DECISION APPROACHES TO RISK CRITERIA

S. L. Derby
Stanford University

B. Fischhoff, S. Lichtenstein, and Paul Slovic
Decision Research, A Branch of Perceptronics

R. L. Keeney
Woodward-Clyde Consultants

INTRODUCTION

Many approaches are advanced for finding acceptable risk criteria. Each has advantages and disadvantages, supporters and detractors. This paper discusses only the use and limitations of decision analysis for determining acceptable risk. Hopefully, this information can provide a better basis for comparing the value of decision analysis with other approaches.

Decision analysis has much to recommend its use. For example, it is, in principle, open to public scrutiny, flexible enough to suit many problems, and supported by a sophisticated methodology. However, when used in isolation, it has a number of technical, conceptual and political weaknesses threatening its usefulness. For example, it may be hard to assess the robustness of its conclusions, to accommodate hard-to-measure effects, or to guarantee meaningful public involvement and scrutiny.

This paper presents a definition of acceptable risk and explains how decision analysis could be used to determine "How *safe* is safe enough?" Practical limitations to the potential of decision analysis are then summarized.

ACCEPTABLE RISK IS A DECISION PROBLEM

An important conceptual clarification comes with realizing that acceptable risk problems are decision problems; that is, they require a choice between alternatives. What distinguishes acceptable risk problems is that all include threats to the life and health of humans, animals, or plants among their consequences and that almost all require very difficult choices among their decision options.

Whether done formally or informally, examination of the alternatives in a decision problem involves the following five interdependent steps:

1. Specifying the objectives (i.e., what one wants),

2. Defining the possible alternatives (including "do nothing"),

3. Identifying the possible consequences of each alternative (including, but not restricted to, risks),

4. Specifying the desirability of the various consequences and the likelihood of their being achieved, and

5. Analyzing the alternatives and selecting the best one.

This final step is prescriptive, in the sense that it prescribes the option that should be selected (given the logic of the analysis); as such, it identifies the <u>most acceptable</u> risk. In addition, if the recommendation of the last step is implemented and the best alternative is adopted, the risks of that alternative are <u>accepted</u>.

It is clear that there are no universally acceptable risks. At the same point in time, different people, using different decision-making methods, defining different alternatives, or having different values, beliefs, or objectives, might hold quite different views on which risks are acceptable.

Similarly, acceptable risk is not a static notion. As the world changes, people's values may change, new safety devices may be invented, additional information about the risks may come to light, and so forth. Any such change may mean that a risk that is acceptable today will not be acceptable tomorrow (or vice versa).

Conceptualizing acceptable risk problems as decision problems is central to the definition of "acceptable risk." It is essential because the outcome of the decision process _is_ the acceptable risk. It is important because whether one believes that none of the available options are acceptable, or many of them are, one should still prefer to choose the most acceptable. This risk may not make one happy; it is, rather, the best option one can now find.

FLAWS IN SIMPLISTIC SOLUTIONS

Viewing acceptable risk as a decision problem helps illuminate the flaws in some simplistic solutions of hazard management. For example, one tempting strategy is to claim that no risk should be tolerated. The decision perspective forces one to ask in response "What is the cost of absolute safety?" Abhorrence of risk can lead to dubious decisions, thereby incurring great cost for a minor reduction of risk. Rather than paying for safety, one might propose doing without the substance or activity or technology in question. A decision perspective requires one to ask what alternative can be chosen in its stead. If an alternative with any risk is substituted, then the gain in safety may prove illusory. If individuals troubled by obesity or diabetes have a need for sweeteners, banning saccharin will eliminate one possible cancer risk in return for increased risks from the consumption of sugar.

A variant to the desire for absolute safety is the unqualified suggestion that the chosen alternative be as safe as possible. An option with less risk may be available only through a large increase in cost. Most people would tolerate some small increase in risk for a large reduction in cost (at least if those bearing the risk also received the cost savings).

Another simplification calls for expressing the answer to "How safe is safe enough?" by a small number (like 10^{-7}), representing the probability that an individual will be killed *in a* particular year due to the cause under consideration. The following example illustrates one situation where this solution appears inappropriate. Suppose that two alternatives lie just on opposite sides of the designated standard, and that the "acceptable" one costs substantially more. In practice, the less costly alternative would be preferred by many people, despite being above the safety standard.

