Risks of Risk Decisions

Chauncey Starr and Chris Whipple

Technology creates many risks. Determining which risks are acceptable is an important national issue. It pervades major sectors of our economy: In food production we face decisions about pesticides and preservatives; transportation risks are increasingly regulated; and a central issue in energy policy is the controversy over the risks from power plants. Regardless of whether the seriousness of technological risk is only now being recognized, or, alternatively, that the preoccupation with risk and regulations is an overreaction, it is clear that the cost to society of the conflict over accepting technological risks is great. These costs stem from the anxiety suffered by those who are dismayed by the conflicting information about these risks, and from the litigation, misplaced investment, retrofits, and costly delays that result from industry's inability to predict the acceptance of risk by the public.

Summary. The analytical approaches utilized for evaluating the acceptability of technological risk originate from analogies to financial cost-benefit risk analysis. These analogies appear generally valid for viewing risk from a societal basis, but are not applicable to individual risk assessments. Conflicts arising from these different views of risk assessment provide insights to the origins of individual, intuitive evaluations. Societal risk decisions made under conflict represent political compromises, and the resulting decision process creates substantial conflict costs. The pragmatic use of quantitative risk criteria (safety targets) may be useful in reducing these costs.

Risk assessment is growing in importance as a system design tool. The final configuration of all technical systems is the outcome of a common design sequence. The first task of a system designer is the development of a workable basic concept. The second task is reducing the vulnerability of the system to failures of component parts, including human participants. The final task is balancing the benefits and risks of the new system, starting with the internalized economic costs. The external effects have rarely been analyzed, and it is only in recent decades that we have become deeply concerned with this difficult but important part of the design process.

Risks created by technical systems arise either from routine external effects considered acceptable at the time of design, or from abnormal conditions that are not part of the basic design concept and its normal operation. Most abnormal events usually impair or stop the operation of the technical system, and may threaten the operators. The usual external effect is the loss of operational benefits to the users of the output. The major internal consequences of failures are borne by the operating institution. The timely diffusion within the institution of information about such failures usually stimulates rapid modifications to reduce the ratio of failure costs to benefits. Less frequently, a failure results in effects outside the institutional boundary, creating a public risk—and a potential cost to the public. These external costs are usually difficult to evaluate, and here the informational mechanism for system modification is usually cumbersome and slow. In recent years such modifications have been made because of an increasing public concern over the inherent risks and costs arising from previously acceptable external effects, both occasional and routine. For these reasons, the importance of risk as a design criterion is increasing.

The basic truisms about risk are readily recognized. First, everyone knows that risk taking is an accepted part of life. Living can be fun, but it is also dangerous (just how dangerous can be difficult to measure). Second, everyone reacts differently to risks taken voluntarily and to risks that are imposed by some outside group. Third, decisions imposing risks on us are being made all the time. This results in the fourth truism: a conflict is inherent when a group imposes a risk on others. Historically, such conflicts have been resolved by compromise, but rarely to everyone's satisfaction.

It is, therefore, characteristic of the functioning of an organized society that conflicts arise from the balancing of public benefits and involuntary risks to the individual. Because such conflict is unavoidable, our problem is how to manage and minimize it.

How should group decision processes operate to minimize social costs and maximize social benefits? Group processes range from anarchy to dictatorship. In most of the industrial world, we enjoy a medium between these, but the processes for decision-making have themselves become contentious issues.

Social costs include intangibles, and the question immediately evident is what costs are included and how are they weighted. It is obvious that if we have a decision process, and if we know how to determine costs to the individual, we still have a problem with the full disclosure of all the social costs. What is full disclosure? Do we include the options for societal risk management as part of full disclosure: that is, the cost of the alternatives for managing the risk? Does it include all present events, future events, the people who get the benefits, and the people who bear the costs? We have geographic distributions, time distributions, demographic distributions—all of these are included by the term full disclosure. Where do we draw the boundaries?

Decisions are not made by institutions; the decision process involves people. The government typically works through agencies and committees, so that, in fact, it is a few people in the agencies and a few people on the committees who really decide what happens. How do we allocate the responsibility and the costs of bad decisions? How do we functionally connect authority, responsibility for outcomes, and costs?

After we establish the social costs, how do we set out priorities? How do we determine the relative merits of various outcomes? That is a subject for a separate study, because, of course, value systems depend on culture, background, economic status, and all kinds of psychological factors.

