A METHODOLOGY FOR AIDING REPOSITORY SITING DECISIONS

August 1985

I. BACKGROUND AND INTRODUCTION

On December 20, 1984, the Department of Energy (DOE) published draft environmental assessments (EAs) to accompany the proposed nomination of five sites as suitable for site characterization for the first geologic repository. The final chapter of the draft EAs (Chapter 7) contained a comparative evaluation of the five sites against the DOE's siting guidelines (10 CFR Part 960). To determine which three sites appeared most favorable for recommendation for characterization, three simple quantitative methods were employed to aggregate the rankings assigned to each site for the various guidelines. These methods were reviewed by several groups commenting on the draft EAs, including the National Research Council's Board on Radioactive Waste Management. Two of the methods (averaging and pairwise comparison methods) were criticized for lacking firm theoretical foundations. The third method, described variously as the utility-estimation, rating, or weighting-summation method, was criticized because its application did not follow the formal procedures suggested by the professional literature. In response to these comments, the DOE has developed a more formal utility-estimation method (hereafter referred to as a decision-aiding methodology) to provide a more defensible overall comparative evaluation of sites. That methodology is described in this document.

Relationship to, and Consistency with, the Siting Guidelines

The decision-aiding methodology must be consistent with the DOE siting guidelines, which consist of implementation guidelines, system guidelines, and technical guidelines. System and technical guidelines are defined for the postclosure and the preclosure periods. The system guidelines contain broad requirements that are based generally on the objectives of protecting public health and safety and the environment during repository construction, operation, closure, and decommissioning and of assuring reasonable costs. The data required for a complete assessment of site performance against the system guidelines, however, will be available only after site characterization and the concurrent socioeconomic and environmental investigations. In lieu of such data and analyses, technical guidelines were defined for each system guideline to give a measure of the potential suitability of a site before detailed studies of the site can be performed.

The postclosure technical guidelines govern the performance of a repository over the long term and are concerned with the physical properties and physical phenomena at a site (e.g., geohydrologic conditions). The preclosure technical guidelines are concerned with the impacts of a repository before it is closed. The preclosure guidelines are divided into three subgroups: (1) preclosure radiological safety; (2) environment, socioeconomics, and transportation; and (3) ease and cost of siting, construction, operation, and closure.

The implementation guidelines establish a number of requirements that constrain the application of the methodology. Briefly, they require that primary significance or weight be given to the postclosure guidelines and that, for the preclosure period, radiological safety; environmental impacts, socioeconomics, and transportation; and the ease and cost of siting, construction, operation, and closure be considered in decreasing order of importance.

The decision-aiding methodology is used primarily to aggregate the performance rankings assigned for the technical guidelines because the data collected to date are insufficient for a conclusive comparison of sites on the basis of the system guidelines.

Role of the Methodology

It has been suggested that the ranking of sites should be based on the results of performance assessments. However, the assessments that can be performed before site characterization are preliminary, inconclusive, and incomplete; for example, they do not account for the effects of heat on the isolation capability of the host rock. Nonetheless, the results of the preliminary performance assessments can be used for consistency checks against the results obtained from the formal methodology, which is more specific.

The decision-aiding methodology is intended to provide a framework for systematically accounting for the professional judgment required in selecting sites for characterization. It should permit the scientific and value judgments to be made explicit to the reviewer. Furthermore, the methodology should permit sensitivity analyses and, if necessary, more-complex uncertainty analyses that can be used to explore the sensitivity of the decision to alternative professional judgments. The methodology is not intended to be used, by itself, to determine which sites should be recommended; its purpose is to provide a technical basis, in conjunction with the provisions in the siting guidelines on the diversity of rock types and other information, for such a decision. The decision as to which sites will be recommended will be made by the Secretary of Energy.

Methodology Overview

The technical name for the decision-aiding methodology is multiattribute utility analysis. The procedures and sequence of application follow those recommended in the professional decision-analysis literature (e.g., Keeney and Raiffa, 1976; Keeney, 1980; Edwards and Newman, 1982; Hobbs, 1982; Merkhofer, in press).

The methodology consists of six steps: (1) identifying and organizing objectives, (2) establishing performance measures and associated scales for measuring the extent to which a site meets the objectives, (3) verifying the independence assumptions necessary for the simple aggregation of assessments against competing objectives, (4) assessing single-attribute utility functions, (5) assigning scaling factors or weights, and (6) performing numerical calculations and sensitivity analyses.

The various steps of the analysis are being conducted by a DOE team consisting of experts in decision analysis, the technical disciplines corresponding to the technical siting guidelines, and repository performance. The technical information for the analysis is being obtained from the final EAs. Value tradeoffs and other judgments necessary for sensitivity analyses are being provided by DOE management and staff.

The next section of this document describes the basic concepts and methods on which the methodology relies. Section III describes the basic steps of the methodology in detail.

II. CONCEPTS AND METHODS USED IN THE DECISION-AIDING METHODOLOGY

This section introduces the basic concepts and methods that provide the logical foundation for the decision-aiding methodology. Readers not concerned with the theory on which the methodology relies or those already familiar with decision theory may wish to skip to Section III, which provides a detailed description and explanation of the decision-aiding methodology.

Basic Structure and Logic of Decision-Aiding Methodology

A fundamental tenet of virtually all decision-aiding methodologies is that understanding can be improved by dividing a decision into its parts, analyzing the parts separately, and combining the results at the end. Common sense suggests that this divide-and-conquer strategy improves the quality of decisions.

Perhaps the most important "decomposition" produced by decision-aiding methodologies is the separation of knowledge from preferences, or value judgments. Decision theory argues that a decision should logically depend on the likelihoods of the possible consequences of each alternative and the relative preferences of decisionmakers for those consequences. Figure 1 shows how decision-aiding methodologies generally separate knowledge and judgment. First, alternatives are characterized in terms of technical factors or descriptors. Next, an assessment is made of the consequences associated with the selection of an alternative with the specified characteristics. This assessment provides measures of the performance of the alternative. Finally, the various performance measures are evaluated and integrated to obtain an overall measure of the desirability of the alternative. Nearly all decision-aiding methodologies have this basic form (Merkhofer, 1983).

An advantage of a methodology of the form shown in Figure 1 is the division of responsibility between technical experts and policymakers. Technical experts are responsible for all aspects of the methodology that deal with information or knowledge. For example, a comprehensive and accurate description of an alternative in terms of technical descriptors requires a detailed understanding of the characteristics of the alternative and is therefore the logical responsibility of those most familiar with the alternatives. Similarly, the assessment of the possible consequences of an alternative--which must be based on all available information, including collected data, models, and professional judgment--is also the logical responsibility of technical experts.* Those aspects of the methodology that

*The same experts, however, need not both characterize the alternatives and estimate their consequences, because the latter task relies more on an understanding of cause-and-effect relations than a detailed understanding of the options. The ability to separate the tasks assigned to such experts not only permits data to be collected from those most qualified to provide it but also helps to reduce the potential for biases.

