ES/ER/TM-134

Decision Support for CERCLA Investigations: An Introduction to Decision Analysis Applications

S. T. Purucker
B. F. Lyon
R. N. Stewart
L. D. Nanstad

ES/ER/TM-134

Decision Support for CERCLA Investigations: An Introduction to Decision Analysis Applications

S. T. Purucker
B. F. Lyon
R. N. Stewart
L. D. Nanstad

Date Issued--September 1994

Prepared by
The University of Tennessee
Knoxville, Tennessee
under direction from the
Environmental Restoration Risk Assessment Council

Prepared for
U.S. Department of Energy
Office of Environmental Restoration and Waste Management
under budget and reporting code EW 20

OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831-6285
managed by
MARTIN MARIETTA ENERGY SYSTEMS, INC.
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-84OR21400

CONTENTS

TABLES

1. INTRODUCTION

The overall objective of this document is to provide the Oak Ridge Operations-Environmental Restoration (ORO-ER) technical community with an introduction to various decision analysis applications and their relevance towards the Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA) process. The long-term goal of investigating the decision analysis literature is to find specific applications that are useful in the collection of data and the selection of alternatives. A secondary goal is to use these techniques to facilitate technical presentations in a manner that produces results which are responsive to the needs and preferences of decision-makers. This will allow the decision-makers to better use the data collected for assessment efforts. These methods will supplement the current implementation of activities such as programmatic prioritization, the Data Quality Objective (DQO) process, and the Remedial Investigation (RI)/Feasibility Study (FS) process.

Environmental decision-making is perceived by many as a process that poses unique problems at levels of complexity and uncertainty that can overwhelm attempts to treat the resulting decisions analytically. Any attempts to reduce the remediation problem to an analytical decision problem are often viewed as involving a morass of mathematics that can cast suspicion on the results. Although there is some validity to this viewpoint, currently there is not enough emphasis on using analytical tools for assisting the decision-making process. Experience from the application of decision techniques in other areas faced with problems of similar scope and uncertainties can provide beneficial effects for decision processes in conducting remediation efforts at the ORO-ER.

Quantitative decision methods can be developed during environmental restoration projects that incorporate stakeholder input and can complement current efforts that are undertaken for data collection and alternatives evaluation during the CERCLA process. These decision-making tools can supplement current United States Environmental Protection Agency (EPA) guidance as well as focus on problems that arise as attempts are made to make informed decisions regarding remedial alternative selection. In this document, background information is provided on decision theory and decision tools that are available for determining environmental solutions which can supplement the existing RI/FS decision framework. In examining the use of such applications, the use of decision analysis tools and their impact on collecting data and making environmental decisions is discussed. This is done by using a broad risk-based conceptual formulation of the CERCLA decision problem and attempts to address the uncertainties inherent in the decision process while allowing for stakeholder input. We will look at the construction of objective functions for quantifying different risk-based decision rules that incorporate stakeholder concerns. Objective functions represent a quantitative method for implementing the DQO process. Based on such defined objective functions, a project can evaluate the impact of different risk and decision selection strategies on data worth and alternative selection. Included are general methods for integrating results of risk analyses with other pertinent information so that an informed decision can be reached for a particular site.

2. THE ROLE OF DECISION ANALYSIS IN ENVIRONMENTAL REMEDIATION

Decision analysis implements decision theory with the aid of concepts from management science, operations research, and economics and differs from the decision theory through the addition of these applied methodologies. These additions make the methods more generally applicable while eliminating their axiomatic qualities. In decision theory, if one accepts the premises and axioms in practice, one *should* make the recommended choices. Decision analysis is less prescriptive.

Definitions of decision analysis are largely a function of convenience:

- "provides the link between the economic framework in which decisions are made and the results of the technical analyses on which decisions are based" (Freeze et al. 1990)

- provides useful ways of analyzing the selection of a choice of action under conditions of uncertainty (Mansfield 1980)

- includes a number of principles and techniques for the systematic study of decision-making with uncertain conditions (Baker and Kropp 1985)

- "a formalization of common sense for decision problems which are too complex for informal use of common sense" (Keeney 1982)

Since its inception, the development of the RI/FS process has used the fundamental concepts of decision analysis to construct a sound framework for environmental decision-making. The goal of the RI/FS is ". . . to gather information sufficient to support an informed risk management decision regarding which remedy appears to be most appropriate for a given site...by identifying and characterizing hazards in a way that will contribute directly to the selection of an appropriate remedy" (EPA 1989). The approach is designed to be a dynamic, flexible process that can and should be tailored to specific circumstances of individual sites; it is not intended to be a rigid step-by-step approach that must be conducted identically at every site. However, all too often, the RI (site characterization) and FS (remedial alternative evaluation) become distinct processes where data analysis and alternative evaluation processes are conducted in manners that fulfill their own independent needs, rather than contributing to an overall technical understanding of the site that can be communicated to the appropriate decision-makers. This is contrary to specific EPA guidance (EPA 1988): "It is important to note that the RI and FS are to be conducted concurrently and that data collected in the RI influence the development of remedial alternatives in the FS..." The end result can be uninformed decisions that are based on individual subjective judgement and/or political processes rather than on the technical understanding that has been or could be developed through the RI/FS process. The RI/FS process is described in various EPA guidance documents and is the methodology that the Superfund program has established for characterizing the nature and extent of environmental risks posed by hazardous waste sites and for evaluating potential remedial options.

The central challenge faced by those implementing the RI/FS process is to determine how best to use the flexibility built into the process to conduct timely, effective, and efficient cleanups. A significant barrier to conducting these investigations is the inherent uncertainties associated with the characterization and remediation of hazardous waste sites. While these uncertainties can foster a scientific impulse to want to know more, this desire competes with the overall objective of performing efficient, timely cleanups. The objective of the RI/FS process must be kept in mind throughout the process; moreover, the objective is not the unobtainable goal of removing all uncertainty but rather to gain information sufficient to support an informed risk management decision concerning which remedial option appears to be the most appropriate for a given site. The appropriate level of analysis needed to meet this objective can only be reached through constant strategic thinking and careful planning concerning the essential data needed to reach a remedy selection decision.

As hypotheses are tested and either rejected or confirmed, adjustments or choices as to the appropriate course for further investigations and analyses are required. These choices, like the remedy selection process itself, involve balancing a wide variety of factors, including technical, practical, and economic considerations. Broadly speaking, there are nine evaluation criteria in the FS process that, to some extent, are used to make the following decisions:

- Is action necessary?
- Is active action necessary? (Is institutional control enough?)
- Which alternative to select?
- Are more data needed, and if so how much?

These are the same questions that are addressed in any decision analysis process. A thorough decision analysis is usually composed of five steps (Fischoff et al. 1981). Table 1 outlines some of the general parallels between the CERCLA process and these steps.

Table 1. Decision analysis and CERCLA

Decision Step	**CERCLA Step**	**Step Description**
1. Structure the Problem	Data Quality Objectives	Define the decision problem by specifying the technical objectives and identifying data needs.
2. Assessing Probabilities	Remedial Investigation	Quantify the uncertainties about the present and future states of the system to be studied through data collection, assessment, and judgement.
3. Assessing Preferences	Risk Management	Consider subjective value judgement and attitude towards risk.
4. Evaluating Alternatives	Feasibility Study	Evaluate potential remedial alternatives against relevant criteria.
5. Sensitivity & Information Value	Data Quality Assessment	Reexamine analysis to determine sufficiency of current information and value of additional data.

3. CONSTRUCTION OF THE OBJECTIVE FUNCTION

Selecting a course of action is often a subjective process that leads to conflict between different stakeholders because different decision-makers sometimes have fundamental value conflicts that prevent them from reaching a consensus. Constructing different objective functions that account for different values can aid in distinguishing where conflicts really exist as opposed to where they simply appear to exist and can also serve as a forum for resolving disagreements. In addition, assuming the decision-maker is able to quantify the amount of allowable uncertainty in the selection process, definition of objective functions can be used to demonstrate when current knowledge is sufficient to support an informed decision.

The purpose of risk analysis activities under CERCLA is to provide information to aid the evaluation of proposed remediation alternatives. The alternatives to be evaluated are often explicitly selected to meet a range of cleanup criteria, land use options, and compliance points. The cleanup criteria may be based on an acceptable human cancer risk level, a level at which no toxic effects to human health is expected, or they may be based on levels that are expected to have no ecological impacts. These criteria may be developed under current land use and reasonable future uses, such as industrial or residential use, and may be evaluated on-site or at a number of other off-site compliance points.

The criteria used for evaluating alternatives in the FS are described in detail in *Guidance for Conducting Remedial Investigations and Feasibility Studies Under CERCLA* (EPA 1988). The first two criteria, protection of human health and the environment and compliance with federal and state regulations, are termed threshold criteria and must be satisfied by all proposed remedial action technologies. All alternatives that pass the primary criteria are weighed against the next five criteria: long-term effectiveness; short-term effectiveness; reduction of toxicity, mobility, and volume; implementability; and cost, so that decision-makers will be able to select the appropriate site remedy. Currently, the final two criteria, state and community acceptance, are then addressed after comments on the RI/FS report have been received. These criteria form the basis of decision rules for alternative selection under CERCLA.

A clear decision rule (or rules) around which quantitative work can be developed is necessary to allow an explicit application of mathematical logic to decision-making. This usually involves defining an objective function that represents the problem to be solved. Objective functions use different indices and are the primary functions used to evaluate environmental problems with respect to a variety of RI/FS criteria. Constructing an objective function can help distinguish where true conflicts exist and serve as a forum for resolving disagreements. In essence, the objective function serves as a measure for comparing various alternatives for remediation. They are usually based on cost (payoff or loss), but other forms of measurement are possible as well. Broadly speaking, the objective function can incorporate three factors:

- Understanding of the problem via predictive models and the data necessary to use these models
- Assessment of the costs involved (this can include risk-aversion; e.g., some costs are unacceptable)
- Preference of decision-makers (relative weighting of which criteria are important to the decision)

The objective function should comprise components that represent important facets of the problem and should be constructed in an interpretable fashion. The objective function is usually constructed so that the optimal alternative will produce a minimum or maximum value. Common elements can be found in various types of objective functions, particularly in the environmental field. Such components include:

- Time horizon—the amount of time which the project will continue, usually measured in years. The time factor is important for long-term projects as certain quantities change over time (e.g., value of the dollar). These quantities need time-scaling factors to estimate a meaningful objective function.