Part of the problem in risk assessment comes from confusions that arise during discussions of the subject—confusions about reality, analysis, and individual

Chauncey Starr is vice chairman of the Electric Power Research Institute, Palo Alto, California 94303. He is also director of its Energy Study Center, where Chris Whipple serves as technical manager.

0036-8075/80/0606-1114$01.50/0

perceptions. Reality is what has happened or what will happen. Analysis is a process based on collected data, anecdotal cases, and statistics, any of which may or may not be correct; and, based on these, we invent simplified models to predict an outcome. The result, of course, is a large uncertainty in the predictions.

What is the intuitive perception of the individuals involved? Involuntary risks are perceived differently by individuals. Their perceptions may be far from reality. So, in discussing public acceptance of risk, we have to distinguish between the uncertain reality of what may occur, the uncertain analysis of predicting it, and the variable perception of its potential. Similar confusions exist, incidentally, over social costs and social benefits, which are also involved. As an illustration, who in the year 1900 could have predicted the social costs and benefits of the automobile?

Finally, people's perceptions of probabilities are frequently in gross error. The accident at Three Mile Island proved very little about probabilities of such events. The inadequacy of such single events for providing probability numbers can be explained analytically, but the political response and the public perceptions are often based on single events. So even if a professional group develops analytic answers, it has difficulty persuading the public to accept them.

Recognizing all these difficulties, it is nevertheless important to explore the subject of risk management in order to improve the quality of decision-making.

Analytic and Judgmental Approaches

A question implicit in the term acceptable risk is "acceptable to whom?" Certainly congressional approval of any method for making risk-benefit decisions establishes its legitimacy, but a public consensus is needed to sustain its use. Defining this consensus is difficult because there are technologies that are favored by a majority, or at least by a plurality, but are opposed by extremely motivated individuals and groups (for example, those who fight water fluoridation and nuclear power). Because of our experience with other political issues in which similar divisions of public opinion occur (abortion, gun control), we know that we should not be optimistic over the prospects that a regulatory approach can neutralize these controversies. Problems such as these raise issues, such as the definition of majority versus minority

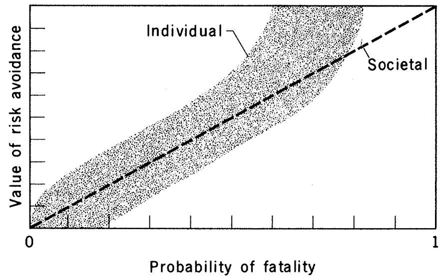

Fig. 1. Value systems for risk.

rights and the scope and limit of due process, that are well beyond those normally associated with risk management.

Congress has not defined "acceptable" risk levels, except for the few cases in which a zero risk approach was mandated. Far more frequently (*1*), Congress delegates responsibility for judging risk acceptability to regulatory agencies with the criteria that protection be provided against "unreasonable" risks. The methods by which these agencies interpret "reasonableness" range from a formal analysis of risk, benefits, and alternatives to purely subjective evaluations.

Analytic Approaches

The attraction of analytic methods (cost-benefit analysis, decision analysis) is their capacity to make explicit the assumptions, value judgments, and criteria used for making a decision. The analytic approaches are considered logically sound and sufficiently flexible to accept any value system. Given a specific set of values and criteria, a cost-benefit analysis could ideally indicate the decisions that would best balance technological risk and benefit (assuming that both tangible and intangible costs and benefits are included). But in reality it is difficult to measure group values, and at best the analytic methods can only be used to reach a rough approximation of the social cost and benefits that characterize a decision.

The debate over the relative merits of these approaches generally focuses on the effects of incomplete information (omitted and uncertain risks, benefits, and values), neglect of distributional effects, and other errors of simplification. It is not our intent to review the merits of these methods as commonly practiced; that has been done elsewhere (*2–6*).

Physical Versus Financial Risk

Because of our use of the term risk as the probability of either financial or physical damage, we may tend to uncritically allow the use of premises about the acceptance of the risk to "life and limb" (*7*) to be based on an analogy to financial risk taking.

From the societal viewpoint, the presumption that risk equals cost may be valid in most cases. For example, the cost of the risk of death is sometimes calculated as being equal to the discounted net earnings of those killed. This method, now out of favor, operates as if the loss of lives were equivalent to the breakdown of productive machines.