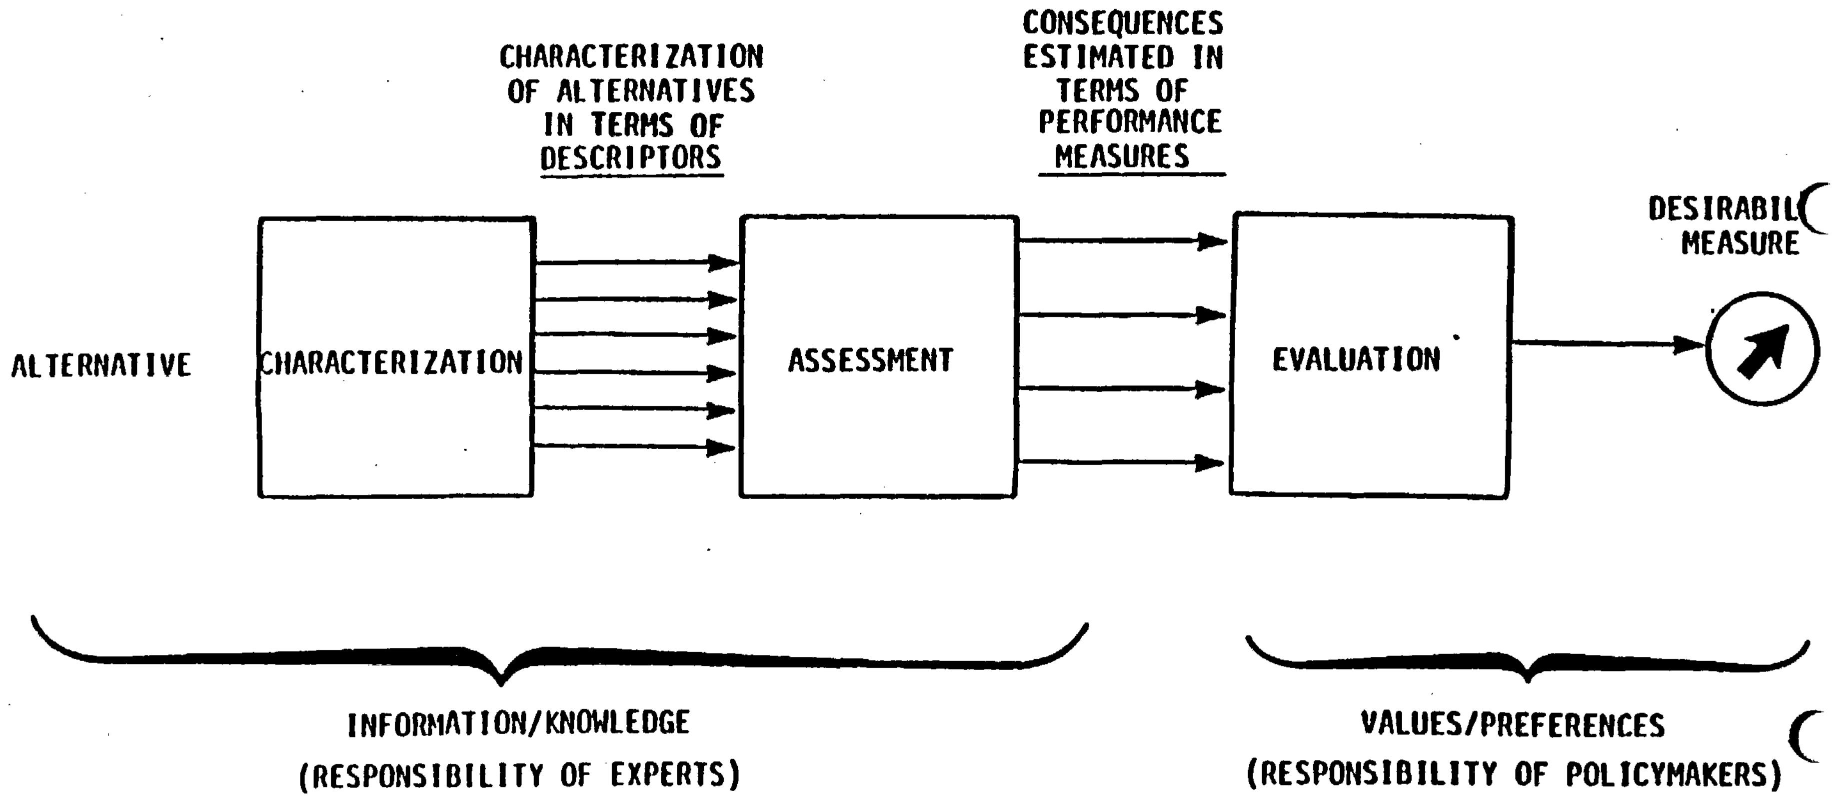

Figure 1. Basic form of most decision-aiding methodologies.

deal with preferences, or value judgments, on the other hand, are assigned to policymakers. To establish preferences, it is necessary to consider the objectives and the values of stakeholders. This is the logical responsibility of policymakers.

To represent and account for knowledge and judgment, the decision-aiding methodology relies on concepts that are well established in the decision-analysis literature. The most important of these concepts are the multiattribute utility theory and probability theory.

Multiattribute Utility Theory

Performance measures provide assessments of an alternative along specific dimensions. The multiattribute utility theory provides a means for making these assessments commensurable in terms of a common scale of value (Holloway, 1979; Keeney and Raiffa, 1976; Keeney, 1980).

According to the theory, the value, or "utility," of an alternative can be expressed as a mathematical function of its performance measures. Thus, if numerical values are assigned for each performance measure for each site, then a numerical utility for each site can be calculated with the property that the more desirable sites will have the higher utility values.

Nearly all practical applications of the multiattribute utility theory include independence assumptions that permit the utility function to be decomposed. In most such cases, the utility function has a linear additive form. Expressed mathematically, if $x_1, x_2, x_3, \ldots, x_n$ are the performance measures of interest and independence holds, then the multiattribute utility can be calculated from an equation of the form

$$U = w_1U_1(x_1) + w_2U_2(x_2) + w_3U_3(x_3) + w_nU_n(x_n)$$

where U_1, U_2, etc., are single-attribute (marginal) utility functions (described below) for each performance measure and w_1, w_2, etc., are weighting factors.

Although independence assumptions often seem difficult to interpret conceptually, procedures for their verification are available. Keeney (1980), for example, gives an illustrative series of questions for verifying an additive form for the utility function. If such procedures indicate that the appropriate form of independence cannot be assumed, then the definitions of performance measures must be changed until independence does apply (or more-complex forms than the linear additive for the utility function must be used).

The advantage of the additive form is that it greatly simplifies the construction of a multiattribute utility function. Although general multiattribute utility functions are difficult to derive, single-attribute utility functions are relatively easy. Therefore, independence permits a multiattribute utility function to be constructed by (1) assessing single-attribute (marginal) utility functions for each performance measure, (2) assessing weighting factors, and (3) calculating the overall utility of a site as a weighted average of the marginal utilities.

Techniques for constructing utility functions are described by Keeney and Raiffa (1976), Keeney (1980), and Changkong and Haimes (1983). A possible form of a single-attribute utility function is shown in Figure 2. Although the utility function in Figure 2 is linear, utility functions are often nonlinear, reflecting, for example, a judgment of the diminishing utility of increments of performance beyond some satisfactory level.

Several techniques can be used to establish weighting factors. The simplest approach is to interpret the weighting factors as the relative importance of the objectives that underlie the performance measures. Subjects may be asked to allocate 100 percentage points among the various objectives, according to their judged importance. Although this method is simple, it is difficult to make declarative statements about the relative importance of competing objectives, and inaccuracies are likely to be produced. A preferred method for determining the weighting factors is to establish a series of "indifference" points between different combinations of performance-measure values. If the points are of equal preference, their utilities are equal, and a series of linear equations relating the utilities of the indifference points can be developed. If the indifference points are established so that only two performance measures vary at a time, the resulting equations can be easily solved for the weighting factors. A simple example is given in Section III. A detailed example that illustrates the assessment and equation-solution process is given by Keeney (1980).

Probability Theory

The concept of probability is used in the decision-aiding methodology to account for uncertainty. Following the perspective of decision analysis, probabilities (numbers between 0 and 1) represent an individual's degree of belief concerning some uncertain quantity. In the decision-aiding methodology, descriptors (e.g., ground-water travel time), performance measures (e.g., the total preclosure costs of the repository), and utilities (numbers between 0 and 100) may be uncertain. Probabilities may therefore be assigned to reflect the uncertainty about the appropriate value for descriptors, performance measures, and utilities. Where possible, historical data and statistics should be used in assigning probabilities, but if such information is not available, expert judgment can be substituted.

Probabilities can be displayed in several ways, depending on whether the uncertain variable is discrete (i.e., it can have only a finite number of possible values) or continuous (i.e., it can have any value within some range). Three alternative displays for an uncertain variable--the tree form, a cumulative probability distribution, and a probability density function--are shown in Figure 3, which illustrates uncertainty about the uncertain descriptor ground-water travel time.