- Alternative cost—the cost of an alternative in a given year. This includes personnel required to operate the technology, maintenance, etc.

- Resource value—the equivalent value of a resource to be protected (e.g., aquifer, endangered species, a human life). Assigning a measurement to such intangibles is often an unpleasant but unavoidable task in constructing meaningful objective functions.

- Utility—the concept of decision-makers weighting factors to better reflect the relative desirability of various outcomes derives from utility theory (discussed in subsequent sections of this report). This is a highly subjective assignment that is unavoidable in many situations. However, there are utility models available to assist in quantifying the utility decision.

- Efficiency—the degree to which an alternative can reduce some parameter of interest (e.g., risk, contaminant concentration, contaminant flux).

- Probability of failure—this is the probability that the alternative chosen will fail to achieve the goal defined by the decision criteria. This value can be numerically estimated as in the output of a probabilistic risk assessment or a subjective estimate based on previous experience.

- Cost of failure—any costs associated with the failure of the chosen alternative to remediate the process. This may include litigation or cost of a second remediation. It may also encompass the loss of resource values discussed previously.

- Benefits—the benefit of implementing a given alternative (e.g., profit). For many environmental remediation scenarios, this is zero.

These components may be derived from secondary objective functions themselves. CERCLA sites usually address these different functions (except for cost and risk) in a qualitative manner when developing the detailed analysis of alternatives in the FS. They are all used, either implicitly or explicitly, in the nine FS criteria upon which decisions are based. However, since these criteria are often not quantitatively identified early in the project, this can result in core problems not being clearly identified and technically resolved. Quantitative evaluation of the more contentious objective functions identified by the decision-makers can assist in alternative selection and ensure that data collection efforts are designed to meet the needs of those decision-makers.

A very general example of an objective function based on a risk/cost/benefit strategy (addressed in the next section) is constructed as

$$\Phi_j = \sum_{t=0}^{T} \frac{1}{(1+i)^t} \left[B_j(t) - C_j(t) - R_j(t) \right]$$

where

t	=	the number of years
$B_j(t)$	=	the benefit of alternative j in year t
$C_j(t)$	=	the cost of alternative j in year t
$R_j(t)$	=	the probabilistic cost of failure for alternative j in year t
T	=	the number of years
Φ_j	=	the objective function for the jth alternative.

Therefore, for each year, one determines the benefits, costs, and risks and scales their differences. The sum of those values is the objective function value for the jth alternative. If the decision-maker was the owner of a new waste management facility, C(t) represents the capital and operational costs associated with the facility. The benefits B(t) represents the profit of running the facility. R(t) represents the probabilistic cost associated with the failure of the facility.

An example can be found in a Finkel and Evans report (1987), where three alternatives were considered for the remediation of a hazardous waste site which poses a potential risk to human health. Each method has varying performance levels and costs of implementation. The objective function measured in dollars is the "total societal cost." In this example, one would choose the alternative that minimizes the objective function. The objective function is constructed as

$$\Phi(a;R) = C_a + v(1 - e_a)R,$$

where

C_a	=	cost of implementing the alternative,
e_a	=	efficiency of alternative a,
v	=	cost of failure (million $/cancer), and
R	=	health risk in the absence of any controls (cancers).

Although one may not be comfortable explicitly addressing the costs of failure in this manner (i.e., assigning a dollar value to human life), it is always present, at least implicitly, in any decision rule if any sort of risk is to be tolerated. Only the implementation of a Zero Risk strategy can avoid this issue.

If each component of an objective function were known with certainty, then the decision would be completely analytical. Typically, at least one of the components is not known with complete certainty thereby introducing uncertainty into the objective function and ultimately into the decision. Issues in objective function analysis are discussed in the next section.

The objective function concept has applications in various decision scenarios and construction of the function can have various forms based on the previously described decision elements and simple mathematical operations. The careful construction of an objective function that represents the problem and alternatives available is an important step in the process. It can help identify areas of uncertainty and disagreement and provides a foundation for other important concepts (e.g., data worth).

4. DECISION MAKING AND OBJECTIVE FUNCTION ANALYSIS

A central feature of objective functions for environmental cleanups is risk, usually represented by the results of a risk assessment or of specification of the probability of failure. Decision-makers must contend with the fact that their decisions cannot be made with certainty for the majority of environmental problems and that the analysis of risk must reflect this uncertainty. For our purposes, the term "risk" is used in a broader fashion than that associated with human health and ecological risk assessments for CERCLA activities. Risk can be generally defined as the consequences of making a "wrong" decision relative to different criteria (e.g., engineering, human health, ecological, economical, regulatory concerns, as well as others) and is measured as the probability and associated costs of an adverse effect. These risks arise from the need to satisfy a number of conflicting technical objectives (and their associated uncertainties) as well as the need to satisfy additional legal and political constraints. The nature of the risk strategy chosen has an influence on how the objective function is constructed to represent the decision. A number of different strategies are available to attempt to mitigate risks (Crouch and Wilson 1982), including *Zero Risk, As Low As Reasonably Achievable (ALARA),* and *Best Available Control Technology (BACT).*

In *Zero Risk*, any alternative that involves an element of risk is rejected. Although perhaps the ideal strategy to implement, it soon becomes clear that it is impossible to carry out in practice. The objective function for this strategy would be based on eliminating the risk element and ignoring the other possible elements of an objective function. However, every action, including no-action, involves some element of risk and this strategy is rarely tenable. In *ALARA*, risks are mitigated to levels that are as low as reasonably achievable based on a decision rule that specifies what is reasonable. However, although workable, the open-ended nature of these decision rules often lead to inconsistent applications of risk reduction. In addition, they are sometimes inefficient in reducing risk since a broader objective function is often able to effect a similar reduction in risk while optimizing other variables (e.g., cost). In *BACT*, the best available technology is used, and an analysis involving an objective function is usually not required. However, identifying this technology is often difficult. *BACT* is usually interpreted to consist of a tested and commercially available design which can be implemented at a reasonable cost. This strategy can also be inefficient in risk reduction from a cost basis, but was nonetheless incorporated into the Clean Air Acts of the 1970s.

The preferred framework in EPA guidance is the *Risk/Cost/Benefit Analysis* approach. This method requires an explicit identification of the values assigned to various risks. This is recommended not because it will give the optimal decision every time, since no method can, but because it requires that all values which influence the decision are explicitly recognized and discussed. Within this framework, the general objective function (or similar variants) introduced in the previous section play a key role in measuring outcomes against decision criteria.

Once an objective function has been determined, it may be approached in either a deterministic or probabilistic manner. In a deterministic approach, all quantities are assumed to be known with certainty and the choice is clear within that decision framework. However, rarely are all the components of a problem known exactly. In a probabilistic approach, likelihoods are assigned to values of components or outcomes and an expected value for the objective function is used. This allows flexibility in the decision framework and leads to a broader analysis of capabilities, including sensitivity issues and data worth.

4.1 DETERMINISTIC METHODS

A large part of formal theory in the social sciences and especially in economics and management sciences is related to the field of decision-making under certainty. This field typically reduces the decision problem down to a set of alternatives from which one (or more than one) must be selected that maximize (or minimize) some given objective function. This leaves the choice of the objective function as the core of the problem. Environmental problems have an economic context and therefore profit and loss are often suitable. But as part of the CERCLA process, the index used can take any form of the objective functions described previously.

To illustrate some general principles for decision-making under certainty, we will use an example consisting of a site with contaminated soil and two possible outcomes (future land uses): residential or industrial. The scenario that occurs controls the costs of implementing the necessary cleanup. The decision-maker has also estimated how well each action would do in these two possible outcomes. These costs also include possible penalties for not completely cleaning up the site initially if residential land use later arises. Assume that the decision-maker has to decide between three remedial alternatives regarding the cleanup of the contaminated site. Alternative 1 involves minimal or no action. Alternative 2 involves intensive cleanup. Alternative 3 is an intermediate cleanup. For this discussion, we neglect the likelihood of each of the outcomes.

In a relatively perfect world, the success of each strategy (called the payoff) could be computed with certainty. In that case, the strategy to pick is the one with the highest payoff (or in this case, the smallest loss). This would be a fairly easy decision to make. Of course, nothing is this simple in reality, especially when decisions regarding environmental problems, remedial actions, and their inherent uncertainty are concerned. The outcomes, which are, in fact, uncertain, are referred to as *states of nature* and are considered uncontrollable. The selection of a particular course of action is controllable. Therefore, in this case we have one controllable variable (the course of action to take) and one uncontrollable variable (the state of nature that occurs). A common way to summarize the decision problem is a table of losses. The statistical decision literature commonly treats losses as non-negative, so we will adopt this practice. The loss table for this example is shown in the first three columns of Table 2. We now describe several non-probabilistic methods for making a decision on the particular alternative to implement. The boldest approach is to choose the alternative with the lowest possible loss. In this case, this would mean choosing Alternative 1 with an estimated loss of 5 if there is an industrial future land use. This is certainly an optimistic approach. This type of decision strategy is called the *minimin* approach: we *mini*mize the *min*imum loss. Of course, this alternative also has the largest possible loss, but this is not the concern with this decision strategy. This strategy is illustrated in the fourth column of Table 2.

Table 2. Non-probabilistic decision strategies

Alternative	State of Nature -Loss $K		Decision Strategies $K		
	Industrial	Residential	Minimin	Minimax	Minimax-Regret
1	5	50	5	50	20=max(5-5,50-30)
2	20	30	20	30	15=max(20-5,30-30)
3	10	40	10	40	10=max(10-5,40-30)

On the other end of the spectrum is the most conservative approach: choose the action with the smallest maximum loss. In this case, the choice would be Alternative 2, with a loss of $30,000. This is called the *minimax* approach: we *mini*mize the *max*imum loss. With this (pessimistic) approach the decision-maker is most concerned with the worst thing that could possibly happen. This is illustrated in the fifth column of Table 2.