Similarly, the value assigned to resilience (*8*) leads to a desire to avoid catastrophic accidents that parallels the strategy in which investments are diversified in order to limit losses under adverse conditions. Perhaps recognizing the differences between these two types of risk, Zeckhauser (*9*) argued that, on a per fatality basis, the social cost of multiple-fatality accidents is lower than that of a single fatality because fewer survivors are affected. For example, the social cost of the loss of a city or a family is less than that of an equivalent number of independent, dispersed fatalities. Although the basis for this argument is apparent, it is also incomplete. For example, it ignores the value placed on the continuation of a family line; the importance of this value is evident in the draft deferment that was given to sole surviving sons. Similarly, Wilson (*10*) noted,

Fig. 2. Comparison of perceived risk with actual risk. [Courtesy of the American Psychological Association]

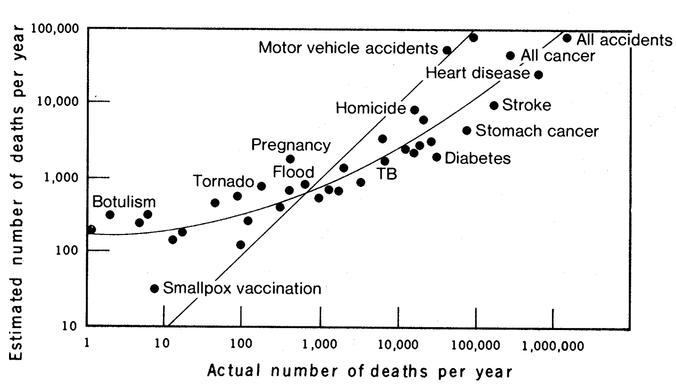

Small accidents throughout the world kill about 2 million people each year, or 4 billion people in 2000 years. This is "acceptable" in the sense that society will continue to exist, since births continually replace the deaths. But if a single accident were to kill 4 billion people, that is, the population of the whole world, society could not recover. This would be unacceptable even if it only happened once in 2000 years.

Another example in which the ability to generalize from financial cost-benefit analysis has been questioned is when physical risks are distributed across time. Arrow (*11*) argued that these risks should be treated as other costs, and discounted accordingly; other analysts of the issue have questioned the validity of this approach, and looked for alternative methods for guidance on how to judge equity in intergenerational risk trade-offs (*12*).

Application of Expected Value to Individual Risk Assessment

Although the above analogy may be valid for the collective view of the cost of risk, it may not apply to the intuitive evaluation of risk. As Fig. 1 illustrates, the individual and societal evaluations of risk are quite different (*8*). In the societal view, the presumption of a linear relation between risk and the cost of that risk may be quite valid. But as Howard (*13*) pointed out, the individual evaluation of that cost is necessarily nonlinear, and becomes infinite as the probability approaches unity.

The use of an expected value or expected utility model is based on the premise that expected cost is simply the product of the probability of an outcome and the evaluation of that outcome. But in the individual's view of the risk of death, this is not valid, for this product is very large or infinite.

It is often presumed that the individual evaluation of the cost of risk is linear over a probability range of interest (*9, 13*), but there is little firm evidence to support any hypothesis about the shape of this curve, as far as we know. Under the common conventions of risk analysis, the slope of the curves in Fig. 1 is referred to as the value of life. The politicians' old saw that life is of infinite value can be reconciled if this refers to the individual evaluation of one's own life. This viewpoint is not inconsistent with the assignment of finite costs to risks; it is the application of an expected value model that is inappropriate for this evaluation.

A second drawback with the application of expected value to individual risk evaluation stems from the tendency by analysts to seek to accommodate differences of opinion entirely within the assignment of utility (*14*). This is because in most decision or cost-benefit analyses, the probability estimates are considered roughly valid because they are based on available data, engineering models (such as fault trees), and expert opinion. But in a study of public attitudes about nuclear power, the bulk of the disagreement was found to be due to different beliefs about accident probability (*15*). Although it may be perfectly valid to base public policy on expert estimates and data, the attempt to reconcile differences in the assignment of costs and values is misdirected if, in fact, the controversy over technological risk is due to divergent beliefs about probability.

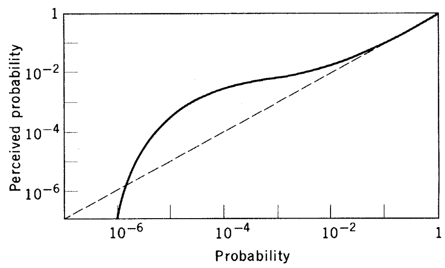

Fig. 3. Perception of probability.

Intuitive Versus Analytical Risk Assessment

We now consider the implications of the premise that risk acceptance is ultimately inseparable from the psychology of risk perception and evaluation. A corollary to this premise is the assumption that when the results of intuitive risk assessments differ significantly from those of the analytical methods, conflict follows.