In practice, probabilities for uncertain variables can be elicited from experts, using probability encoding techniques (Spetzler and Stael von Holstein, 1975). Experience has shown a number of encoding procedures to be effective. The three basic types of encoding methods are (1) probability methods, which require the subject to respond by specifying points on the probability scale while the values remain fixed; (2) value methods, which require the subject to respond by specifying points on the value scale while

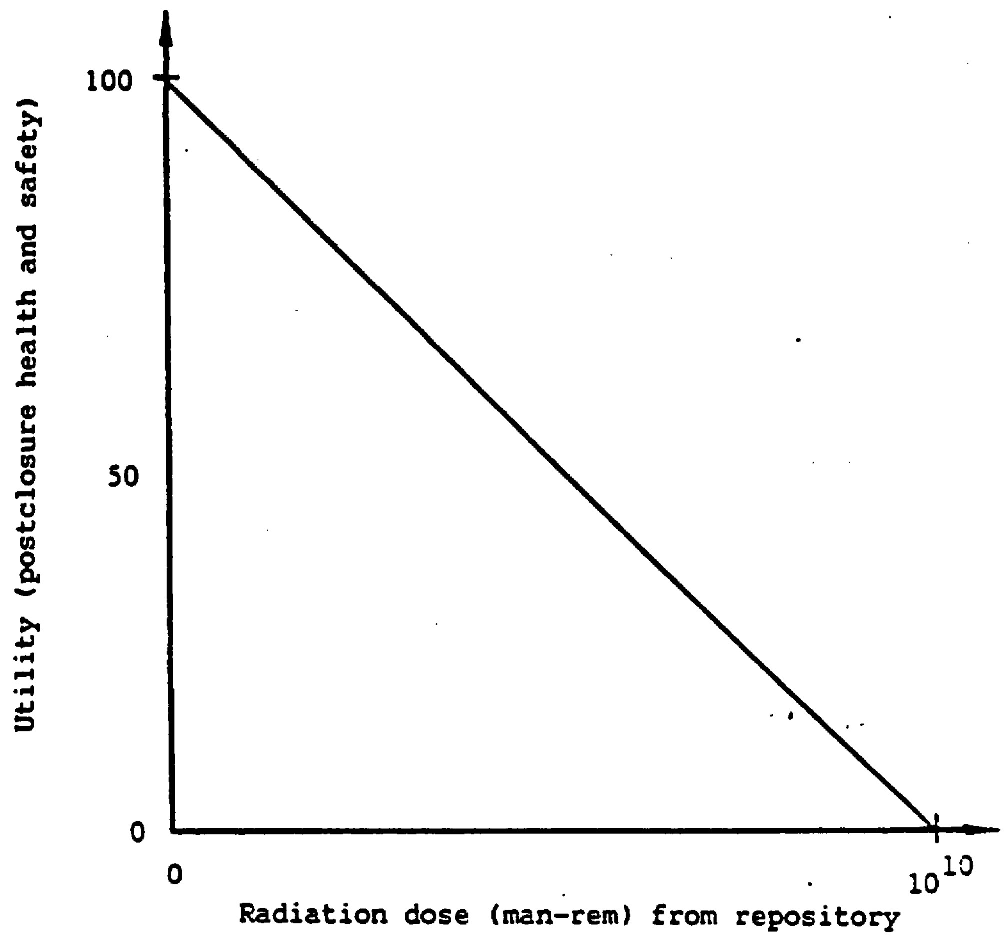

Figure 2. Sample utility function.

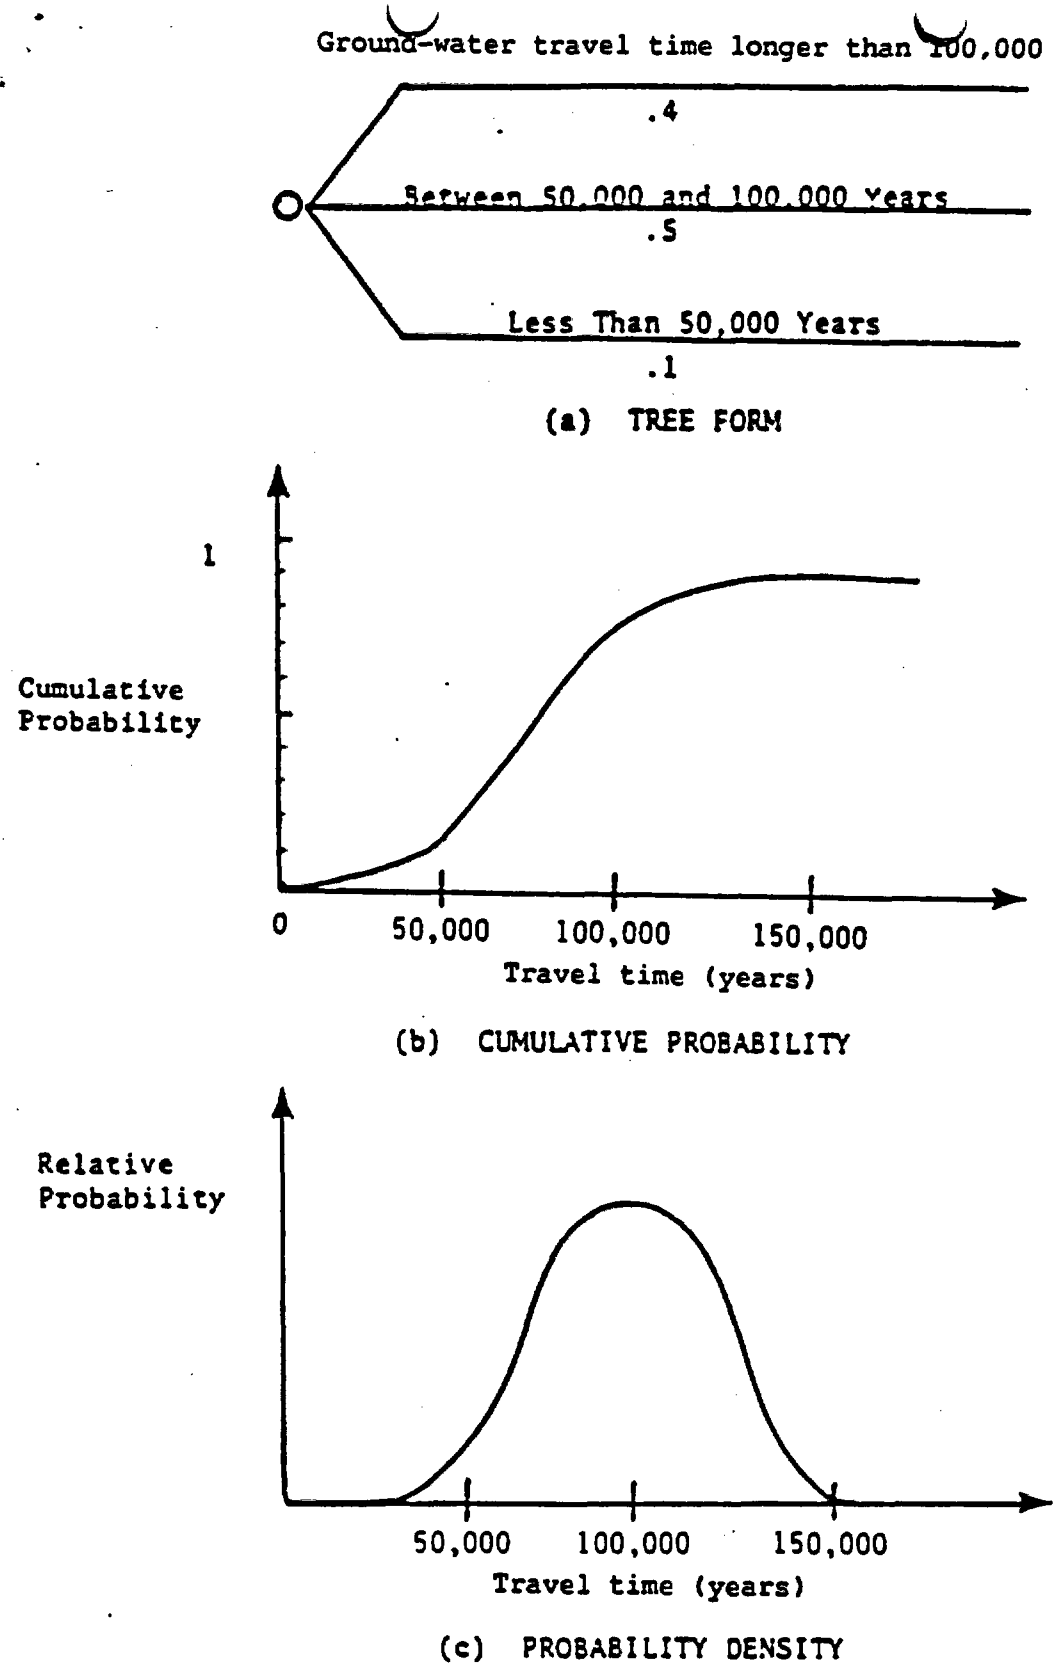

Figure 3. Alternative representations of probabilities: (a) tree form, (b) cumulative probability, and (c) probability density.

