confidence in the possibility of ever understanding or satisfactorily resolving these issues.

Natural disagreements in areas of incomplete knowledge are aggravated by the feeling that "bad evidence drives out good." A two-handed scientist ("on the one hand . . . while on the other . . .") may be bested by a two-fisted debater intent on acquiring converts. Decisions about controversial technologies might be improved if all participants publicly subscribed to an established code of behavior. Some possible rules might be:

- ✓ Never cite a research result without having a complete, accessible reference.
- ✓ Never cite as fact a result supported only by tenuous research findings
- ✓ Acknowledge areas in which you are not an expert (but are still entitled to an opinion).

Like rules of parliamentary procedure, this code would formalize values that many people espouse but have difficulty upholding in practice (fairness, mutual respect, etc.).

Muddling Through Intelligently

No approach to acceptable risk is clearly superior to the others. To exploit the contributions each of these methods can make, careful consideration must be given to the social and political world in which they are used and to the natural world in which we all live. Our social world is characterized by its lack of orderliness. Since hazards are not the only consideration in hazard-management decisions, the best we can hope for is some intelligent muddling through. Recognizing this, we should develop and apply the various approaches to hazard management not as inviolate ends in themselves but as servants to that process. The openness of formal analyses must be assured in order to avoid suspicion and rejection of whatever conclusions are finally reached. When the available numbers are not trustworthy, we should content ourselves with digitless structuring of problems. When good numbers are available, but the issues are

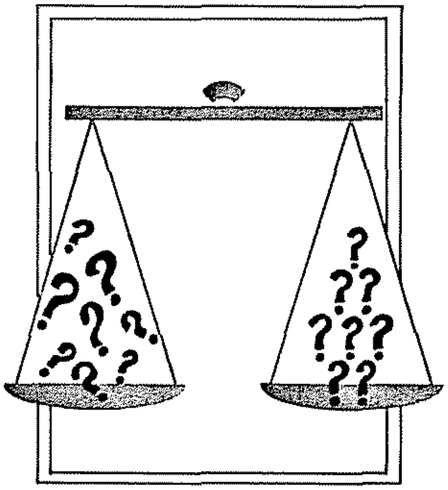

MUDDLING THROUGH INTELLIGENTLY

Uncertainty about facts and values in a disorderly social world means the various decision-making approaches must be viewed as tools rather than ends in themselves.

unfamiliar, great care must be taken in designing suitable presentations. When we do not know what goal we want to reach, value issues should be framed in a variety of ways and their implications carefully explored.

A distinctive characteristic of our natural world is that it typically is not and cannot be known to the desired degree of precision. We must not only acknowledge this uncertainty but also devote more of our efforts to determining its extent. The most critical input to many hazard management decisions may be how good our best guess is. The real alternatives may be: "If we don't understand it, we shouldn't mess with it" and "If we don't experiment, we'll never know what it means."[39]

Uncertainty about facts and uncertainty about values both imply that determining the acceptability of a hazard must be an iterative process, partly because, as time goes on, we learn more about how a hazard behaves and how much we like or dislike its consequences. In other words, it takes experience which acknowledges the experimental nature of life to teach us what the facts are and what we really want.

Iteration is essential to any well-done formal analysis. A measure of the success of any analysis is its ability to inform (as well as to reflect) our beliefs and values. Once the analysis is completed, we may then be ready to start over again, incorporating our new and better understandings. In this light, many of the non-political critiques generated by the *Reactor Safety Study* (the "Rasmussen report")[40] reflect its success in deepening the respondents' perspectives. As an aid to policy, the study's main weakness was in attempting to close the books prematurely and thereby failing to take adequate account of these criticisms.

While a good analysis should be insightful, it need not be conclusive. At times, it may not be possible to reach any analytic conclusion, for example, when inter- and intra-personal disagreements are too great to be compromised. If people do not know what they want or if a topic is so politicized that no solution will ever be acceptable, analysis should perhaps best be treated as a process for deepening knowledge and clarifying positions. Performing the sort of calculations that lead to a specific recommendation would, in such cases, only create an illusion of analyzability.