A more sophisticated solution is to specify fixed tradeoffs between cost and risk. For example, one could adopt any safety measure costing less than one million dollars per expected life saved. This, too, can be an oversimplification. For example, if an alternative is too risky, such a standard may require a more costly option than actually necessary for acceptance. The tradeoff may depend on the overall cost, as well as, how much safety can be purchased for increased cost.

In summary, the answer to the question "How safe is safe enough?" means choosing between alternatives. The risk associated with the chosen alternative is, by this definition, the most acceptable risk. However, decision makers using different alternatives, believing different information, and having different value judgments can arrive at conflicting conclusions on what risks are acceptable.

DECISION ANALYSIS

Decision analysis has its origins in the theory of individual decision making developed by von Neumann and Morgenstern[1] and Savage[2]. Decision theory is an axiomatized, prescriptive theory for addressing decision situations involving uncertainty; if you accept the axioms and their interpretation in practice, you ought to make choices as the theory prescribes. The key elements of decision theory are uncertainties, expressed as probabilities, and values, expressed as utilities.

Decision analysis combines decision theory with techniques for modeling complex problems, many of which were originally developed in operations research and management science. In the past two decades, this marriage of axiomatic theory and applied methodology has evolved into a broadly applicable form of analysis[3,4,5,6,7].

A thorough decision analysis has five main steps:

(1) Structuring the problem. In this stage, the analyst defines the decision problem, identifies the alternatives and their possible outcomes, selects the variables suitable to the problem, and builds a structural model. These tasks involve artfulness and common sense; the goals of the structuring and many of the tools available to the analyst are the same as for cost-benefit analysis. In this structuring, decision analysis

can accommodate any consideration that the decision maker deems appropriate. Thus, in principle, "soft" values like aesthetics or "satisfying Senator X" can be included as easily as "hard" values like monetary costs.

(2) Assessing probabilities. The model developed in Step 1 typically contains many uncertainties about the present and future state of the world. In Step 2, these uncertainties are quantified as probabilities. Decision analysts take a subjectivist view of probabilities: probabilities are not characteristics of things, but are expressions of individuals' beliefs. Thus, the probabilities needed for the decision-analysis model are elicited from the decision maker or from professionals and experts.

(3) Assessing preferences. Whereas cost-benefit analysis seeks the numerical expression of preferences in market data, decision analysis uses subjective value judgments. Monetary data can, however, enter a decision analysis if the decision maker finds them relevant.

When alternatives have several consequences, cost-benefit analysis simply adds together the various costs and benefits. Although such additive, linear models are also used in decision analysis, other combination rules are available. For example, a multiplicative rule may be appropriate when the utility of one consequence depends on the level of another.[7]

Decision analysis also addresses attitudes toward risk. Risk aversion means that the desirability of an alternative with uncertain outcomes is less than the desirability of its expected value (i.e., its outcomes weighted by their probabilities). Thus, for example, an analysis could reflect the decision maker's belief that a safety device having a .5 chance of saving 100 lives is less desirable than one that will surely save 50 lives. Risk proneness is the reverse.

(4) Expected utility. After all the parts of the model have been assessed, the attractiveness of each alternative is summarized in its expected utility. The alternative with the greatest expected utility is the indicated choice.

(5) Sensitivity analysis and value of information. Throughout the analysis, the indicated choice is reexamined from two perspectives: (a) Can the analysis be simplified by omitting components that do not affect the final outcome? For example, if

one alternative were inferior to another in all aspects, it could be dropped. (b) Are there critical places in the analysis where a reasonable change in the structure, a value, or a probability could lead to the selection of a different alternative? Two tools have been developed for these reexaminations. Sensitivity analysis examines the vulnerability of the decision analysis to changes in its parts. Value-of-information analyses assess the value of gathering further information to clarify the critical components of the decision. These two techniques can clarify the important contributions to a decision and concentrate resources on obtaining meaningful information. Most important, value of information analyses can prevent meaningless data gathering and experimentation.