the probabilities remain fixed; and (3) probability/value methods, which ask questions that must be answered on both scales jointly (the subject essentially describes points on the cumulative distribution). Each of these encoding procedures can be presented either in a direct-response mode, in which the subject is asked questions that require numbers as answers, or in the indirect-response mode, in which the subject is asked to choose between two or more lotteries. The lotteries are adjusted until the subject is indifferent to choosing between them. Either external reference events (alternative lotteries defined on some external event, such as a probability wheel) or internal reference events (events defined on the same value scale as the uncertain quantity) can be used in the indirect mode.

Uncertain variables are often dependent on one another in the sense that knowledge of one influences information about the others. In such cases, the probability assigned to any one variable must be conditional on the values of the others. The tree form is useful for displaying such conditional probabilities. To illustrate, Figure 4 shows a probability tree with conditional probabilities that might be assigned to reflect the dependences between two descriptors--the average fault density in the vicinity of a site and the average rate of faulting. Gathering conditional probability assignments amounts to asking such questions as, "What are the odds that the rate of faulting exceeds X cm/year, given that the current density is Y cm^2/m^3?"

An important question involving the algebra of probability theory is how to compute the probabilities associated with an uncertain variable that is assumed to be related to other uncertain variables. Occasionally, an equation may be defined that permits a performance measure to be approximately calculated from values provided for descriptors. Similarly, an equation may be defined for relating utilities to performance measures. If probabilities can be assigned to the uncertain variables that serve as the inputs to such equations, then techniques exist for computing the probabilities for the output variables. The two principal techniques are Monte Carlo analysis and probability-tree analysis--well-known techniques that are discussed by Holloway (1979). When properly applied, both methods give essentially the same result.

The extent to which these techniques will be required in the application of the methodology to the problem of determining which sites should be recommended for characterization will depend on, among other things, whether simpler techniques like sensitivity analyses prove adequate.

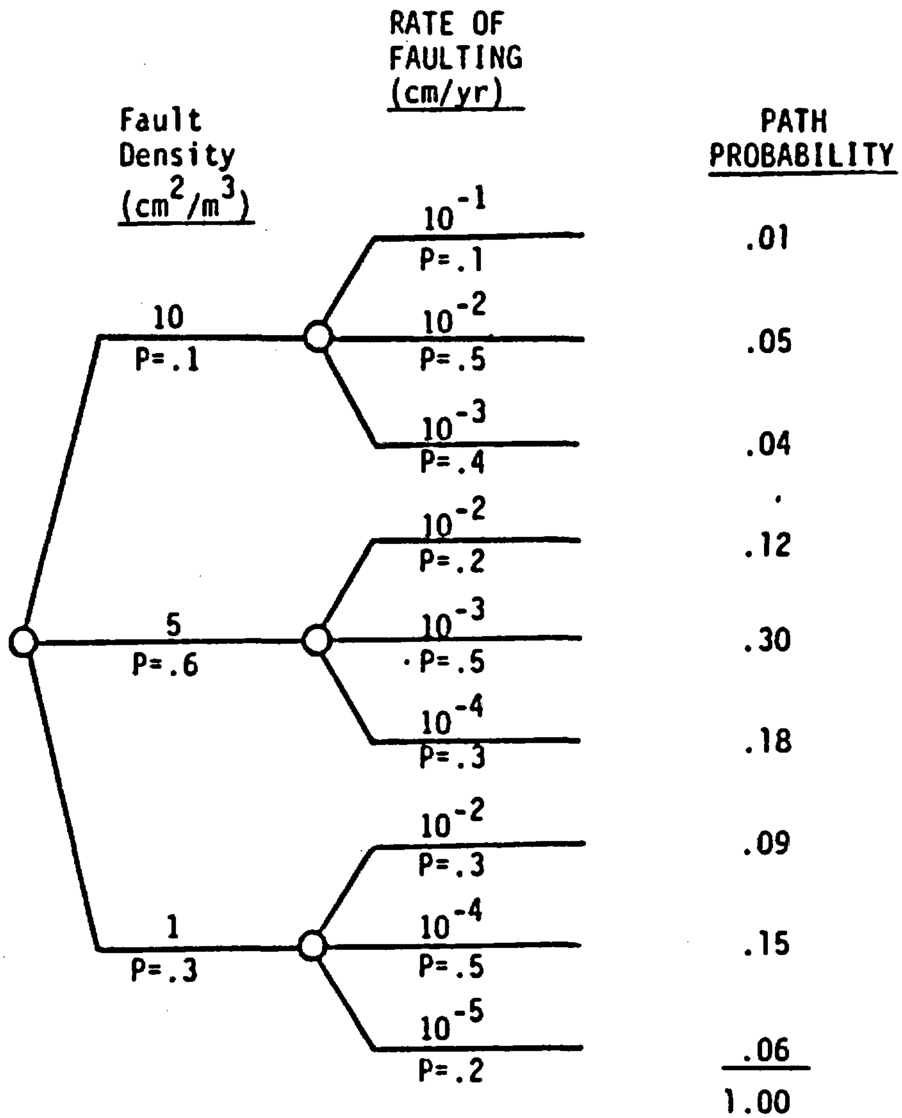

Figure 4. Probability tree illustrating probabilistic dependences.

III. DETAILED DESCRIPTION OF THE METHODOLOGY

This section describes and illustrates the steps required to apply the decision-aiding methodology being used by the DOE in the site-recommendation process. To simplify the presentation, a full theoretical justification of some of the steps has been omitted. Where such omissions occur, references to discussions of the theory in the literature are provided.

Step 1: Identify and Organize Objectives

The relative desirability of a candidate site is assumed to depend on the extent to which the selection of that site for recommendation would achieve the various objectives of site selection. Thus, the first step in the analysis is to explicitly identify siting objectives. These objectives are being generated iteratively, beginning with generic top-level objectives and proceeding with the various lower-level objectives that provide the means for achieving the higher-level objectives. The identification of objectives is based on the siting guidelines.

Objectives are being organized in a hierarchy to show the relationship between overall objectives and more-specific subobjectives. The process is being continued until specific technical guidelines or considerations represented within guidelines are identified. An illustration of a possible hierarchy of objectives is given in Figure 5, which shows "minimize impacts of the repository" as the overall objective and various lower-level subobjectives. Figure 5 will be used as the basis for generating examples for illustrating the remaining steps of the methodology. The reader should bear in mind, however, that the objectives hierarchy of Figure 5 is under revision and is provided for illustration only.

With the illustrative objectives hierarchy of Figure 5, the overall objective of minimizing the impacts of the repository (relative to the available and comparable siting options) is related to five lower-level objectives:

1. Maximize the protection of postclosure health and safety.
2. Maximize the protection of preclosure health and safety.
3. Minimize impacts on the environment.
4. Minimize adverse socioeconomic impacts.
5. Minimize economic costs.

The objectives dealing with postclosure and preclosure health and safety and the objective dealing with economic costs are divided further. For postclosure health and safety, three subobjectives are identified:

1. Minimize the health effects associated with nondisruptive geologic processes and events.

2. Minimize the health effects associated with disruptive geologic processes and events.

3. Minimize the health effects associated with human interference.

III. DETAILED DESCRIPTION OF THE METHODOLOGY

This section describes and illustrates the steps required to apply the decision-aiding methodology being used by the DOE in the site-recommendation process. To simplify the presentation, a full theoretical justification of some of the steps has been omitted. Where such omissions occur, references to discussions of the theory in the literature are provided.