It seems clear that intuitive and quantitative risk-benefit assessments can produce quite different results, even given the capacity of the analytical approaches to accommodate complex values relating to different risk attributes. The differences of opinion over probability assignments are not limited to those risks for which data are not available; many people intuitively fear travel by airplane more than by automobile, yet aviation is safer. Explanations of this effect focus on the degree of individual control over risk (*16*), the conditional probability of survival given an accident, and the catastrophic nature of airplane accidents (*17*).

The difficulty that arises from these differences in assessment stems from the dual meaning of acceptable risk. The analytical methods help regulators set standards that implicitly define acceptable risk. But the intuitive individual assessments of acceptability can overrule these decisions through the political process. The repeal of the seat belt interlock regulation and recent congressional action to prevent a ban on saccharin are cases in which public opinion resulted in a policy change.

Intuitive Risk-Benefit Analysis

Given the role of individual judgments of (physical) risk and benefit in determining the political acceptability of specific technologies, it seems particularly valuable to try to understand intuitive risk-benefit analysis. Efforts to develop this understanding were made by Starr (*8, 18, 19*), whose approach was based on a study of historically accepted risk (revealed preferences) and by Fischhoff *et al.* (*17*) and Slovic *et al.* (*20*), whose approach was usually based on risk-taking behavior as determined by questionnaire (expressed preferences). An additional source of information is the study by Lawless (*21*) of many controversies over technology. If we assume that many of these controversies arose because of intuitive estimates of unreasonably high risk (not true in all the cases described; some cases, such as the thalidomide tragedy, were due to late identification of a risk), then the common characteristics of risk and benefit in these controversies may indicate important factors in the intuitive risk process. Lawless did this, and his findings confirm those of other studies in identifying catastrophic potential and lack of individual control over risk as "factors that influence the impact of the threat."

Understanding the Intuitive Process

There is an attraction to try to develop an understanding of intuitive risk-benefit decisions by constructing parallels to the analytical methods. This approach leads to a model of intuitive decision-making in which subjective judgments of the probability and consequence of undesirable outcomes are somehow combined to produce a perceived risk; parallel judgments provide a perceived benefit; the two are then compared to provide intuitive judgment of acceptability.

This model is quite broad; it does not specify the intuitive procedures for arriving at either perceived risk or benefit, or for their comparison. Even so, the available evidence suggests that this model may be incorrect. First, studies of intuitive decision-making in general (not limited or applicable to physical risks alone), have identified numerous decision-making rules that do not follow the model described above (*22*). Second, there is evidence to indicate that benefits are not intuitively evaluated independently from risks.

In the survey of subjective risk and benefit by Fischhoff *et al.* (*17*), perceived risk and perceived benefit were negatively correlated, due principally to the subjective evaluation of a number of things as high in risk and low in benefit (handguns, cigarettes, motorcycles, alcoholic beverages, nuclear power). When subjects were asked to judge "the socially acceptable level of risk," those who first took the benefits into consideration consistently reported higher levels of acceptability than did subjects who first evaluated risk, which reinforces the view that risks and benefits are not evaluated independently.

Despite the limitations of the perceived risk–perceived benefit view of deciding risk acceptability, we know of no better way to attempt to understand the intuitive processes for risk decisions. Support for this approach stems from the fact that the acceptability of a risk has been found to increase with increasing benefit both by Starr (*18*) and Fischhoff *et al.* (*17*).

Benefits

Little work has been done to characterize the perceived benefits of technological activities. Starr (*18*) found a correlation between risk and "benefit awareness," which he described as a crude measure of public awareness of social benefits. This measure was based on the relative level of advertising, the percentage of the population involved in the activity, and a subjective judgment of the usefulness of the activity. The survey by Fischhoff *et al.* (*17*) included a subjective ranking of benefits, but no attempts were made to relate perceived benefit with any characteristics of that benefit.

Probability Perception

By far the most studied and best understood component of intuitive risk-benefit analysis is risk perception and evaluation. There is excellent literature on the subjective estimation of probability (*23*).

One aspect of the interpretation of probability that has been noted repeatedly is the intuitive handling of very low probabilities. As Mishan (*7*) noted: "One chance in 50,000 of winning a lottery, or of having one's house burned down, seems a better chance, or greater risk, than it actually is." The same observation was made by Selvidge (*24*). Lichtenstein *et al.* (*25*) found similar results (Fig. 2) when they asked people to estimate the number of fatalities from specific causes annually in the United States: "The full range of perceived risk is only about 10,000 while the corresponding actual range is closer to 1,000,000." Similar results were found in another survey in which risk was ranked subjectively (*17*).