USING DECISION ANALYSIS TO SET ACCEPTABLE RISKS

If an acceptable risk problem receives a thorough decision analysis, the alternative with the greatest expected utility is an alternative whose risk is acceptable to the decision maker. A problem arises, however, when a single decision maker cannot be identified. Since the key elements of a decision analysis (probabilities, utilities, problem structure) are all subjective, they must come from someone. However, in societal risk decisions, there is rarely one individual whose values and probabilities are the only appropriate ones to use in the analysis. Several approaches are possible when more than one set of utility or probability inputs seem relevant.

At times, it may turn out that the values and probabilities of the various parties lead to the same ordering of the available alternatives. Gardner and Edwards[8] found that when two opposing groups, realtors and conservationists, used only their own intuitions for ranking alternative solutions to a coastal zoning problem, they were in strong disagreement. However, when their rankings were generated by a simplified form of decision analysis, the disagreements generally disappeared.

In some situations, however, the use of formal analysis may lead to polarization, rather than convergence of views. The act of specifying one's values and beliefs may harden one's commitment to them. Compromise may be more difficult when the opposing parties' views are "on record". Leaders may assume extreme positions to ensure followers' allegiance. Finally, as constituent groups gain experience with formal analysis, they may exaggerate their positions in order to bias the analysis in their favor or gain advantage in later stages of conflict resolution.

When the parties are willing to work together, it may be possible to generate agreement on the judgments needed. Such agreement might come from compromise (I'll give up here if you give up there; put it to a vote; let's take an average), or from the gradual formation of genuine consensus. Keeney and Raiffa[7] have provided a justification for such solutions by positing the existence of a hypothetical person, the Supra Decision Maker, whose values and probabilities are those of the group consensus.

Even if the various parties cannot arrive at agreement, Keeney and Raiffa[7] recommend selecting a real or hypothetical Supra Decision Maker for the analysis. The Supra Decision Maker can then use theoretically justified techniques for incorporating the probabilistic judgments of various experts into his or her own beliefs.[9] As for integrating different values, regulators and public policy makers already view their job as just that, trying to reflect the values of the entire society they serve. One could formalize that process by eliciting the values of various stakeholders (environmentalists, politicians, manufacturers, impactees, etc.) in a consistent fashion, letting the Supra Decision Maker determine the relative weight to place on each.[10]

Finally, if all else fails, the analyst can prepare several complete decision analyses, each one incorporating the views of one of the disagreeing parties. The analyses are clearly only aids in establishing acceptable risk; other procedures are needed to resolve the disagreements.

PRACTICAL LIMITATIONS

Problem Definition

Because of resource constraints, a formal model cannot include everything. It must simplify and omit. Acceptable risk problems have frequently been defined too narrowly, acknowledging only technological options and economic consequences.

Although critics have typically complained about inadequate scope, breadth also holds dangers. An analysis may become so large as to be unwieldy and unworkable, the structure so complex that its interrelationships cannot be understood, the needed inputs too numerous to measure adequately. Time pressures may justify deliberate omission. For example, a government regulatory body receiving a flurry of complaints about severe side-effects from a recently licensed drug may, in order to respond quickly to a potentially dangerous situation, choose to omit from their analysis otherwise important considerations such as the effect of a recall on pharmaceutical innovation. However, a persistent narrowing of focus leaves larger issues perpetually unaddressed.

There is no inherently correct problem definition. Since technological risk problems change over time, their problem definitions keep changing. New alternatives are added, additional objectives become relevant, and new measures for achieving objectives are devised. However, these uncertainties about problem structure are not modeled themselves. The structure of any analysis is fixed, at least until one iteration is complete. Changing the structure is expensive and time consuming, and in practice is rarely done.

Knowing the Facts

One reason for using decision analysis techniques is to organize the relevant facts effectively. Ideally, the analyst seeks to assemble all the significant events and their consequences. Unfortunately, many of the necessary factual data are hard to come by. Despite the enormous scientific progress of recent decades, we still do not understand all of the possible physical, biological, psychological, and social consequences of any large technological project. No consideration of the facts of an analysis is complete without a consideration of the uncertainty surrounding them.