Step 1: Identify and Organize Objectives

The relative desirability of a candidate site is assumed to depend on the extent to which the selection of that site for recommendation would achieve the various objectives of site selection. Thus, the first step in the analysis is to explicitly identify siting objectives. These objectives are being generated iteratively, beginning with generic top-level objectives and proceeding with the various lower-level objectives that provide the means for achieving the higher-level objectives. The identification of objectives is based largely on (1) the siting guidelines and other applicable regulations and (2) the knowledge and experience of the task team.

Objectives are being organized in a hierarchy to show the relationship between overall objectives and more-specific subobjectives. The process is being continued until specific technical guidelines or considerations represented within guidelines are identified. An illustration of a possible hierarchy of objectives is given in Figure 5, which shows "minimize impacts of the repository" as the overall objective and various lower-level subobjectives. Figure 5 will be used as the basis for generating examples for illustrating the remaining steps of the methodology. The reader should bear in mind, however, that the objectives hierarchy of Figure 5 is under revision and is provided for illustration only.

With the illustrative objectives hierarchy of Figure 5, the overall objective of minimizing the impacts of the repository is related to five lower-level objectives:

1. Maximize the protection of postclosure health and safety.
2. Maximize the protection of preclosure health and safety.
3. Minimize impacts on the environment.
4. Minimize adverse socioeconomic impacts.
5. Minimize economic costs.

The objectives dealing with postclosure and preclosure health and safety and the objective dealing with economic costs are divided further. For postclosure health and safety, three subobjectives are identified:

1. Minimize the health effects associated with nondisruptive geologic processes and events.

2. Minimize the health effects associated with disruptive geologic processes and events.

3. Minimize the health effects associated with human interference.

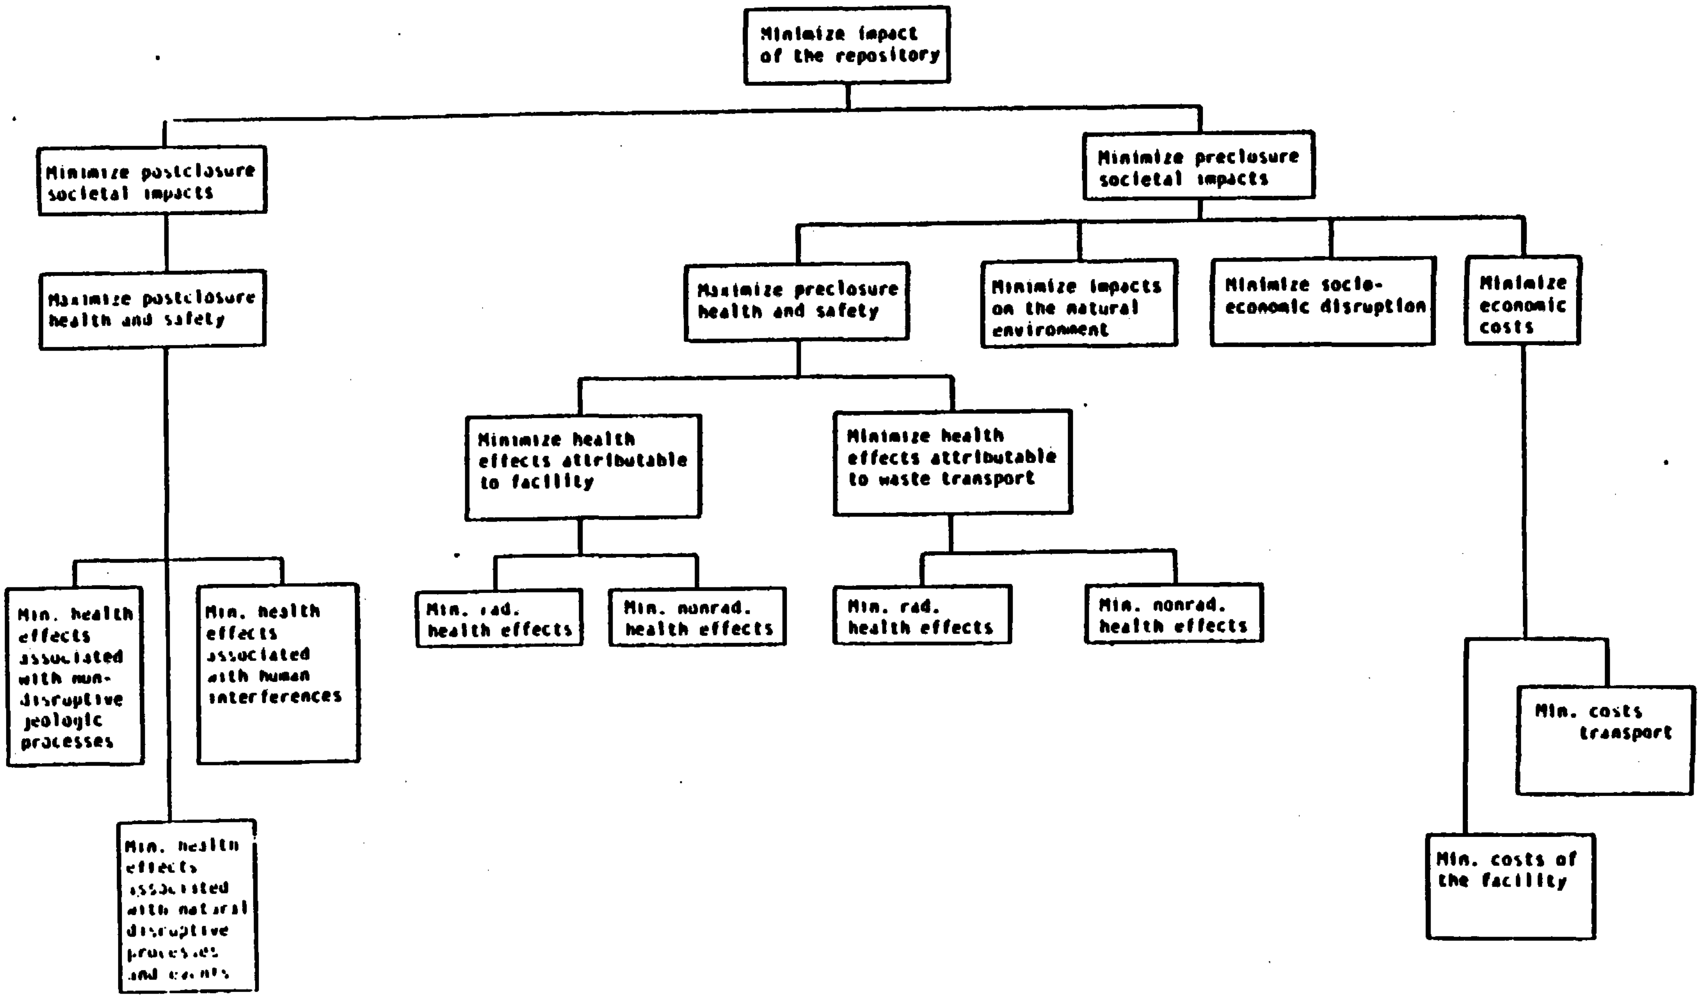

Figure 5. Objectives hierarchy showing various major and lower-level siting objectives.

For preclosure health and safety, four subobjectives are identified:

1. Minimize the health effects attributable to the repository.

2. Minimize the nonradiological health effects incurred by workers from the repository.

3. Minimize the radiological health effects attributable to waste transportation.

4. Minimize the nonradiological health effects attributable to waste transportation.

For costs, two subobjectives are identified:

1. Minimize the total economic costs associated with the repository.

2. Minimize the total economic costs associated with waste transportation.

Constructing a hierarchy of objectives, such as the example of Figure 5, aids the development of performance measures in several important ways. Performance measures need be defined only for the subobjectives at the bottom of the hierarchy. Because these lower-level subobjectives are more specific, it is easier to identify reasonable performance measures for them. Systematically constructing the hierarchy helps to ensure completeness and helps to eliminate situations where overcounting or undercounting might result (because omissions and redundancies should be fairly easily identified). The hierarchy puts the various subobjectives in perspective and provides a qualitative basis for screening out lesser concerns as not important to the overall goal.