Another (less straightforward) approach is based on the assumption that relative payoffs are more significant than absolute payoffs (Baker and Kropp 1985). This approach is a step up in sophistication from the minimin and minimax approaches. To illustrate this approach, suppose we choose Alternative 2 and yet an industrial land use occurs. We may then feel "regret" because we could have limited our losses to $5000 had we decided on Alternative 1. However, if residential land use occurs we would feel no regret because alternative 2 is the best choice if residential land use occurs. With this in mind we define the *regret* for any alternative/state pair to be the difference between the loss for this choice and state and the best possible loss for that state. In our example, it is the difference between the loss for a given land use and the smallest possible loss for that particular land use. We illustrate this by calculating the maximum regret for each alternative for the decision at hand in the last column of Table 2. Alternative 3 is preferred because it is the action with the smallest possible regret. This is called the *minimax-regret* approach, because it minimizes the maximum regret.

Our example illustrates that the type of strategy employed can influence the decision that is made. This process also operates at the subjective decision-making level. If more than one decision-maker is involved, different decision-makers may arrive at conflicting conclusions based on different strategies. This is frequently realized in environmental cleanup situations where two primary decision-makers are often the owner-operator [known under CERCLA as the Potentially Responsible Party (PRP)] and the regulator(s). The owner-operator has a financial interest in the outcome and often feels that the "uncontrollable" state in terms of risk is actually somewhat controllable. Therefore, they may tend to choose a minimin approach, either consciously or unconsciously, to simultaneously ensure protection and minimize financial losses. The regulator, on the other hand, having no assurances that the outcome is in fact "controllable" and having less of a financial interest in the selection of an alternative, would lean towards a minimax approach to ensuring protection given a worst-case outcome at some point in the future (e.g., return to residential land use). This can lead to more cost, but it is often the best way to ensure the protection against risk in a non-probabilistic framework. The fact that these decision strategies are non-probabilistic and that the majority of environmental restoration activities are inherently uncertain make these types of strategies less applicable to cleanup decisions. Non-probabilistic decision strategies are intuitive and applicable to decisions of the sort that people make everyday; however, when these strategies are applied to more complicated decision problems, they often come up short because they are not able to take all the technical information that exists into account. With this in mind, we turn to potentially more useful, though less intuitive, probabilistic decision strategies for evaluating environmental problems.

4.2 PROBABILISTIC METHODS

A natural desire in decision-making is to somehow incorporate the *likelihood* of each possible outcome. This is the key to the probabilistic approach. A method of probabilistic estimation is the expected value approach. Next, we discuss this method and the role of utility theory in estimation.

The first step in the expected-value approach is an estimation of the likelihood of each possible outcome. Assume that in our previous example the decision-maker is willing to assign probabilities and state that there is a 75% chance that there will be an industrial land use and that there is a 25% chance future land use will be residential. The best manner for incorporating this information is to calculate the expected value for a given decision. This is defined as the sum of the products of the loss for each outcome times the probability of that outcome. For example, the expected value for the contamination problem is:

$$\textit{loss if residential} \times \textit{probability that residential occurs}$$
$$+$$
$$\textit{loss if industrial} \times \textit{probability that industrial occurs}$$

Substituting the losses and probabilities:

Expected loss of Alternative 1 implementation $= 5 \times 0.75 + 50 \times 0.25 = 3.75 + 12.5 = 16.25$

Therefore, the expected loss of Alternative 1 is $16,250. Table 3 illustrates the expected value of all the alternatives given the above assumptions.

Alternative 1 represents the preferred alternative with the expected value approach because it represents the minimum expected loss. One may observe that if Alternative 1 is selected, the losses are either $5000 or $50,000; actual losses will not be $16,250. The expected value approach is useful to estimate how the strategy would do in the *long run* (i.e., over many trials). If an industrial land use is to occur 75% of the time, then, over the course of many future scenarios, Alternative 1 would be the best selection because losses are minimized most of the time. The expected-value approach is useful when

Table 3. Expected value table

Alternative	State of Nature		
	Industrial $K	Residential $K	Expected Value $K
1	3.75	12.5	16.25
2	15 (20*.75)	7.5 (30*.25)	22.5
3	7.5 (10*.75)	10 (40*.25)	17.5

one decision guides many other decisions (i.e., the decision will be "tried" many times). The same procedure can be done with regret: we define the expected regret for a particular course of action and choose the strategy with the smallest expected regret. Even though these values will be different from those obtained using the expected value approach, one can show that the decision reached is the same. This is illustrated in Table 4.

Table 4. Expected regret table

Alternative	State of Nature		
	Industrial $K	Residential $K	Expected Value $K
1	0	5	5
2	11.25	0	11.5
3	7.5	2.5	7.5

There are serious limitations to the expected value approach when applied to environmental problems. For example, if one of the possible outcomes involved the extinction of a species or a catastrophic type of impact, then these unquantifiable types of outcomes lose their immediacy in an expected-value framework. It may be that a particular alternative is the best to choose using an expected-value approach but the decision-maker would still be hard-pressed to consider this alternative seriously unless the probability of extinction or failure was "low enough." The question of irreversibility and of the occurrence of an adverse effect, that is too important to quantify in dollar terms, raises the question of whether a profit/loss index is the appropriate one for certain decision problems (e.g., the kind encountered within the RI/FS process). In instances where the consequences of an adverse effect are of a priceless magnitude, a more appropriate index must be selected to minimize (or maximize). Even if the expected value approach is not used, just the act of identifying and expressing the actions and their potential payoffs in a payoff table is a significant step towards making a better decision, no matter what the decision. As illustrated with the expected value approach, whether or not you will want to maximize expected gain (or, equivalently, minimize expected regret) depends on *your attitude towards risk*. The incorporation of attitudes toward risk leads us to yet another step up in sophistication.

The foundation of the study of attitudes towards risk and their relevance in decision-making was laid in von Neumann and Morgenstern (1953). The authors' ambitious goal was to "find the mathematically complete principles which define "rational behavior" for the participants in a social economy and to derive from them the general characteristics of that behavior." The "mathematically complete principles" take the form of relatively innocuous axioms that allow one to treat aversion to risk separately from probability judgements:

- Preferences are transitive (i.e., if a decision-maker prefers A to B and B to C, then she/he prefers A to C).
- A decision-maker will always choose the higher probability of winning an identical prize.
- The laws of probability govern a decision-maker's attitude towards a lottery. She/he is just as willing to play the lottery 100 times as playing its probabilistic equivalent once (also known as "no fun in gambling").
- There is a "certain equivalent" between the extreme outcomes of a lottery. In other words, if someone is offered a chance to receive either an item of small value or of great value dependent upon the flip of a coin, she/he would be willing to accept an intermediate sure thing (the "certain equivalent").
- The decision-maker is indifferent between the certain equivalent and the coin flip in the previous axiom. In other words, there is no "rush" to take the chance.

Assuming these axioms, von Neumann and Morgenstern showed that a rational decision-maker's preferences can be assigned specific values relative to each other. This number attached to a possible outcome is called the utility of the outcome. There are two fundamental implications of these axioms:

- A decision-maker's preferences can be encoded in terms of a *utility function*. This utility function represents a scaling of the values the decision-maker assigns to outcomes that captures his/her attitude toward accepting risk.

- A decision-maker's preference for various alternatives may be measured by calculating expected utility defined as the sum of utilities of possible outcomes weighted by their probabilities of occurrence). The preferred alternative will have the highest expected utility.

von Neumann and Morgenstern's fundamental result was that a rational decision-maker will maximize expected utility. That is, they will choose the action whose expected utility is the highest. Certainly people do not always adhere to the axioms above, and therefore they do not actually always maximize utility unconsciously. This theory is not designed to explain or incorporate irrational human behavior; rather, it is designed to indicate how people should make choices if their decisions are to be in accord with their own preferences. The difficulty in its everyday use, however, is that although utility theory can clarify the meaning of economic consequences, utility values are abstractions not represented by any physical measures of value that can be interpreted easily.

The theory presented in von Neumann and Morgenstern show that, assuming the axioms, every individual possesses a utility function with regard to comparable outcomes. In practice, the utility function is estimated by a trial-and-error approach whereby the utilities of non-extreme outcomes are gradually refined. Once the utility of possible outcomes is calculated, it is a simple matter to maximize the utility, assuming that the probability of the outcomes is known (itself not an easy task).

Limitations on the application of utility theory to environmental remediation problems are similar to those discussed previously for the expected value approach when one of the potential outcomes is totally unacceptable. However, for other environmental problems, the utility approach does allow one to construct risk-based analyses that are not necessarily dollar-based and that incorporate the decision-maker's attitude towards risk. Table 5 summarizes the deterministic and probabilistic decision strategies discussed in the previous two sections.

Table 5. Summary of decision strategies

Decision Type	Strategy Name	Strategy Description
Deterministic	Minimin	Select the course of action with the smallest minimum loss among all possible outcomes. *Gamble on the best possible outcome.*
	Minimax	Select the course of action with the smallest maximum loss among all possible outcomes. *Avoid the worst possible outcome.*
	Minimax-Regret	Select the course of action with the smallest maximum regret among all possible outcomes. Where regret is the loss due to not picking the best course of action.
Probabilistic	Expected-Value	Select the course of action with the minimum expected loss over the long run (i.e. over many trials).
	Expected-Regret	Select the course of action with the smallest expected minimum loss among all possible outcomes.
	Expected-Utility	Select the course of action that has the highest expected utility.