The influence of this perception is important when we recall that the expected value or expected utility model calculates that a change in event probability by a factor of 1000 produces a change in expected value or utility by a like amount. If the probability is perceived as having changed by a much smaller amount, then it would not be surprising to find that an intuitive evaluation of risk is less sensitive to probability changes. This can be extremely important for low-probability, high-consequence risks, because probabilities lying below an intuitively understandable range may be overestimated.

We postulate that this is only true to a point (see Fig. 3). Although we selected these scales judgmentally, their chief purpose is to illustrate that, at some low level of probability, the intuitive interpretation goes from "low" or "unlikely" to "negligible" or "impossible." This hypothesis can be used to explain behavior regarding seat belt use and perhaps smoking. In a study of seat belt use, Slovic *et al.* (*26*) noted that if the decision to wear seat belts is approached on a per trip basis, "we might expect that many motorists would find it irrational to bear the costs (however slight) of buckling up in return for partial protection against an overwhelmingly unlikely accident." They observed that "change of perspective, towards consideration of risks faced during a lifetime of driving, may increase the perceived probabilities of injury and death and, therefore, induce more people to wear seat belts. . . . Such differing perspectives may trigger much of the conflict and mutual frustration between public officials and motorists, each believing (with some justice) that their analysis of the situation is correct." Similarly, Jacobson (*27*) referred to carcinogenic "chemicals which pose minuscule hazards to individuals, but significant hazards to the population as a whole." This last point supports our premise that much conflict over technological risk is due to differences between intuitive and analytical risk-benefit analyses. If the hypothesis that perceived probability is effectively zero for some risks is valid, then the perceived risk of a short automobile trip without seat belts or of one cigarette may be zero.

This nonlinearity in probability perception indicates that even something apparently as basic as the unit of exposure used to evaluate risk can be influential. In his analysis, Starr (*18*) commented, "The hour-of-exposure unit was chosen because it was deemed more closely related to an individual's intuitive process in choosing an activity than a year of exposure would be."

Accepting, at least tentatively, the relation between perceived and actual probability (Fig. 3), we can see a basis for the controversy over catastrophic risks. As mentioned above, high-consequence, low-probability risks are of particular concern if their probabilities are overestimated subjectively. But when part of the public believes the probability is low, and another part believes it to be negligible, these beliefs lead to radically different evaluations. This may be the case with nuclear power and other risks of this type, and may be a key reason for the controversies over these risks.

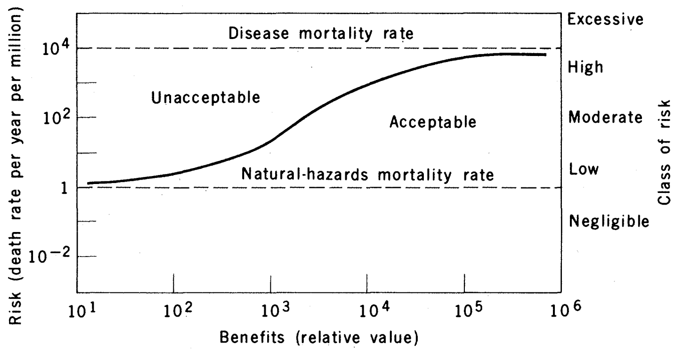

Fig. 4. Risk-benefit pattern for involuntary exposure.

Risks Distributed over Time

Given the apparent nonlinearities in risk evaluation depending on the unit of measurement, it seems reasonable to look for other perceptual factors related to the units in which risks are expressed.

A number of distinctions can be considered: risks can be immediate or delayed, cumulative or ephemeral, and can affect future generations or our own or both. There is little evidence to indicate how these factors are handled. Fuchs (*28*) cited evidence that individual discount rates for financial and physical risk are positively correlated. But the fact remains that benefits and risk may be discounted at different rates. For decisions with very long-term implications, the use of a variable discount rate, declining with time, may more accurately reflect the value given to future risks and benefits than a constant discount rate (*29*). This is an area that seems particularly worthy of attention, for many risk controversies are about risks that are persistent or cumulative, such as carcinogens.