The system guidelines provide a good starting point in developing the higher-level objectives. Most of the technical guidelines, however, cannot be directly used as subobjectives in the multiattribute-utility approach because of dependences among the guidelines. As the full hierarchy of objectives is being developed, it is being checked against the technical guidelines to ensure that all the objectives implied by the guidelines are included.

Step 2: Establish Performance Measures

The second step in the decision-aiding methodology is to establish performance measures for indicating how well each subobjective is met. Defining performance measures and their scales is essentially a creative process requiring professional judgment, knowledge, and experience. If the objectives hierarchy of Figure 5 were used, for example, three postclosure and eight preclosure performance measures would be needed. These might be denoted by the following symbols:

Symbol	Postclosure measure
x_1	Performance with respect to nondisruptive geologic processes
x_2	Performance with respect to disruptive geologic processes and events
x_3	Performance with respect to human interference
	Preclosure measure
y_1	Radiological safety of repository operation
y_2	Nonradiological safety of repository workers
y_3	Radiological safety of waste transportation
y_4	Nonradiological safety of waste transportation
y_5	Performance with respect to the natural environment
y_6	Performance with respect to socioeconomics
y_7	Performance with respect to repository costs
y_8	Performance with respect to transportation costs

To help establish the factors that must logically be represented by performance measures, influence diagrams are being constructed. An influence diagram is a directed graph displaying relationships (influences) among various factors (see, for example, Howard and Matheson, 1980, and Owen, 1978). The influence diagrams make explicit the relationship between the siting objectives and the guidelines (or considerations represented in the siting guidelines). Figures 6, 7, and 8 show sample influence diagrams for several of the siting objectives shown in Figure 5.

The process being used to construct influence diagrams involves both analysts and technical experts. Starting with a given siting objective--for example, minimize the postclosure public health effects resulting from nondisruptive geologic processes--the analyst asks the expert to identify the key variables whose values influence the degree to which this objective is met. In Figure 6, for example, the key variable is expected radionuclide releases to the accessible environment. Factors strongly influencing this variable are the effectiveness of the natural barriers and the effectiveness of the engineered barriers. These factors are in turn influenced by the pre-waste-emplacement characteristics of the host rock and the reactivity of the waste package and other engineered barriers. This filling-out process continues until all the factors on the bottom tiers can be readily assessed or until the point at which further decomposition is unlikely to facilitate assessment.

The bottom-tier factors basically determine the degree to which a particular objective is likely to be met. They represent considerations that are addressed by various technical guidelines. For example, Figure 6 shows that the guidelines on geohydrology, geochemistry, and rock characteristics (natural barriers) are of primary importance in determining the extent to which a site achieves postclosure subobjective x_1 (minimize the health effects due to nondisruptive geologic processes). Figure 7 shows that both these natural-barrier guidelines and the guidelines on climatic changes, dissolution, erosion, and tectonics influence the ability of a site to meet subobjective x_2 (minimize the health effects due to disruptive geologic processes and events). Figure 8 indicates that three groups of guidelines--those on natural barriers, disruptive geologic processes, and natural resources and site ownership and control--influence the achievement of subobjective x_3 (minimize the postclosure health effects due to human interference).

Because the influence diagrams indicate the factors that must logically be taken into account in judging the degree to which a site achieves each siting objective, they show the guidelines that are relevant to the various objectives and the logical relationships between the scores a site achieves on technical guidelines and the degree to which that site meets siting objectives. Coupled with the hierarchy of objectives, the influence diagrams help avoid overcounting and undercounting the importance of the various considerations represented in the guidelines because the logical significance of factors can be inferred from the relationships between these factors and the lower- and higher-level objectives that they influence. Figures 6, 7, and 8 show, for example, that considerations represented by the natural-barrier guidelines (rock characteristics, geochemistry, and geohydrology) have great importance because these considerations influence all three postclosure subobjectives.

After the construction of influence diagrams, it is necessary to specify the attributes that define the performance measures and the associated scales. Technical experts familiar with the objectives and goals of repository siting are undertaking the development of performance measures as a joint effort with analysts who are experienced in the development of such measures and knowledgeable in the role and purpose of performance measures in decision-aiding methodology. Careful attention is being given to establishing the performance measures because they serve as criteria for representing how well a particular site meets the objectives of the repository program. Care must be taken to ensure that, to the extent practicable, performance measures are complete (to cover all repository siting objectives), operational, nonredundant (to avoid doublecounting possible impacts), and minimal (to reduce the time and cost of their application). The influence diagrams show the basic site characteristics that must be logically reflected in the performance measures and provide the basis for relating a site's score on a performance measure to its scores on various guidelines.

In theory, performance measures can be either direct or indirect measures of objectives, and either natural or constructed scales can be used. Natural scales are established scales that enjoy common usage and interpretation. For instance, the objective to "minimize construction costs" might be associated with the direct performance measure of total costs. The appropriate natural