4.3 DATA WORTH AND THE VALUE OF PERFECT INFORMATION

Viewing the objective function in a probabilistic manner can be useful in evaluating the amount and worth of data needed to support a decision. A fundamental concept in the statistical decision theory literature is that of data worth—the value of data in making decisions between alternative courses of action. In particular, we will look at the quantification of upper bounds on how much one should be willing to spend to obtain data. Data worth can have broad applications for the environmental restoration process. These concepts can influence two problem-solving processes central to environmental cleanup activities: data collection and alternatives selection. As in any type of formalized decision approach, there is more to gain than just the quantitative output; walking through the process is itself advantageous in understanding the process of evaluating historical data, collecting new data, and making decisions based on current knowledge.

Complete information about a problem is an ideal but often excessively expensive status to achieve. In some applications, complete characterization would cost more than the penalty for choosing the wrong alternative. The purpose of the "value of perfect information" concept is to address the quantification of the maximum price one should pay to remove all uncertainty from the decision process. In practice, the value of perfect information can serve as a rough upper bound on the price one should pay for any data collection and is therefore a measure of data worth. In the following text, we illustrate its application to the quantification of upper bounds on data worth, as well as its use as a means of exploring the sensitivity of a decision to uncertainty. We also discuss potential applications of these methods to the types of decisions required under CERCLA when considering remedial alternatives.

In the decision science literature, the definition of perfect information depends on the conceptualization of processes involved, and it does not always necessarily mean the knowledge about what is going to happen. In particular, perfect information is sometimes referred to as the exact knowledge of the parameter values in the model of reality used to compute the objective function. This is not always the same as knowing the future. For example, if the problem depends on the results of a risk assessment, and probability distributions are assigned to represent the exposure variables and the dose-response relationship, then perfect information is the knowledge of the parameters in the stochastic process, which may themselves be used to estimate the *probability* that the risk exceeds some regulatory limit, not whether it will or not.

Assume we had to make a decision now and we have some method of choosing the alternative. Further assume that the objective function is constructed so that its minimization is preferable. If we had perfect data, then we would know the parameter values in the model. In this case we would choose the alternative with the smallest costs based on the perfect data. The difference between the costs associated with the decision we made using the data we have now and the smallest cost possible for the perfect data is called *the value of perfect information*. The value of perfect information is the most that any data could be worth within the constraints of the objective function we are using. Of course, we don't know what the perfect data *are*; if we did, we would choose the best alternative based on the perfect data. This does not mean that we cannot investigate the value of perfect information as a function of what the perfect data might be. There are several manners in which this can be carried out. Before discussing these methods, we introduce some notation.

Suppose that one must choose from a set of available alternatives. These alternatives may be discrete (e.g., finite number of distinct alternatives), continuous (e.g., depth of piling to use for a bridge), or a mixture of both (e.g., finite number of distinct alternatives, some of which have "optimizable" design parameters). In evaluating the objective function, one has a set of input parameters (e.g., risk, benefit, cost, etc.) for each of the alternatives, $\mathbf{x}$ will denote this set. For a given alternative $\mathbf{a}$, and input parameters $\mathbf{x}$, let $\mathbf{F(a;\, x)}$ denote the value of the objective function. Let $\mathbf{f(x)}$ denote the probability distribution of the input parameters $\mathbf{x}$ that will be used to evaluate the objective function. This represents our uncertainty regarding the "true" values of the input parameters. Let $\mathbf{a^*}$ denote the alternative with the smallest expected value of the objective function (i.e., smallest expected cost)[1] given the uncertainty specified.

Let $\mathbf{x_T}$ denote the "true" value of the input parameters $\mathbf{x}$, and let $\mathbf{a_T}$ denote the alternative that minimizes $\mathbf{F(a;\, x_T)}$; i.e.,

$$\mathbf{F(a_T;\, x_T) \leq F(a; x_T)} \qquad \text{for all alternatives } \mathbf{a}.$$

If $\mathbf{a^*}$ is the alternative selected given our current understanding of the problem, then the value of perfect information is the difference in the objective function for the alternative chosen and the correct alternative given the true state of nature. This is denoted by $\mathbf{VPI(a^*;\, x_T)}$ and is given by

$$\mathbf{VPI(a^*;\, x_T) = F(a^*;\, x_T) - F(a_T;\, x_T).}$$

Therefore, implicit in this definition is that the value of perfect information depends on:

1) our current uncertainty regarding the true state of nature described by $\mathbf{f(x)}$, and

---

[1]Finkel and Evans (1987) state that a decision maker can do no better than to choose the strategy with the lowest expected cost without the benefit of perfect information. However, other decision rules are possible.

2) the true state of nature.

In most cases, we will not have the luxury of perfect information. However, this does not mean that we should abandon the concept. In particular, we can explore various quantities related to perfect information:

- What the value of perfect information is expected to be
- What the value of perfect information distribution looks like
- How the value of perfect information depends on the input parameters, for each alternative (not necessarily the one that would be selected now)
- How the value of perfect information depends on uncertainty in the input parameters for each alternative

The expected value of perfect information for a fixed alternative $\mathbf{a^*}$, denoted by $\mathbf{EVPI(a^*)}$, is the sum over all input parameters of the value of perfect information weighted by the probability that the parameter values represent the true state of nature, and is given by

$$EVPI(a^*) = \int_x [\Phi(a^*; x) - \Phi(a_x; x)]\, f(x)\, dx$$

where

$\mathbf{a_x}$ = the alternative selected if $\mathbf{x}$ is the state of nature
$\mathbf{f(x)}$ = the probability distribution for the input parameter(s) $\mathbf{x}$.

Although any statistic of the distribution of values of perfect information could be potentially useful, the expected value is often suggested as the upper bound on data worth (Freeze et al. 1992, Davis and Dvoranchik 1971, Merkhofer 1987). An obvious drawback to this method, as well as any method that uses **f(x)**, is that it depends on what the distribution **f(x)** is that describes our current uncertainty. This distribution is also sometimes called the "prior" distribution. If one has little confidence in this distribution, then one may not be comfortable with results that are computed based thereon. For this reason, it may be useful to either explore the value of perfect information deterministically (e.g., study dependencies of the value of perfect information on the various input parameters) or even analyze the dependency of the expected value of perfect information on the *uncertainty* in one or more parameters. Although these methods may be prohibitive for objective functions with many input parameters, screening sensitivity analyses can be performed that may help exhibit a smaller set of "most important" parameters suitable for such an analysis.

We will illustrate some of these methods with a simple example inspired by those in Finkel and Evans (1987). We are to decide between three alternatives in the remediation of a hazardous waste site. These alternatives consist of 1) do nothing, 2) implement an alternative that will reduce risk by 70% at a cost of $8 million, and 3) implement an alternative that will reduce risk by 99% at a cost of $40 million. The objective function is the "total societal cost" and is calculated by

$$\Phi(a;R) = C_a + v\,(1 - e_a)R$$

where

C_a = cost of the alternative *a* (million \$)

e_a = efficiency of the alternative *a*

v = cost of failure (million \$/death)

R = health risk in the absence of any controls (deaths)

For the purpose of this example, we assume that the costs are a one-time cost and that the value of life has been determined to be \$1 million. The only input parameter is then the health risk *R*. Our uncertainty as to its true value is assumed to be represented by a probability distribution *f(R)*. As in Finkel and Evans (1987), we assume that *f(R;*m,s*)* is adequately represented by a lognormal distribution with geometric mean and standard deviation of m and s, respectively.

In this case, the expected value of perfect information for a given alternative *a* is determined by

$$EVPI(a;\mu,\sigma) = \int [\Phi(a;\,R) - \Phi(a^{*};\,R)]\, f(R;\mu,\sigma)\, dR$$

where a^* is the alternative that minimizes Φ if *R* is the true risk.

Figure 1 (Appendix A) shows a graph of the objective function as a function of risk for each alternative. This graph facilitates a deterministic visualization of the impact of risk in the objective function. It is useful in identifying break points where the minimization of the objective function changes alternatives. This minimization of social cost is given by the dashed line. If one knew the exact value of the risk *R*, then it would be a relatively simple matter to choose the alternative. Figure 2 (Appendix A) shows graphs of the value of perfect information for each alternative as a function of *R*. This is the difference between the value of the objective function for an alternative and the value of the *lowest* objective function for a given risk. We reiterate that these graphs are independent of the distribution for the risk.

The expected value of perfect information depends on the uncertainty in the risk *R* described by the lognormal distribution. At early stages of the remediation process, one may not be comfortable with specifying the parameters of this distribution. For this reason, it may be desirable to consider the expected value of perfect information as a function of these parameters for each alternative. Figure 3 (Appendix A) shows contour maps of the expected value of perfect information as a function of m and s for each alternative. These graphs illustrate regions of m and s where an alternative can be expected to produce large values of expected value of perfect information and regions where the expected value of perfect information becomes very sensitive to values of the parameters. It may become therefore easier to identify "stable" alternatives (i.e., ones in which the expected value of perfect information does not radically change in any region of m and s) and "dominating" alternatives (one in which the expected value of perfect information is comparatively low in all regions).

Figure 4 (Appendix A) shows how the alternative selected depends on the uncertainty in the risk, *R*. This graph is a probabilistic analogy to Fig. 1 as it also reveals "break" lines where the decision changes. The results of this graph agree with our intuition: as the uncertainty in the risk grows (i.e., s increases), eventually Alternative 3, the most expensive yet most effective, is the one selected. Such a graph could prove quite valuable if one can simply *bound* the uncertainty and/or mean of the risk. For example, if one is relatively sure that the geometric mean of the risk is above 4, then no matter what the uncertainty, Alternative 3 should be the one selected. Therefore, there would be no point in reducing the uncertainty. However, if one may suspect that the geometric mean is less than 4, then there could be sufficient justification for attempting to better estimate the uncertainty.

Frequently, one may only want to (or only be able to) consider the reduction in uncertainty in a few of the input parameters. Although one can formally define the concept of the perfect information of certain data, a simpler method is to explore the dependence of the expected value of perfect information as a function of the uncertainty in the various input parameters. This may aid in not only providing bounds on certain types of data but also can be used to compare the worth of some data relative to the remaining uncertainty in other variables. Indeed, it may be that there is no advantage in reducing the uncertainty in one parameter past a certain point without reducing another parameter's uncertainty as well. This insight could prove quite valuable.