Predicting Risk Controversies

Because of the work to define the factors influencing perceived risk, it is now possible to anticipate the kinds of risk likely to generate controversy. Catastrophic potential and lack of individual control, particularly once an accident occurs or a risk is identified, are apparently the most important risk characteristics. When the uncertainty associated with risks is great, data concerning the uncertainty not forthcoming, and expert opinion apparently divided, apprehension by the public is understandable. Haefele (*30*) termed these risks hypothetical, and described nuclear power as the "pathfinder" for these risks. Certainly there are many risks with the characteristics described above (for example, toxic chemicals and recombinant DNA research). Whether decisions can be made about these risks without the high degree of controversy and the resulting high social cost associated with the nuclear debate remains to be seen.

Quantitative Criteria for Risk Acceptance

In May 1979, the Advisory Committee on Reactor Safeguards (ACRS) recommended "that consideration be given by the Nuclear Regulatory Commission [NRC] to the establishment of quantitative safety goals for overall safety of nuclear power reactors" (*31*). The ACRS further recommended that "Congress be asked to express its views on the suitability of such goals and criteria in relation to other relevant aspects of our technological society. . . ." A similar suggestion, accompanied by proposed criteria, was made by Farmer (*32*) in 1967; the criteria were expressed by a curve relating acceptable accident frequency with accident magnitude. Subsequent proposed criteria for acceptable risks, not necessarily limited to nuclear power, have been made by Starr (*18*), Bowen (*33*), Rowe (*34*), Okrent and Whipple (*35*), Wilson (*36*), and Comar (*37*). Currently efforts are under way within the NRC, the ACRS, and elsewhere to develop quantitative criteria for risk acceptance and to consider the many issues raised by this approach.

Incentives to Develop Quantitative Criteria for Acceptable Risks

The dissatisfaction with current regulatory systems for risk management provides impetus to develop new methods. Theoretically, quantitative criteria for acceptability would resolve many specific criticisms. One criticism stems from the fact that in several cases, a zero-risk goal has been established. This denies the concept of a trade-off between risk and benefit, and ignores the difficulty or impossibility of reaching zero risk. Further, improvements in technology have permitted identification and estimation of risk at levels far below those that were possible when specific zero-risk laws were passed; risks we might consider negligible are not treated in the regulatory process differently from much higher risks. As Hutt (*38*) argued,

> Until quite recently, a no-risk food safety policy was widely thought to be an achievable goal. . . . It is now clear that it is literally impossible to eliminate all carcinogens from our food. Moreover, many of the substances which pose a potential risk are part of long-accepted components of food, and any attempt to prohibit their use would raise the most serious questions both of practicality in implementation and of individual free choice in the marketplace.

A suggested way of handling this problem would be to set a level below which risks would be ignored, provided some benefit were associated with the risk. This low level would serve as a quantitative standard for acceptability of the risk.

A second criticism of regulatory approaches is that decisions are often made arbitrarily. Such a charge is not surprising considering that several regulatory agencies have a mandate to protect the public from "unreasonable" risk, without congressional guidance on how to judge reasonableness. The objections are enhanced when regulators are believed to be overly accommodating or hostile to the regulated industry. Certainly, one way to reduce the influence of bias and arbitrariness is to institute a numerical definition of "reasonable." Perhaps the time required for risk decisions would also be reduced by the availability of clear, relatively simple criteria.

Often, regulatory authorities specify the technology for meeting risk targets, rather than the targets themselves. The drawback of this approach is that there are no incentives to develop more efficient methods of controlling risk. The establishment of risk targets alone could stimulate the development of a variety of creative methods for risk control.

Finally, another criticism of current risk management is that the effort required to control risk (as measured by the cost per life saved) varies considerably from one risk to another; this wastes both lives and money (*9*). Assuming that the total funds allocated for risk reduction could be transferred freely between different risk reduction opportunities (which is certainly not always possible), the maximum number of lives that could be saved nationally is found when the marginal cost of saving a life is uniform among the opportunities. Thus the comparative marginal cost-effectiveness of each opportunity for saving lives would become the guiding principle in the allocation of resources, and the value of life would be implicit in the total national allocation of funds. There would, of course, need to be a national allocation of resources to such "life-saving" endeavors, but as with military budgets, a common-sense consensus judgment is likely to be as reliable as any analytic formula.

Applications for Risk Criteria

One of the pitfalls in trying to develop regulatory approaches for managing risk is the desire to use the same method to tackle a number of different risks. There are different types of risk decisions, and no single regulatory method seems applicable to all of them.