A Combined Approach

The disciplinary training of scientists shows them how to get the right answers to a set of specially defined problems. The problems raised by hazard management are too broad to be solved by any one discipline. No one knows how to get the right answer. All we can do is avoid making the particular mistakes to which each of us is attuned. The more scientific and lay perspectives applied to a problem the better chance we have of not getting it wrong.

Just as no single discipline has all the answers, no one of the approaches discussed above provides a sufficient basis for determining what levels of safety are acceptable. In attempting to solve the problems inherent in the other methods, each approach engenders problems of its own.

Are better approaches likely to come along? Probably not, for it seems as though all attempts to rule on the safety of particular hazards share common conceptual and operational difficulties whose source lies in the very attempt to reduce the problem to manageable size. What we can hope for is to understand the various approaches well enough to be able to use them in combination so that they complement one another's strengths rather than compound each other's weaknesses.

ACKNOWLEDGEMENTS

This research was supported by the National Science Foundation under Grant ENV77-15332 to Perceptronics, Inc. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation. We are grateful to Bob Kates and Richard Wilson for comments on an earlier draft of the manuscript and to Michael Enbar for much of the analysis in Figure 2.

NOTES

1. Business Week, December 25, 1978. pp. 54-68.

2. T. Bick and R. E. Kasperson, "Pitfalls in Hazard Management: The Consumer Product Safety Commission Experiment," Environment, October 1978, p. 30.

3. B. J. Culliton, "Toxic Substances Legislation: How Well Are Laws Being Implemented?" Science, 201 (1978): 1198-1199.

4. C. Sinclair, P. Marstrand, and P. Newick, Innovation and Human Risk, Centre for the Study of Industrial Innovation, London, 1972.

5. R. A. Howard, J. E. Matheson, and D. L. Owen, "The Value of Life and Nuclear Design." in D. Okrent and E. Cramer (eds.), Probabilistic Analysis of Nuclear Reactor Safety, American Nuclear Society Publication, LaGrange Park, Ill., 1978; IV, 2-1 to IV 2-9.

6. Other discussions of acceptable risk criteria may be found in W. Lowrance, Of Acceptable Risk, Wm. Kaufman, Los Altos, Calif., 1976, and W. Rowe, An Anatomy of Risk, Wiley, New York, 1977.

7. See B. Fischhoff, "Cost-Benefit Analysis and the Art of Motorcycle Maintenance," Policy Sciences, 8 (1977): 177-202, or E. Stokey and R. Zeckhauser, A Primer for Policy Analysis, Norton, New York, 1978.

8. R. V. Brown, A. S. Kahr, and C. Peterson, Decision Analysis for the Manager, Holt, Rinehart & Winston, New York, 1974; R. A. Howard, J. E. Matheson, and K. E. Miller, Readings in Decision Analysis, Stanford Research Institute, Menlo Park, Calif., 1976.

9. B. Fischhoff, C. Hohenemser, R. E. Kasperson and R. W. Kates, "Handling Hazards," Environment, September 1978, p. 20.

10. T. Najarian. "The Controversy over the Health Effects of Radiation," Technology Review, 81 (1978): 74-82.

11. P. Slovic, B. Fischhoff, and S. Lichtenstein. "Rating the Risks," Environment, 21 April 1979, p. 14.

12. J. Linnerooth. The Evaluation of Life-Saving: A Survey, Joint IAEA/IIASA Research Report, International Atomic Energy Agency, Vienna, Austria, 1977.

13. R. Holmes, "On the Economic Welfare of Victims of Automobile Accidents," American Economic Review, 60 (1970); R. Kidner and K. Richards, "Compensation to Dependents of Accident Victims," Economic Journal, 84 (1974).

14. J. Linnerooth, note 12.

15. R. Thaler and S. Rosen, "The Value of Saving a Life: Evidence from the Labor Market," in N. Terleckyj (ed.), Household Production and Consumption, Columbia University Press, New York, 1976.