Continuing with the example discussed in the previous section, we briefly illustrate these methods. As in Finkel and Evans (1987), we assume that the risk is calculated by the formula $R=D\,P$, where D is the total population dose (average individual dose in mg/kg/day times population exposed) and P is the potency (incremental probability of cancer incidence per mg/kg/day). What we wish to explore is how the expected value of perfect information, as well as the alternative selected, depends on the uncertainty in these parameters. Since the product of lognormal distributions is lognormal, assuming that both dose and potency are lognormally distributed makes the mathematics tractable. In particular, the geometric standard deviation of R can be calculated in terms of the those of D and P. Assuming that one has predetermined that the geometric mean is approximately 1.4, Fig. 5 shows how the alternative selected depends on the uncertainty in the dose and potency. From Fig. 5 (Appendix A) we can deduce that if the uncertainty in the potency cannot be reduced below 2, there is little use in reducing the uncertainty in dose below about 0.8. Similar comments can be made regarding reducing the uncertainty in potency.

Often in environmental problems, several alternatives exist which all satisfactorily pass regulation and risk-based criteria. The decision boils down to one of cost. Frequently, there are large uncertainties in the cost estimates of remediation alternatives. This is due in large part to contingencies, some of which data may eliminate. Assuming that the objective is to simply minimize this cost, the concept of the value of perfect information can be used to justify further collection of data. The objective function in this case is then just the cost of the alternative, which may be decomposed into various factors (e.g., cost of implementation, cost of failure).

Even for more complicated problems, similar methods can be applied, thereby potentially offering valuable insight that can be made available at the early stages of any remediation and/or data collection effort.

4.4 SUBJECTIVE INPUT

The use of some subjective judgement is always necessary in the evaluation of a decision. Often, a useful objective function will require subjective judgement as an input variable. A decision framework can incorporate human judgement by calling on qualitative expertise to estimate the value of an element in an objective function when empirical data are judged insufficient or are unobtainable. This can be accomplished by using either formal or informal expert elicitation processes. The sensitivity and importance of the outcome versus the increased cost and effort of formal methods is the determining factor in selecting between formal and informal methods.

Formal processes require the presence of an interviewer and a considered expert (or experts) in the field for which the variable needs to be estimated. The formal processes have the advantage of being able to better minimize many types of bias documented in the decision literature. These biases have been generally categorized into two groupings: subject bias and assessment bias. Subject bias is the tendency for the expert to introduce a cognitive bias into the estimates through the thought process of assigning numerical values to events. The assessment bias is the result of systematic error introduced by the assessment method.

Subject bias falls into two categories: management bias and expert bias. Management bias is the tendency for the subject to treat a variable as a goal rather than as an uncertain variable. This results in the over- or underestimation of a variable that is of concern based on whether the proposed management strategy is to minimize or maximize the quantity. Expert bias is based on the observation in numerous studies that an expert will consistently underestimate the amount of uncertainty that exists in the variable. This is usually the result of the subject attempting to meet the perception that an expert is expected to be certain about their specialty.

Assessment bias is often inherent in the methods used to aggregate the subjective individual inputs from a group. Simple averaging and Monte Carlo analyses based on discrete probability distributions can compound individual biases and do not account for differences in knowledge between members of the group. Group dynamics can also influence efforts to come to consensus probabilities.

In addition to countering biases, formal processes have a number of side benefits that include clarification of the issues central to the decision, greater confidence in the results, and facilitation of communication between involved decision-makers (Merkhofer 1987). It can be difficult to justify the necessity of formal methods for a certain decision, but given the continued necessity for decision-makers to make decisions on limited knowledge, these elicitation techniques can assist in making science-based policy decisions when hard evidence is unavailable.

5. POTENTIAL APPLICATIONS FOR ER ACTIVITIES

Experts from many fields have called upon government to take a more scientific approach to making decisions. As we have seen, the assessment and alternatives evaluation methodology that underlies CERCLA has roots in decision analysis. However, the long-term answer to the question of whether more quantitative decision methods have applicability in the environmental field can only be answered by the continued translation of decision theory into more practical approaches. These practical approaches have to promote solutions to the complex, highly uncertain, and politically sensitive problems that are at the core of environmental cleanups on government-owned land. Of course, the intent of this document is to identify key areas where decision analysis tools can make a contribution to decision processes for ORR-ER; this cannot be done by summarizing the decision literature alone. Practical investigations and applications of some of these techniques must be conducted to determine their usefulness.

Environmental decision-making operates at a number of different levels. At the programmatic level, making equitable and defensible budget decisions that maximize risk reduction for resources expended is an annual challenge. At the site level, data must be collected to support the selection of alternatives between competing technologies. Technologies must be chosen and funded and alternative ways of cleanup established to achieve long-term risk reduction. The practical application of decision analysis methods to these diverse problems lies in the ability to incorporate subjective judgement into objective functions that represent the problem to be solved. An area where decision analysis applications have already met with a degree of success is in fiscal resource allocation. Fair resource allocation decisions are difficult to make due to increasing competition for federal funds, increasing ES&H requirements, the lack of sufficient quantitative data, and the wide range of issues that must be addressed and balanced. Because of these complexities, the best means of approaching resource allocation decision-making is to incorporate subjective judgement. One method being successfully applied to support fiscal budget decisions at the DOE-ORO ER Program is the Environmental Restoration Benefit Assessment Matrix (ERBAM).

In the past, subjective judgement has played a major role in the ER fiscal funding decision-making process. However, the elicitation and use of management and expert opinion to justify yearly resource requests has been unstructured, informal, and undocumented. Recently, the need for defensibility and accountability of such subjective decisions regarding the activities undertaken by the DOE ER Program has increased. Stakeholders are increasingly interested in the processes and criteria DOE uses for making decisions regarding which sites merit action first. In addition, increasing competition for federal resources heightens the pressure on the DOE ER Program to defend decisions regarding the use of scarce federal dollars. Such pressure has precipitated the need to have a formalized approach for eliciting and using subjective judgement in a technically defensible manner to prioritize and justify fiscal funding decisions. A viable analytical system for making decisions regarding which projects, of a given set of diverse and fiscal activities (i.e., operational, compliance, improvement, programmatic coordination/planning, etc.) are to be funded and implemented within budget constraints requires the use of human judgement. In addition, when quantitative data are not consistently available for a set of candidate projects competing for resources, the integration of human judgement into a technically defensible process can support a fast, simple, and accurate comparative analysis of fiscal budget decisions.

Decision analysis techniques have been applied to successfully develop and implement a decision-support tool used to evaluate and prioritize a set of candidate fiscal ER projects for resource allocation. The tool is embedded in a qualitative risk management process that relies on management judgement and technical expertise. The decision model has been designed for managers and technical experts to use in evaluating fiscal budget decisions (e.g., ER projects at all five sites, the Off-site Program, and Central ER) within a common framework that incorporates a hierarchy of relevant ER objectives, the appropriate decision parameters (e.g., scales for measuring performance), and a suitable objective function (e.g., a means of estimating performance).

The ERBAM is a risk-based prioritization tool based on multiattribute utility analysis (MUA). MUA is useful as a basis for priority systems similar to ERBAM, but like many formal decision analysis techniques, can be difficult to apply correctly. We now discuss the elements and application of the ERBAM briefly, followed by a discussion of decision analysis techniques that can serve to reduce bias and application errors. The ERBAM is best described through a brief overview of its four primary elements:

- Decision criteria
- Procedure for generating a score
- Procedure for combining scores
- Rule/Output

These elements are illustrated in Fig. 6 (Appendix A).

The decision/selection criteria are used to define the risks and benefits associated with funding decisions. Six impact criteria are addressed in the model, including public health, environmental protection, site personnel safety, stakeholder preference, mission, and cost-effectiveness.

The procedure for generating a score involves the development of an objectives hierarchy, the establishment of performance scales, and the development of an objective function. Simply, an objectives hierarchy is assembled by determining the relative importance of the decision criteria and assigning them the appropriate weights. To establish performance scales, an ER project is viewed as an effort to reduce either 1) the magnitude of a unwanted consequence or 2) the likelihood of an unwanted consequence occurring. The ERBAM measures risk as:

$$W_i \; x \; S_i \; x \; L_i$$

where

W_i = weight of impact
S_i = severity of impact
L_i = likelihood of impact occurring

The ERBAM is designed to measure the reduction of risk, both in terms of severity and likelihood of occurrence, for a given project by requiring value judgement using a *no action* scenario and comparing that judgement to a valuation of residual risks *after* a project has been implemented. The objective function is

$$\text{Project Benefits} = (W_i \; x \; S_{ib} \; x \; L_{ib}) - (W_i \; x \; S_{ia} \; x \; L_{ia})$$

where

W_i = weight of impact
$S_{i\,b}$= severity of impact before project implementation
S_{ia} = severity of impact after project implementation
$L_{i\,b}$= likelihood of impact occurring before implementation
L_{ia} = likelihood of impact occurring after implementation

The procedure for combining scores involves the evaluation of a candidate project with respect to each of the six ERBAM decision criteria. This process yields six independent numerical values that represent the "benefits" a project provides in a given area of concern. The summation of the values represents the "net benefit score" that can be used to generate an initial ranking of candidate fiscal budget decisions.

The rule/output involves the ranking of budget decisions by net benefit score (from highest to lowest), enabling managers to distribute a projected amount of fiscal funding among projects beginning with those that provide the greatest or most valuable benefits overall. To determine which of these projects provide the greatest return on ER investment, a benefit-cost analysis can be conducted.

The identification of decision criteria, the development of an objectives hierarchy, the construction of performance scales, and designation of an objective function can be handled in a variety of ways, including the use of influence diagrams and other decision analysis techniques. Such techniques help reduce the bias and support the quantification of uncertainty in this type of decision-making process.

Relative to ER's programmatic and budget prioritization objectives, this approach to making fiscal resource allocation decisions provides many benefits and is illustrated in Table 6.