The use of cost-effectiveness criteria serves as an example. This issue arises when, a priori, the technology is found acceptable but the specific operating point is left to be decided. An example of this type of decision is the determination of allowable levels of a pollutant in automobile exhaust. In this case the issue involved is not the relative risk and benefit of transportation, nor the selection of a transportation technology (automobile versus mass transit). For this simplified type of decision the only issue is the marginal trade-off between the social cost of the risk and the cost of controlling it. For these cases, two kinds of quantitative criteria can be considered; the first is the standard for judging cost-effectiveness described above. There is nothing new in this approach; it is simply cost-benefit analysis in which the metric for judging the social cost of risk has been specified. The second quantitative criterion is more pragmatic: it is a lower risk limit below which no regulatory action would be taken. This could be useful in allocating a regulator's time and would help prevent the highly visible cases in which the nuisance aspects of regulation are intuitively greater than the benefits of that regulation.

The next level of difficulty in risk decisions is the choice of the best method for obtaining a specific benefit. In these cases the benefits need not be analyzed, it is presumed that the benefits are sufficiently great to justify any of several alternatives. For example, in the often heated debates over the selection of energy production technologies, it is generally assumed that, under any proposed policy, energy services will be provided (such services include conservation). For this decision, the dominant issues are the costs and risks associated with each alternative. The difficulty in making these choices is often due to the qualitatively dissimilar character of the risks (for example, air pollution risks from coal mining and burning versus nuclear reactor accident risks). It is difficult to see a role for quantitative criteria in making comparisons of the type needed. One could establish a maximum permissible risk level that would serve to screen out excessively risky alternatives, but the selection of a technology generally depends on some aggregation, either explicitly or implicitly, of the components of the social cost of each alternative. Presumably, after one alternative is selected, the decision is reduced to the determination of the preferred operating point discussed above.

A complete risk-benefit decision requires that the relative social cost of the risk be compared with the associated benefit. A pragmatic application of quantitative criteria for these cases was suggested by Starr (*19*) and by Starr *et al.* (*8*), and is illustrated in Fig. 4. This risk-benefit curve reveals the commonly proposed characteristics for risk criteria: a lower limit for concern about risk (in this case, the natural-hazards mortality rate), an upper limit for acceptability (set by the average disease rate), and provision for risk-benefit trade-off between these limits.

We should not overestimate the capacity of simple criteria, such as those illustrated in Fig. 4, to reduce risk conflict costs. Many, if not most, risk estimates include significant judgmental inputs. There are often substantial disagreements over risk estimates, the methods used to arrive at risk estimates, and the competency, integrity, and motivation of the experts providing subjective risk estimates. What is needed for the application of quantitative criteria for risk acceptance is a standard of proof for determining whether the criteria have been met. Although many different approaches to this issue have been recommended (including peer review, scientific courts, and quantitative methods for resolving differences between experts), the ultimate responsibility for judging the competency of risk analysis still resides with the regulatory agency responsible for managing the specific risk.

Conclusion

Analytical approaches to decide risk-benefit issues ideally come closer to maximizing net social benefits than any other approach. The usefulness of these methods in making assumptions and values explicit justifies their application. But a necessary condition for applying their results to specific decisions is a social consensus on the relative benefits and costs of the proposed actions. For specific types of risk, in which intuitive evaluations of risk and benefit contradict analytical evaluations, the necessary consensus may not develop, but rather a conflict requiring political resolution is likely to result.

When the conflict arises from a disagreement over the level of risk rather than the value assigned to that risk, efforts to reduce the cost of conflict by incorporating values into an expected utility approach will be unsuccessful. Quantitative risk criteria appear quite attractive in this respect, because the key to the acceptability of a technology under the proposed method is the level of risk. Assuming that the estimated risk became the central point in the debate, the public might have more confidence in the regulatory systems if their concern were directly addressed.

We see significant value in trying to understand the intuitive risk-benefit process. The evaluation of its outcome could reduce anxiety and cost if used as a tool in the design of technical systems. This is already the case, as when we use more stringent criteria for nuclear power and commercial aviation than for a more commonplace risk (*39, 40*). The balance between individual and group risk-benefit decision methods is fundamental to the development of national policies on risk acceptance. It is customarily achieved through the political process, and is not amenable to quantitative analysis.