16. E. Rappoport, unpublished Ph.D. dissertation, Department of Economics, University of California, Los Angeles, 1977.

17. These assumptions are essentially those underlying the revealed preference approach described below.

18. B. Fischhoff, "Clinical Decision Analysis," Operations Research, in press.

19. See P. Slovic, B. Fischhoff, and S. Lichtenstein, note 11.

20. C. Starr. "Social Benefit versus Technological Risk," Science, 165 (1969): 1232-1238.

21. B. Fischhoff, P. Slovic, S. Lichtenstein, S. Read, and B. Combs. "How Safe Is Safe Enough? A Psychometric Study of Attitudes towards Technological Risks and Benefits." Policy Sciences, 8 (1978): 127-152.

22. One complication of this latter relationship is that the degree of voluntariness of some activities proved to be rather ambiguous (e.g., handguns). A second is that double standards were also observed with other qualitative aspects of risk such as perceived control, familiarity, knowledge, and immediacy.

23. B. Fischhoff, P. Slovic, and S. Lichtenstein. "Knowing What You Want: Measuring Labile Values," in T. Wallsten (ed.), Cognitive Processes in Choice and Decision Behavior, Erlbaum, Hillsdale, N. J., in press.

24. Not to mention the well being of other species. L. H. Tribe, C. S. Schelling, and J. Voss, When Values Conflict: Essays on Environmental Analysis, Discourse and Decision, Ballinger Publishing Co., Cambridge, Mass., 1976.

25. J. Ericson, H. Shirahata, C. Patterson. "Skeletal Concentrations of Lead in Ancient Peruvians," New England Journal of Medicine, in press.

26. K. Z. Morgan, "Present Status of Recommendations of the International Commission on Radiological Protection, National Council on Radiation Protection and Federal Radiation Council," in A. M. F. Duhamel (ed.), Health Physics, 2, Pergamon Press, New York, 1969.

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

28. R. C. Harriss, C. Hohenemser, and R. W. Kates, "Our Hazardous Environment," Environment, September 1978, p. 6.

29. T. Bick and R. E. Kasperson, note 2; T. H. Maugh, "Chemical Carcinogens: The Scientific Basis for Regulation," Science, 201 (1978): 1200-1205.

30. D. M. Berman, "How Cheap Is a Life?" International Journal of Health Sciences, 8 (1978): 79-99.

31. Federal Register, February 6, 1978, pp. 1865-1868.

32. L. McGinty and G. Atherly, "Acceptability versus Democracy," New Scientist, 74 (1977): 323-325.

33. Chemical and Engineering News, January 22, 1978, p. 7; L. J. Carter, "An Industry Study of TSCA: How to Achieve Credibility," Science, 203 (1979): 247-249.

34. S. K. Fairfax, "A Disaster in the Environmental Movement," Science, 199 (1978): 743-748.

35. B. N. Casper, "Technology Policy and Democracy," Science, 194 (1976): 29-35.

36. P. Slovic, B. Fischhoff, and S. Lichtenstein, note 11.

37. H. R. Piehler. A. D. Twerski, A. Weinstein, and W. A. Donaher, "Product Liability and the Technical Expert," Science, 186 (1974): 1089-1093.

38. A. Mazur, "Disputes between Experts." Minerva. 11 (1973): 243-262.

39. R. Goodwin, "Uncertainty as an Excuse for Cheating Our Children: The Case of Nuclear Wastes," Policy Sciences, 10 (1978): 25-43.

40. Atomic Energy Commission, Reactor Safety Study: An Assessment of Accident Risks in U. S. Commercial Power Plants, WASH-1400, The Atomic Energy Commission, Washington, D. C., 1975.

This is the eighth in a series of articles on technological hazards and their management edited by Christoph Hohenemser and Robert W. Kates of Clark University. Other articles in this series are: "Our Hazardous Environment" and "Handling Hazards," September 1978; "Pitfalls of Hazard Management," October 1978; "Mercury: Measuring and Managing the Risk," November 1978; "Target: Highway Risks" (two parts), January/February and March 1979; and "Rating the Risks," April 1979.