Table 6. Benefits of ER prioritization tool

ER OBJECTIVES	HOW ERBAM SUPPORTS ER OBJECTIVES
Use limited resources more effectively	Helps eliminate decision errors and biases Promotes consistency and "level playing field" Reduces duplicity of effort Controls the role of politics in decision making
Improve decision-making efficiency	Provides framework for organizing information and exploring issues Facilitates communication among parties Serves as a catalyst for action
Improve decision-making defensibility	Documents underlying assumptions and logic Promotes consensus

Source: Merkhofer, 1994

Any process that relies on subjective judgement can fall prey to decision bias. Both formal and informal decision analysis techniques can be used to reduce both subject bias (management and expert) and assessment bias. One such formal technique includes the development of influence diagrams that graphically display the influences between key choices and uncertainties in a decision problem. The diagrams are intuitively digestible by non-experts and relatively simple to generate. An informal means of reducing decision bias is to establish an expert panel that is responsible for approving the evaluating and scoring of the activities.

Once a funding decision has been made to conduct a remedial action at a contaminated site, an additional decision must be made concerning which alternative to implement. Data must be collected and evaluated against certain decision rules as embodied in the RI/FS process. This is more complex than the resource allocation decision because there are more than two alternatives (to fund or not to fund), and the decision-makers and decision criteria vary from site to site.

Under CERCLA, the decisions about how much data to collect and the criteria for decision-making are established during the DQO process. This process consists of developing qualitative and quantitative statements that help specify the quality of data required to support decisions during remedial activities. Data of known or acceptable precision, accuracy, completeness, representativeness, and comparability are necessary for the construction of defensible decisions. Indeed, the development of DQOs is an important step in assuring quality data for site or facility characterization, fate and transport modeling, and exposure estimation. These statements are established before data collection during the project scoping and sampling and analysis planning phases. Construction and evaluation of objective functions that represent the environmental problem can be used to quantitatively implement the DQO process from project inception to remedial design and ensure that the goal of directing the data collection process towards alternative selection is achieved in a defensible manner.

The quantitative evaluation of DQOs through the construction and use of objective functions and data worth concepts is an effective means for implementing the intent of the CERCLA process. Objective functions can be developed for action and no-action scenarios and used for optimizing the collection of site data in the RI.

As remedial alternatives are identified and refined, the objective function also should be reworked so that additional data is collected for the purpose of reducing uncertainties in the developed objective function. The DQO process becomes the forum for carrying this information through the process and evaluating the specific decisions embodied by the RI and FS. "The value of obtaining additional data or increasing data quality has traditionally been based on professional judgement for RI/FS projects. The intent of the DQO process is to provide a systematic approach for the evaluation of the risk associated with a wrong decision and for determining levels of uncertainty associated with decisions to provide a framework for the RPM (Remedial Project Manager)" (EPA 1987). Decision analysis applications that can give a rough estimate of an upper bound on data worth, such as the value of perfect information described earlier, will assist in meeting the goals of the DQO process. However, the standardization of an objective function or decision selection strategies for consistent implementation at all contaminated sites on the ORO-ER would not serve the needs of decision-makers since the flexibility needed to make sound decisions would be eliminated. The use of objective functions for alternative selection at contaminated sites must be conducted on a site-by-site basis, based on the conceptual model for the particular site and the needs of the involved decision-makers. *No single algorithm or strategy fits all approaches to environmental decision-making*.

6. CONCLUSION

A systematic decision approach cannot remove the subjectivity from the decision process since different individuals (facing the same situation and having the same information) can arrive at different "optimum" decisions. Further, as noted by Finkel and Evans (1987), even an "optimal" decision may not lead to an "optimal" outcome. This effect and the lack of a guarantee of satisfactory results should not be construed as a failing of formal decision methods. The value gained by implementing a systematic approach is in the clarification of the logic behind the decision and the identification of the real issues that can delay consensus. Although the quantitative results of a decision analysis will not actually make the decision relative to which alternative is most appropriate, it is worthwhile to evaluate how the results of different decision-making strategies may affect the alternative selected. It should be emphasized that no approach to decision-making is fool-proof, and, as noted by Baker and Kropp (1985), "decision science has not evolved a universally accepted methodology for analyzing social decisions involving risk."

The decision-making methods based on the calculation of objective functions rely more heavily on the calculation of numbers for making decisions than more commonly used subjective practices. This is often unavoidable if the project goal is a thorough development of the assessment problem. Despite this emphasis, the decision-makers will require a clear, concise communication of the numerical results, the uncertainties, and the simplifying assumptions upon which results are based. The tendency to overburden the decision-makers with numerical details must be avoided. Equally important for the decision-making process is to clearly identify areas where no numerical results are possible and value judgements are required.

Although formal decision-making tools place a lot of power in the hands of assessors, it must be kept in mind that decision-making for environmental problems should not be performed by assessors: "Balancing the benefits against the risks belongs not in the domain of science but to society. The judgement is a value judgement-a social rather than a scientific decision." (Commoner 1977). However, it is the job of the assessor to present unbiased descriptions of the benefits and risks of possible solutions for an environmental problem in a manner that is conducive to decision-making. The word risk implies uncertainty, and therefore any environmental decision that involves risk requires a thorough evaluation and presentation of the uncertainties present. The methods presented here can provide a firm foundation for investigating environmental decision problems that involve risk while conforming with and maintaining the intent of the environmental laws and guidance that govern assessment activities.

An obvious drawback to all of these methods is that it can seem tedious to bear the technical burdens involved, and it can be difficult to communicate the results. Voices from the past have anticipated resistance to formal approaches, as evidenced by R. Howard's words from almost 30 years ago (Howard, 1967):

> *"...it is inevitable that in the future both technical and managerial decision makers will employ formal logical methods in decision making. The transition probably will be painful."*

However, with the limited resources available for environmental cleanup, with the costs associated with some remediation alternatives exceeding hundreds of millions of dollars, and with the likelihood that all decisions will be subject to intense scrutiny, an investment into formal approaches may reap rewards that easily absorb the costs of any such "painful transition."

7. REFERENCES

Baker, K.R. and D.H. Kropp. 1985. *Management Science: An Introduction to the Use of Decision Models*. John Wiley & Sons, Inc., New York.

Commoner, B. 1977. *New York Times Magazine* (September 25): 73.

Crouch, E.A.C., and R. Wilson. 1982. *Risk/Benefit Analysis.* Ballinger Publishing Company, Cambridge, MA.

Davis, D.R. and W.M. Dvoranchik. 1971. *Evaluation of the worth of additional data*. Water Resources Bulletin. 7, No.4, 700-707.

Finkel, A.M. and J.S. Evans. 1987. *Evaluating benefits of uncertainty reduction in environmental health risk management*, JAPCA 1164-1171.

Fischoff, B., S. Lichtenstein, P. Slovic. S.L. Derby, and R.L. Keeney. 1981. Acceptable Risk. Cambridge University Press, Cambridge.

Freeze, R.A., J. Massman, L. Smith, T. Sperling, and B. James. 1990. "Hydrogeological Decision Analysis: 1. A Framework." *Ground Water,* v. 28, pp. 738-766.

Freeze, R.A., B. James, J. Massman, T. Sperling, and L. Smith. 1992. *Hydrogeological Decision Analysis: 4, The concept of data worth and its use in the development of site investigation strategies*. Ground Water 30: 574-588.

Howard, R.A. 1966. *Information Value Theory*. IEEE Transactions on Systems Science and Cybernetics: 1, 22-26.

Keeney, R.L. 1982. *Decision Analysis: An Overview*. Operations Research. Vol 30, No.5, 803-833.

Mansfield, E. 1980. Statistics for Business and Economics. W.W. Norton & Company, Inc., New York, NY.

Merkhofer, M.W. 1987. Decision Science and Social Risk Management, A Comparative Evaluation of Cost-Benefit Analysis, Decision Analysis, and Other Formal Decision-Aiding Approaches. D. Reidel Publishing, 1987.

Merkhofer, M.W. 1994. *Understanding Formal Priority Systems*. Presentation at the Energy Facilities Contractors Group/ Prioritization Methodology Users Group First Annual Meeting, February 8, 1994.

U.S. EPA. March 1987. Data Quality Objectives for Remedial Response Activities, Development Process. EPA/540-G-87/003, OSWER Directive 9355.0-7B. U.S. EPA, Washington, D.C.

U.S. EPA. October 1988. Guidance for Conducting Remedial Investigations and Feasibility Studies Under CERCLA. EPA/540/G-89/004, OSWER Directive 9355.3-01. U.S. EPA, Washington, D.C.

U.S. EPA. 1989. Risk Assessment Guidance for Superfund. EPA/540/G-89/004, OSWER Directive 9355.3-01. U.S. EPA, Washington, D.C.

von Neumann, J. and O. Morgenstern. 1953. Theory of Games and Economic Behavior, Princeton University Press.