References and Notes

1. W. W. Lowrance, *Of Acceptable Risk* (Kaufmann, Los Altos, Calif., 1976).
2. *Symposium/Workshop on Nuclear and Non-Nuclear Energy Systems: Risk Assessment and Governmental Decision Making* (Mitre Corp., Bedford, Mass., 1979).
3. L. Lave, in (*2*), p. 181.
4. S. Jellinek, in (*2*), p. 59.
5. M. Green, "The faked case against regulation," Washington *Post*, 21 January 1979, p. C-5.
6. P. H. Schuck, "On the chicken little school of regulation," Washington *Post*, 28 January 1979, p. C-8.
7. E. J. Mishan, *Cost-Benefit Analysis* (Praeger, New York, 1976).
8. C. Starr, R. Rudman, C. Whipple, *Annu. Rev. Energy* **1**, 629 (1976).
9. R. Zeckhauser, *Public Policy* **23** (No. 4), 419 (1975).
10. R. Wilson and W. J. Jones, *Energy, Ecology, and the Environment* (Academic Press, New York, 1974).
11. K. Arrow, in *Energy and the Environment: A Risk-Benefit Approach*, H. Ashley *et al.*, Eds. (Pergamon, Elmsford, N.Y., 1976), p. 113.
12. National Planning Associates, *Social Decision-Making for High Consequence, Low Probability Occurrences* (PB-292735, National Technical Information Service, Springfield, Va., 1978).
13. R. Howard, in *Societal Risk Assessment: How Safe is Safe Enough?*, R. C. Schwing and W. A. Albers, Eds. (Plenum, New York, in press).
14. W. Haefele, in (*11*), p. 141.
15. H. Otway, *J. Br. Nucl. Energy Soc.* **16**, 327 (1977).
16. S. Montague and E. Beardsworth, unpublished manuscript.
17. B. Fischhoff, P. Slovic, S. Lichtenstein, S. Read, B. Combs, *Policy Sci.* **9**, 127 (1978).
18. C. Starr, *Science* **165**, 1232 (1969).
19. ——, in *Perspectives on Benefit-Risk Decision Making* (National Academy of Engineering, Washington, D.C., 1972), p. 17.
20. P. Slovic, B. Fischhoff, S. Lichtenstein, in *Societal Risk Assessment: How Safe is Safe Enough?*, R. C. Schwing and W. A. Albers, Eds. (Plenum, New York, in press).
21. E. W. Lawless, *Technology and Social Shock* (Rutgers Univ. Press, New Brunswick, N.J., 1977).
22. I. Janis and L. Mann, *Decision Making* (Free Press, New York, 1977).
23. See, for example, A. Tversky and D. Kahneman, *Science* **185**, 1124 (1974).
24. J. Selvidge, thesis, Harvard University (1972).
25. S. Lichtenstein, P. Slovic, B. Fischhoff, M. Layman, B. Combs, *J. Exp. Psychol. Hum. Learn. Mem.* **4**, 551 (1978).
26. P. Slovic, B. Fischhoff, S. Lichtenstein, *Accid. Anal. Prev.* **10**, 281 (1978).
27. M. F. Jacobson, *Science* **207**, 258 (1980).
28. V. Fuchs, "The economics of health in a post-industrial society," *Public Interest No. 56* (1979), p. 3.
29. J. M. English, in *Trends in Financial Decision Making*, C. van Dam, Ed. (Nijhoff, The Hague, Holland, 1978), pp. 229–247.
30. W. Haefele, *Minerva* **10**, 302 (1974).
31. M. W. Carbon, letter to J. M. Hendrie, 16 May 1979.
32. F. R. Farmer, *Nucl. Saf.* **8**, 539 (1967).
33. J. Bowen, in *Risk-Benefit Methodology and Application: Some Papers Presented at the Engineering Foundation Workshop, September 22–26, 1975, Asilomar, California*, D. Okrent, Ed. (Rep. PB-261920, School of Engineering and Applied Science, University of California, Los Angeles, 1975), pp. 581–590.
34. W. D. Rowe, *An Anatomy of Risk* (Wiley, New York, 1977).
35. D. Okrent and C. Whipple, *An Approach to Societal Risk Acceptance Criteria and Risk Management* (Rep. UCLA/ENG-7746, School of Engineering and Applied Science, University of California, Los Angeles, 1977).
36. R. Wilson, testimony before Occupational Safety and Health Administration hearings, Washington, D.C., February 1978, OSHA docket H-090).
37. C. L. Comar, *Science* **203**, 319 (1979).
38. P. B. Hutt, *Food Drug Cosmet. Law J.* **33**, 558 (1978).
39. C. Sinclair, P. Marstrand, P. Newick, *Innovation and Human Risk* (Centre for the Study of Industrial Innovation, London, 1972).
40. R. Schwing, *Gen. Mot. Res. Lab. Rep. GMR-2353* (1977).