Despite an appearance of objectivity, risk assessments are inherently subjective. Precise statistical data are rarely available. Most often, especially with new substances and technologies, risks must be estimated by expert judgment or extrapolations from data on related mechanical or biological systems. Probabilistic models are substituted for direct experience. Although sophisticated, such models include a large component of human judgment. The structure of the problem, the consequences to be considered, and the probabilities of consequences are all determined to some degree by educated intuition.

In sum, the objectivity of "the facts" is an illusion. For most important analyses, we do not know and perhaps cannot make estimates of present and future states with the desired degree of precision. Experts are forced to go beyond (even far beyond) the available data and rely on their educated intuitions.

Assessing Values

The strength of decision analysis is that many value issues are given explicit quantitative expression and are placed at the focus of societal decision making. In addition, conflicts among different members of society can be clarified by specifying which issues underlie people's disagreements and whether such disagreements make any difference in the final decision.

The attention that decision analysis places on values has increased awareness of a number of troublesome issues in measuring values.

A central problem in measuring values is that our preferences may be neither stable over time nor well articulated at any moment.

Stating labile values in numerical form may lead people to believe they are solid, reliable, and unchanging, when they are in fact diffuse, uncertain, and changing rapidly.

Decision Quality

Analysts normally realize that the inputs to their analyses may be inaccurate. The standard practice for acknowledging and accommodating this uncertainty is sensitivity analysis. To the extent that these reanalyses produce similar results, the analyst gains confidence in the quality of the analysis.

However, there are no established guidelines for identifying which inputs might be in error or for choosing the appropriate range of possible values to be tested. Further, sensitivity analyses often cannot identify when some analytical procedure introduces a bias in all aspects of an analysis in which it is used. For example, a value-elicitation method may persistently invoke but one perspective on respondents' labile values.

Whether or not formal procedures like sensitivity analysis are used, determining the quality of an analysis is a matter of judgment. The decision maker must judge which inputs are dubious and how sound the problem statement is. Essentially, that person must decide how good his or her own best judgment is.

SUMMARY

The great strengths of decision analysis are its openness and conceptual soundness. In some sense, this thoroughness is also its downfall, for it makes any failings and assets visible and documented. By detailing every step of their work, from problem definition through value and fact assessment to final selection of an alternative, good analysts maximize the possibility for peer review and political review. Presently, the main weakness of decision analysis is the failure to realize this potential for practical applications.

Decision analysis appeals to regulators because it appears to be a value-free, even objective, guide to decision making for risk criteria. Yet, like any other decision process, it mixes issues of fact and value in complex and subtle ways.

Given the uncertainty surrounding both the facts and values in acceptable risk decisions, the process must be an iterative one. As time goes on, we learn more about how a hazard behaves and how much we like or dislike its consequences. Indeed, a measure of the success of an analysis is its ability to inform our beliefs and values as well as to reflect them.

REFERENCES

1. von Neumann, J. & Morgenstern, O., Theory of Games and Economic Behavior, Princeton, N.J.: Princeton University Press, 1947.

2. Savage, L. J., The Foundations of Statistics. New York: Wiley, 1954.

3. Howard, R. A., The Foundations of Decision Analysis. IEEE Transactions on Systems, Science and Cybernetics, SCC-4(3), Sept. 1968, 393-401.

4. Readings in Decision Analysis. Decision Analysis Group, SRI International, Menlo Park, California, 1976.

5. Raiffa, H., Decision Analysis. Reading, Mass.: Addison-Wesley, 1968.

6. Schlaifer, R. O., Analysis of Decisions Under Uncertainty. New York: McGraw-Hill, 1969.

7. Keeney, R. L. & Raiffa, H., Decisions with Multiple Objectives: Preferences and Value Tradeoffs. New York: John Wiley, 1976.

8. Gardner, P. J. & Edwards, W., Public Values: Multiattribute-Utility Measurement for Social Decision Making. In M. F. Kaplan & S. Schwartz (Eds.), Human Judgment and Decision Processes. New York: Academic Press, 1975.

9. Morris, P.A., Decision Analysis Expert Use. Management Science, 1974.

10. von Winterfeldt, D., A Decision Theoretic Model for Standard Setting and Regulation. RM-78-7. Laxenburg, Austria: International Institute for Applied Systems Analysis, February 1978.