APPENDIX A

Figures

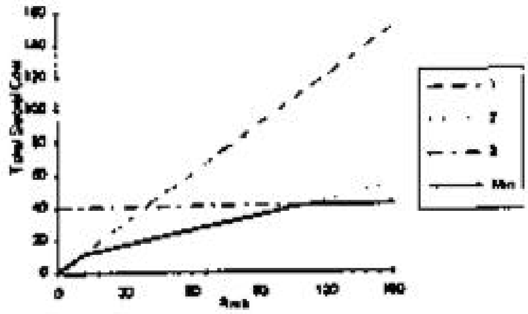

Figure 1. Total social cost for each alternative and minimum social cost as a function of risk

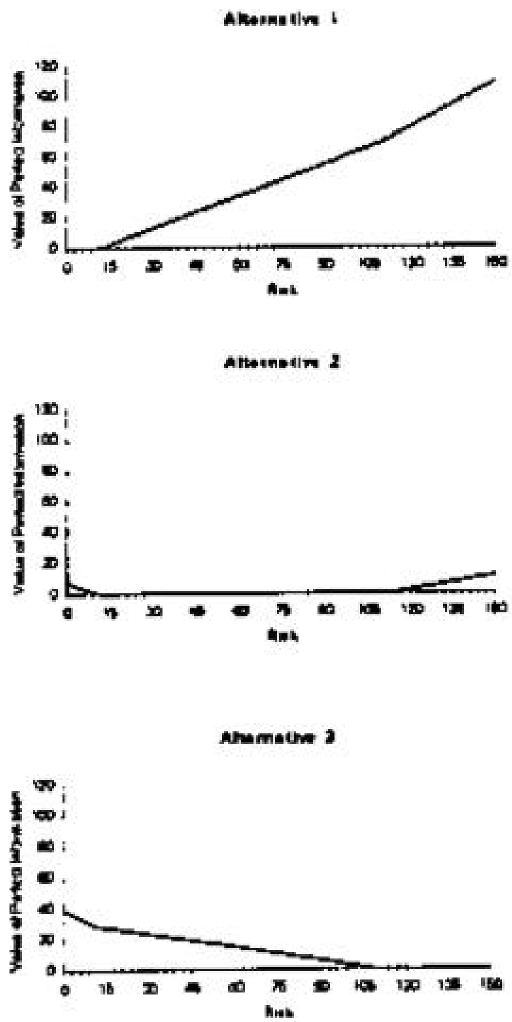

Figure 2. Value of perfect information for each alternative as a function of risk

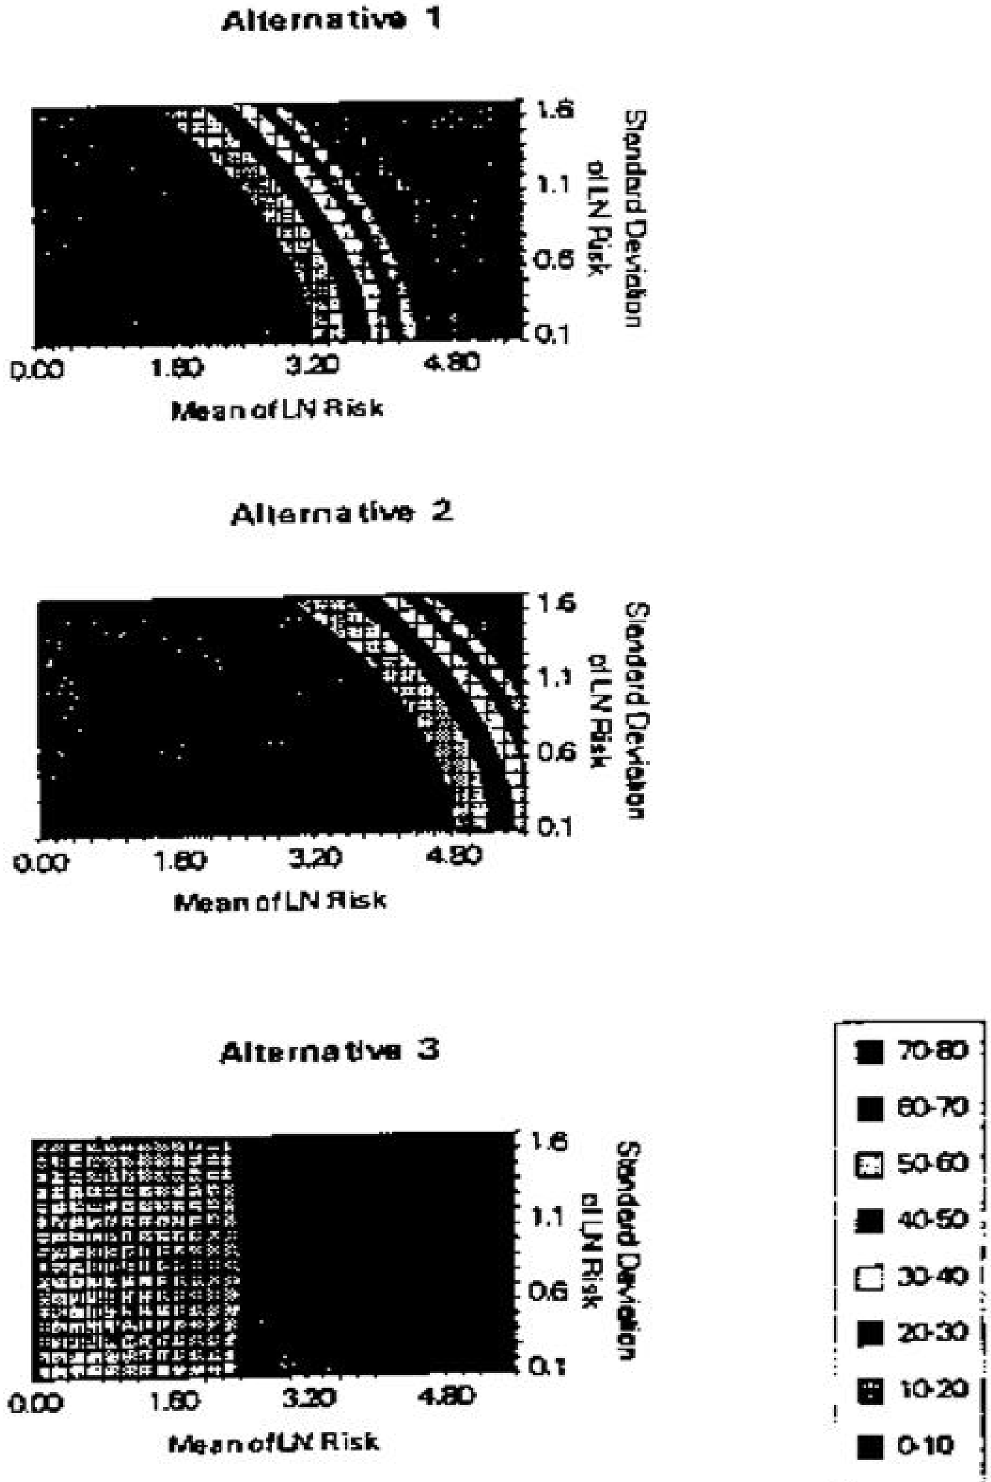

Figure 3. Expected value of perfect information for each alternative as a function of the uncertainty in risk