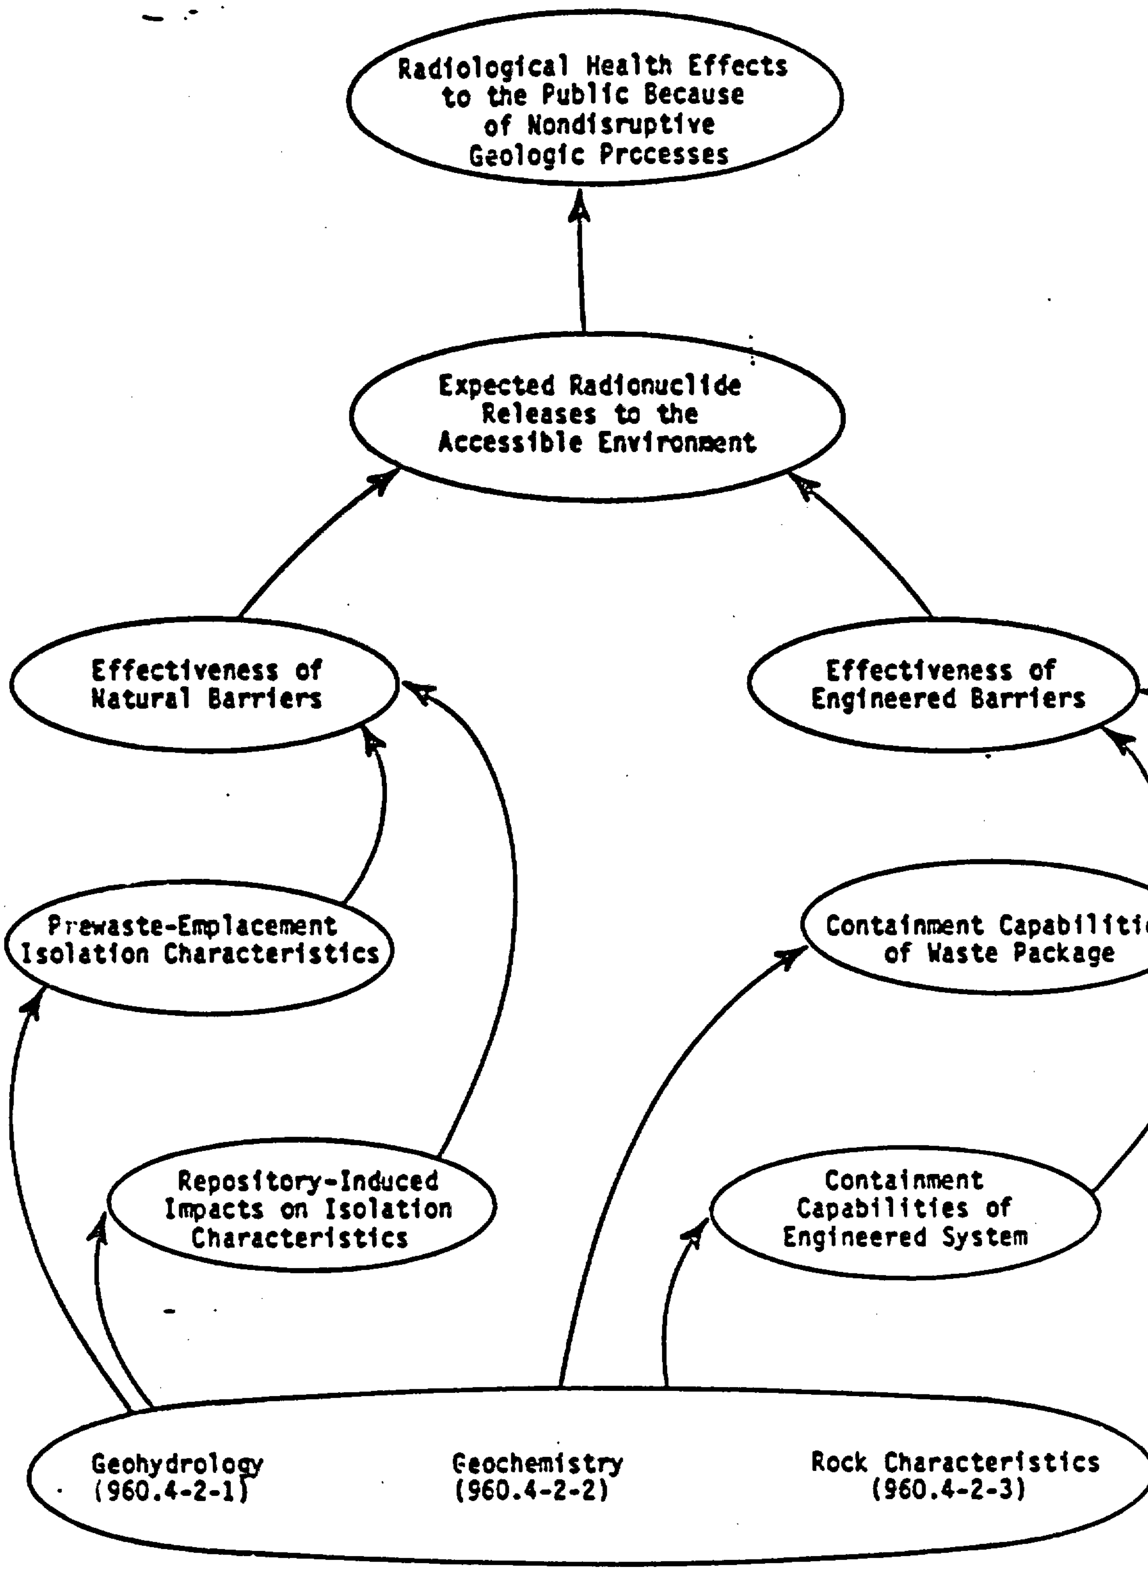

Figure 6. Factors influencing postclosure health effects due to nondisruptive geologic processes.