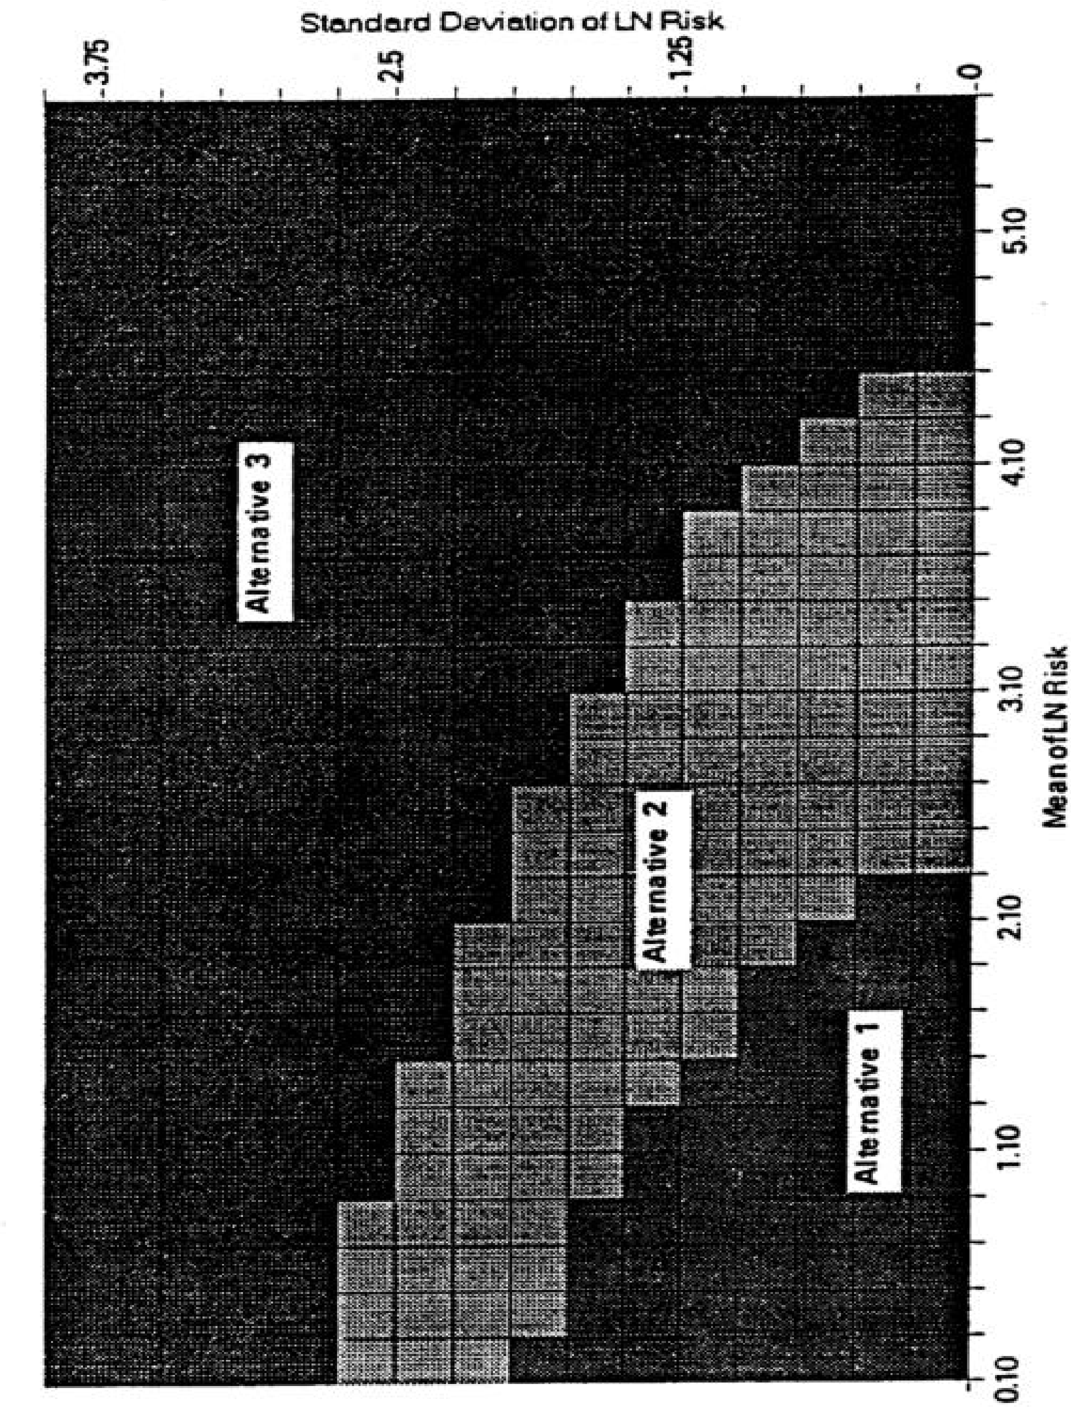

Figure 4. Sensitivity of alternative selected to uncertainty in risk

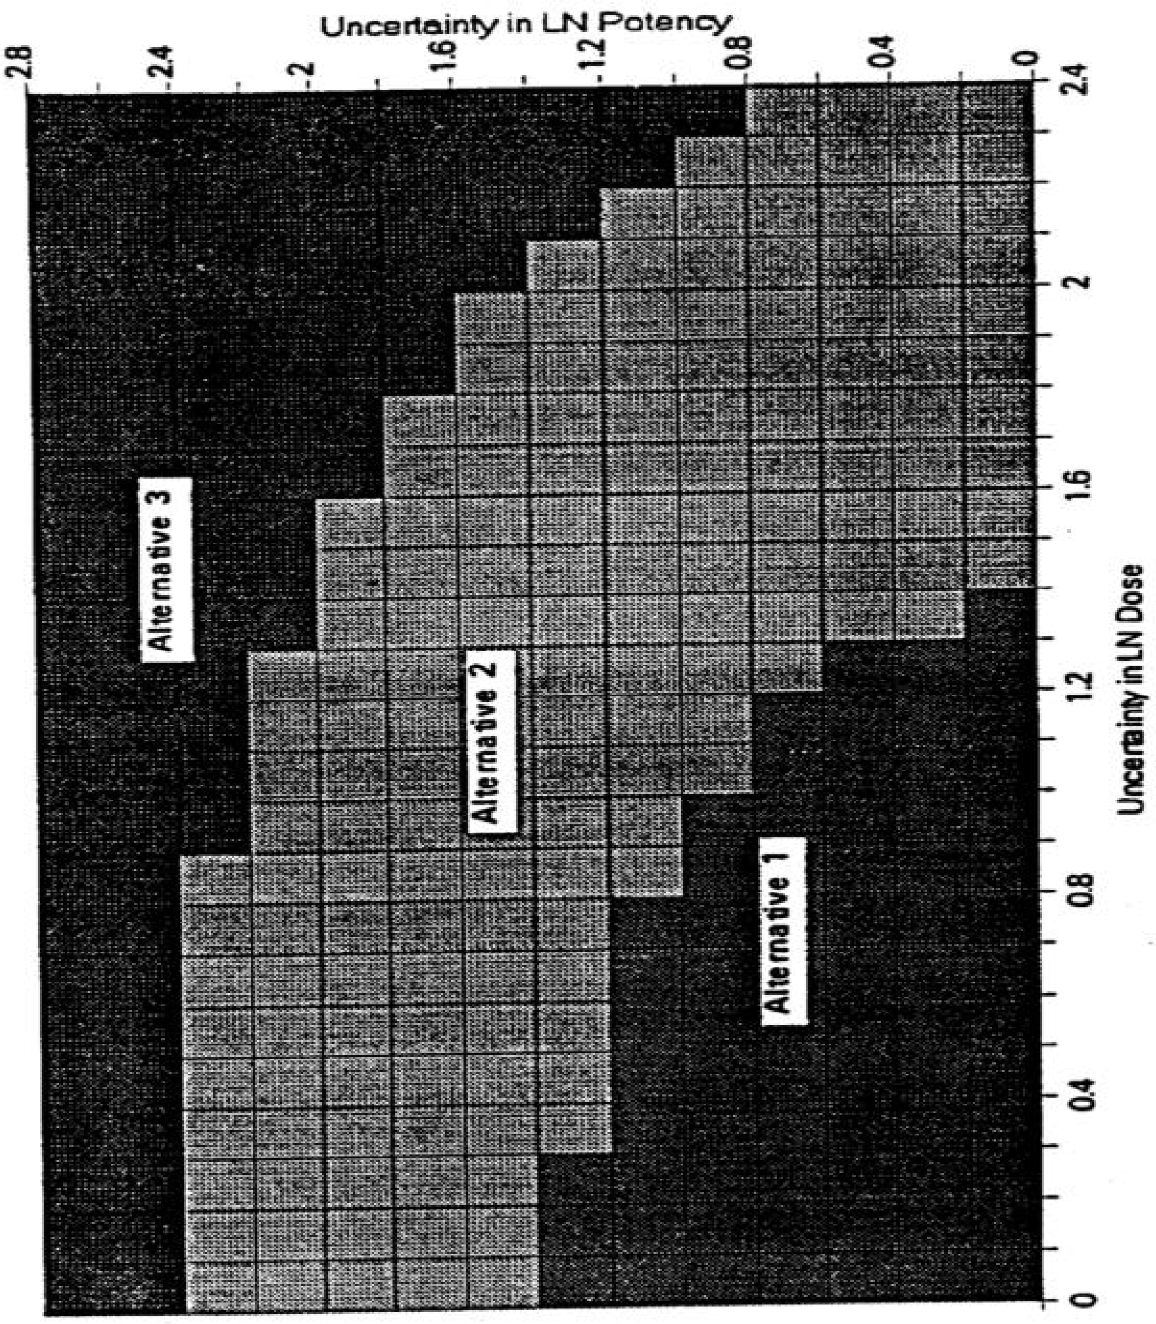

Figure 5. Sensitivity of alternative selected to uncertainty in dose and potency

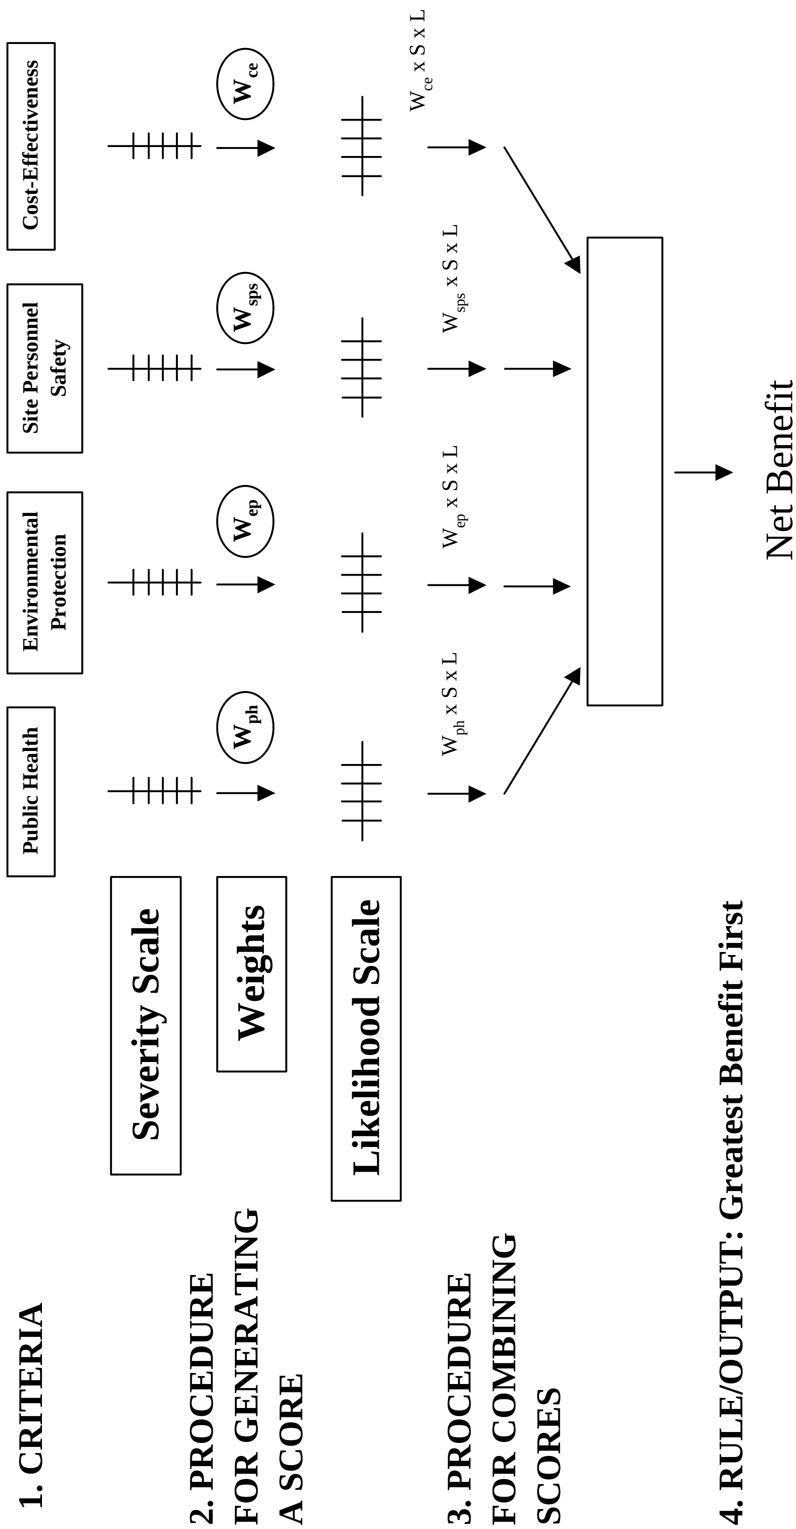

Figure 6. Elements and applications of ERBAM, a risk-prioritization tool