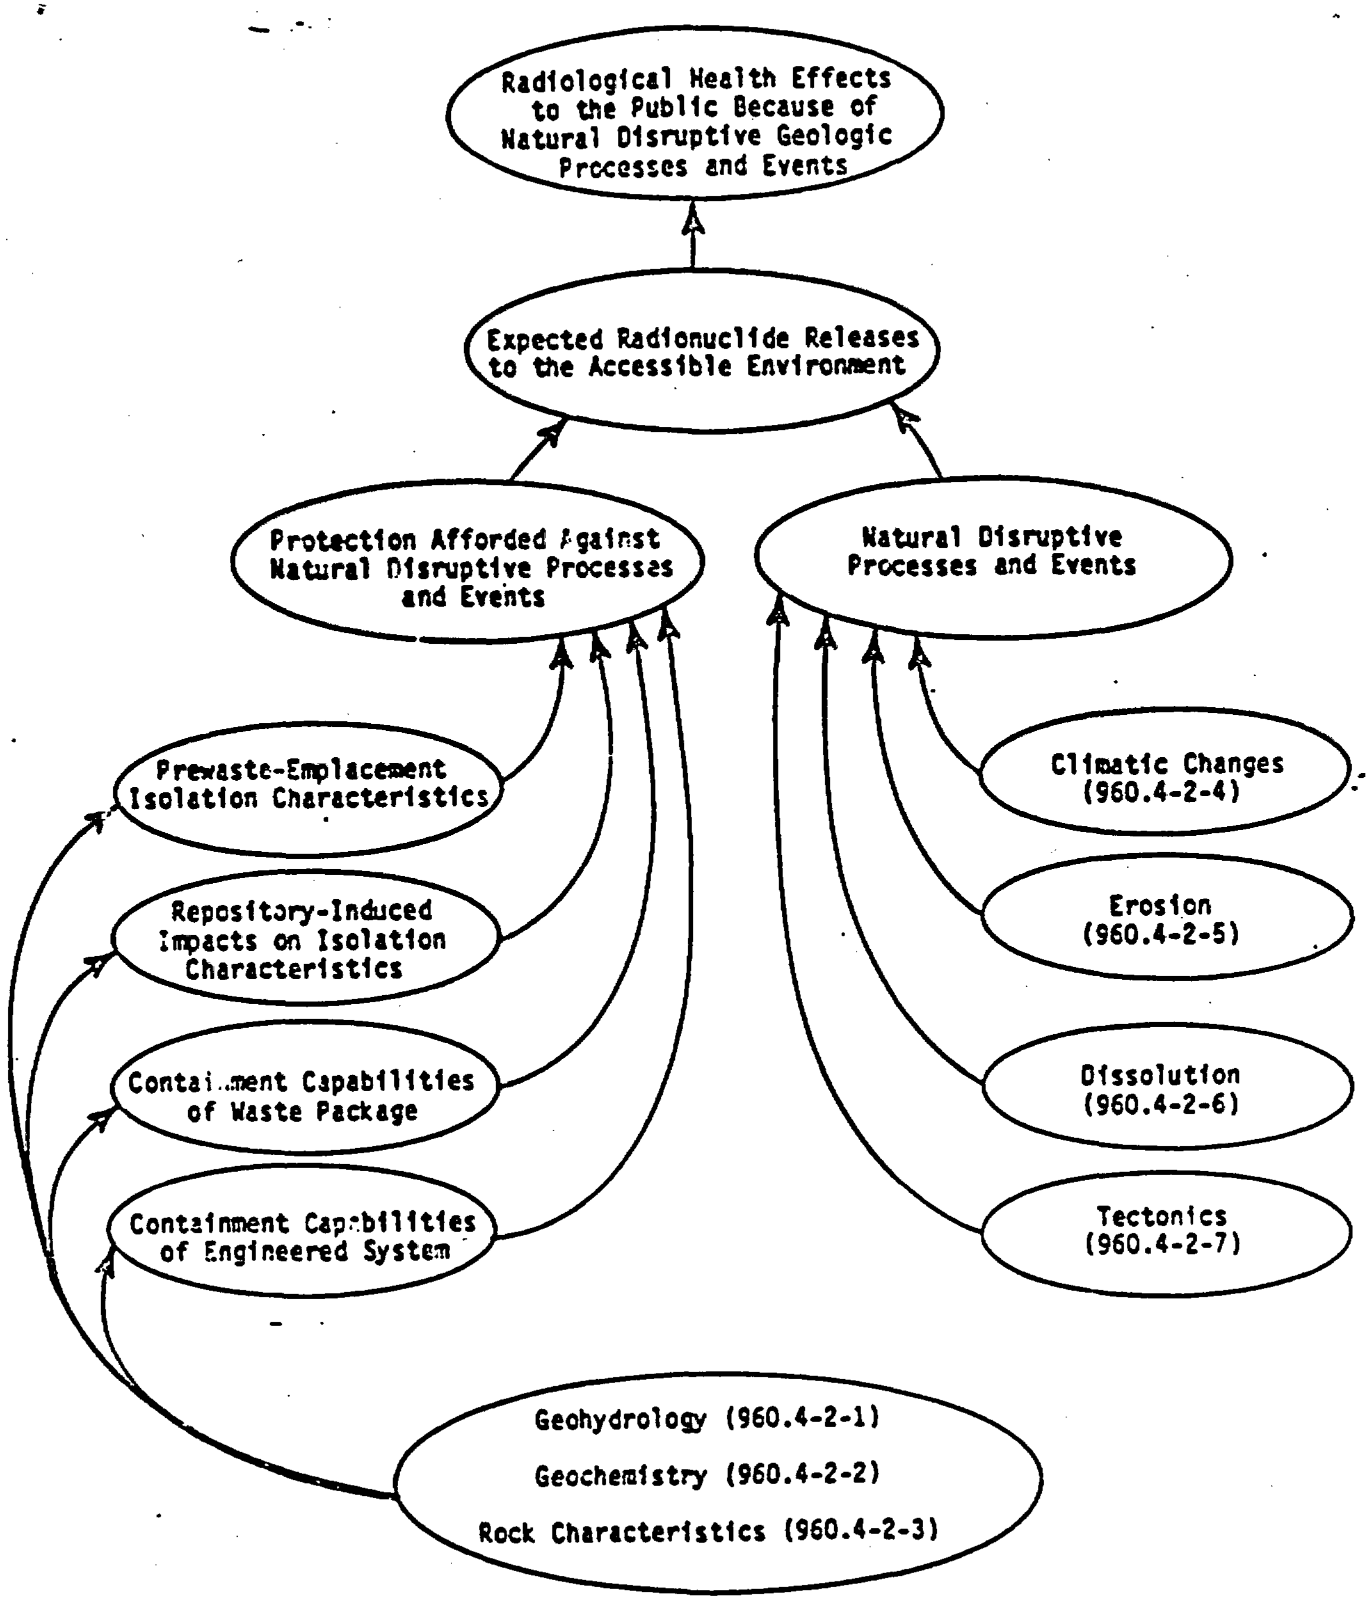

Figure 7. Factors influencing postclosure health effects due to disruptive geologic processes and events.

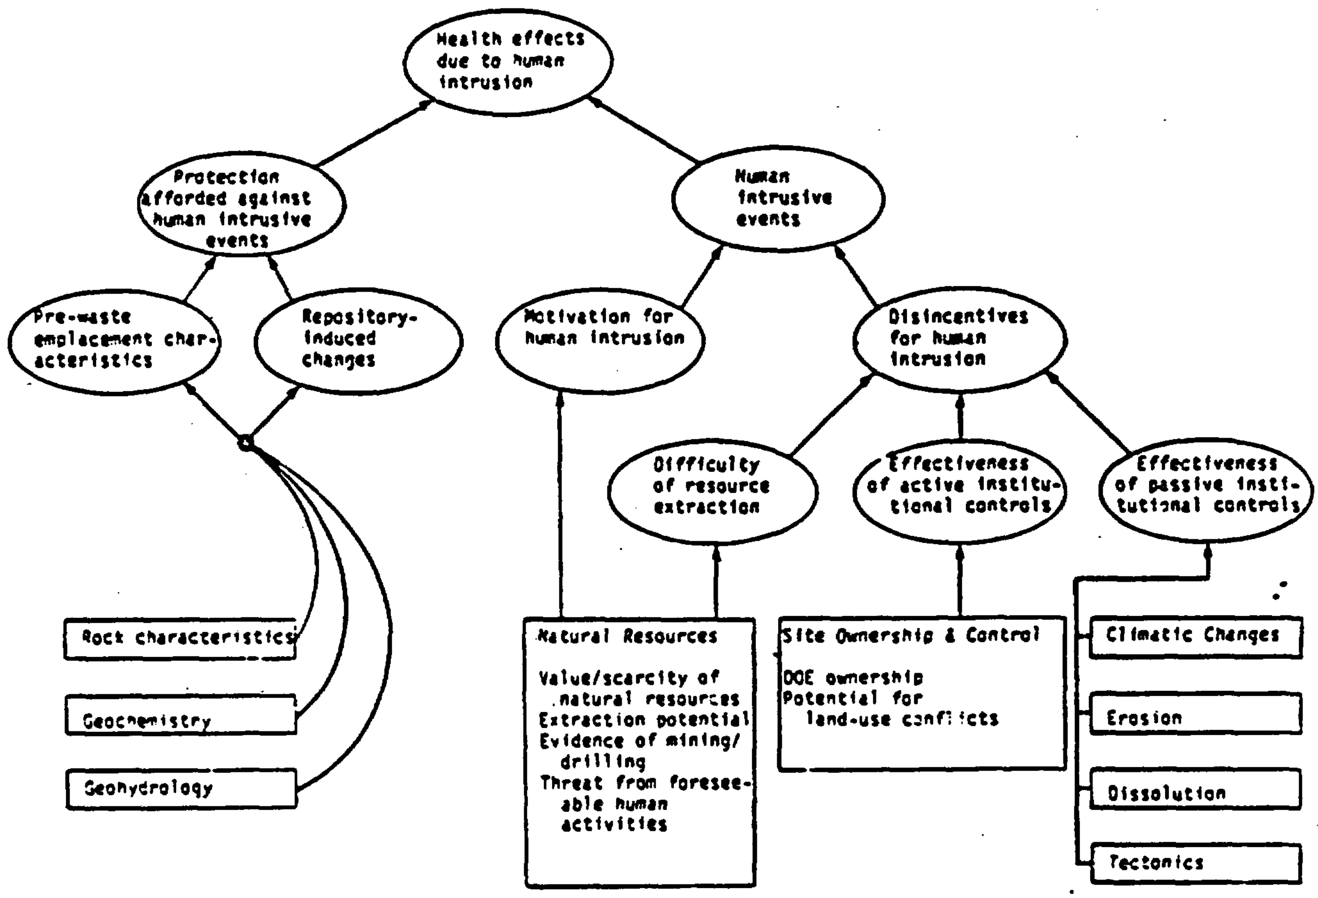

Figure 8. Factors influencing postclosure health effects due to human interference.

scale of measurement would be dollars. Constructed scales, on the other hand, are developed specifically for the problem at hand and are necessary when no natural scale of impact is available.

For maximum consistency with the aggregation method used in the draft EAs, constructed scales of 1 to 10 are being defined for each measure. These constructed scales are being defined in terms of either natural measures like dollars (e.g., a score of 4 on the performance measure "costs of the repository" (y_7) might mean that repository costs are estimated to be $6.33 billion) or in terms of collections of qualitative and quantitative descriptions (e.g., a score of 3 on the performance measure "environmental impact" (y_5) might mean "no significant conflicts with environmental requirements, but many environmental impacts, a few of which are difficult to mitigate").

Figure 9 shows sample definitions for two possible performance measures. In general, scores of 1 and 10 represent, respectively, the worst and the best levels of performance judged to be reasonably conceivable.

The performance measures in Figure 9 are described in terms of radiation releases to the accessible environment. Surrogates for these particular radiological performance measures and for others will be developed in terms of site characteristics traceable to individual technical guidelines. For example, a score of "1" on the performance measure "performance with respect to nondisruptive geologic processes" might represent a site with very short ground-water travel times and a complex geologic setting that could be extremely difficult to model (guideline on geohydrology), strongly oxidizing ground-water conditions and poor sorption characteristics (guideline on geochemistry), and thermal properties such that the heat generated by the waste could decrease the isolation provided by the host rock (guideline on rock characteristics), etc. Such steps are necessary because the data required to calculate reliably cumulative releases and release rates are not available before site characterization.

Step 3: Verify Independence Assumptions

As described in Section II, independence assumptions are necessary for an accurate overall evaluation of a site to be obtained by weighting and adding evaluations against distinct performance measures. The general approach for verifying the necessary independence assumptions is to consider special cases that would contradict the assumption. If none are found, independence is taken as a reasonable assumption.

One condition that permits the additive form to be valid is that the performance measures are "additive independent" of one another. Performance measures Z_1 and Z_2 are said to be additive independent if the "preference order for lotteries (gambles in which possible values for Z_1 and Z_2 occur with specified probabilities) does not depend on the joint probability distributions of these lotteries, but depends only on their marginal probability distributions" (Keeney, 1980, page 231). To illustrate this condition in more concrete terms, suppose that Z_1 and Z_2 are performance measures representing environmental impacts and economic costs, respectively, and suppose that there are two possible lotteries that are compared. The

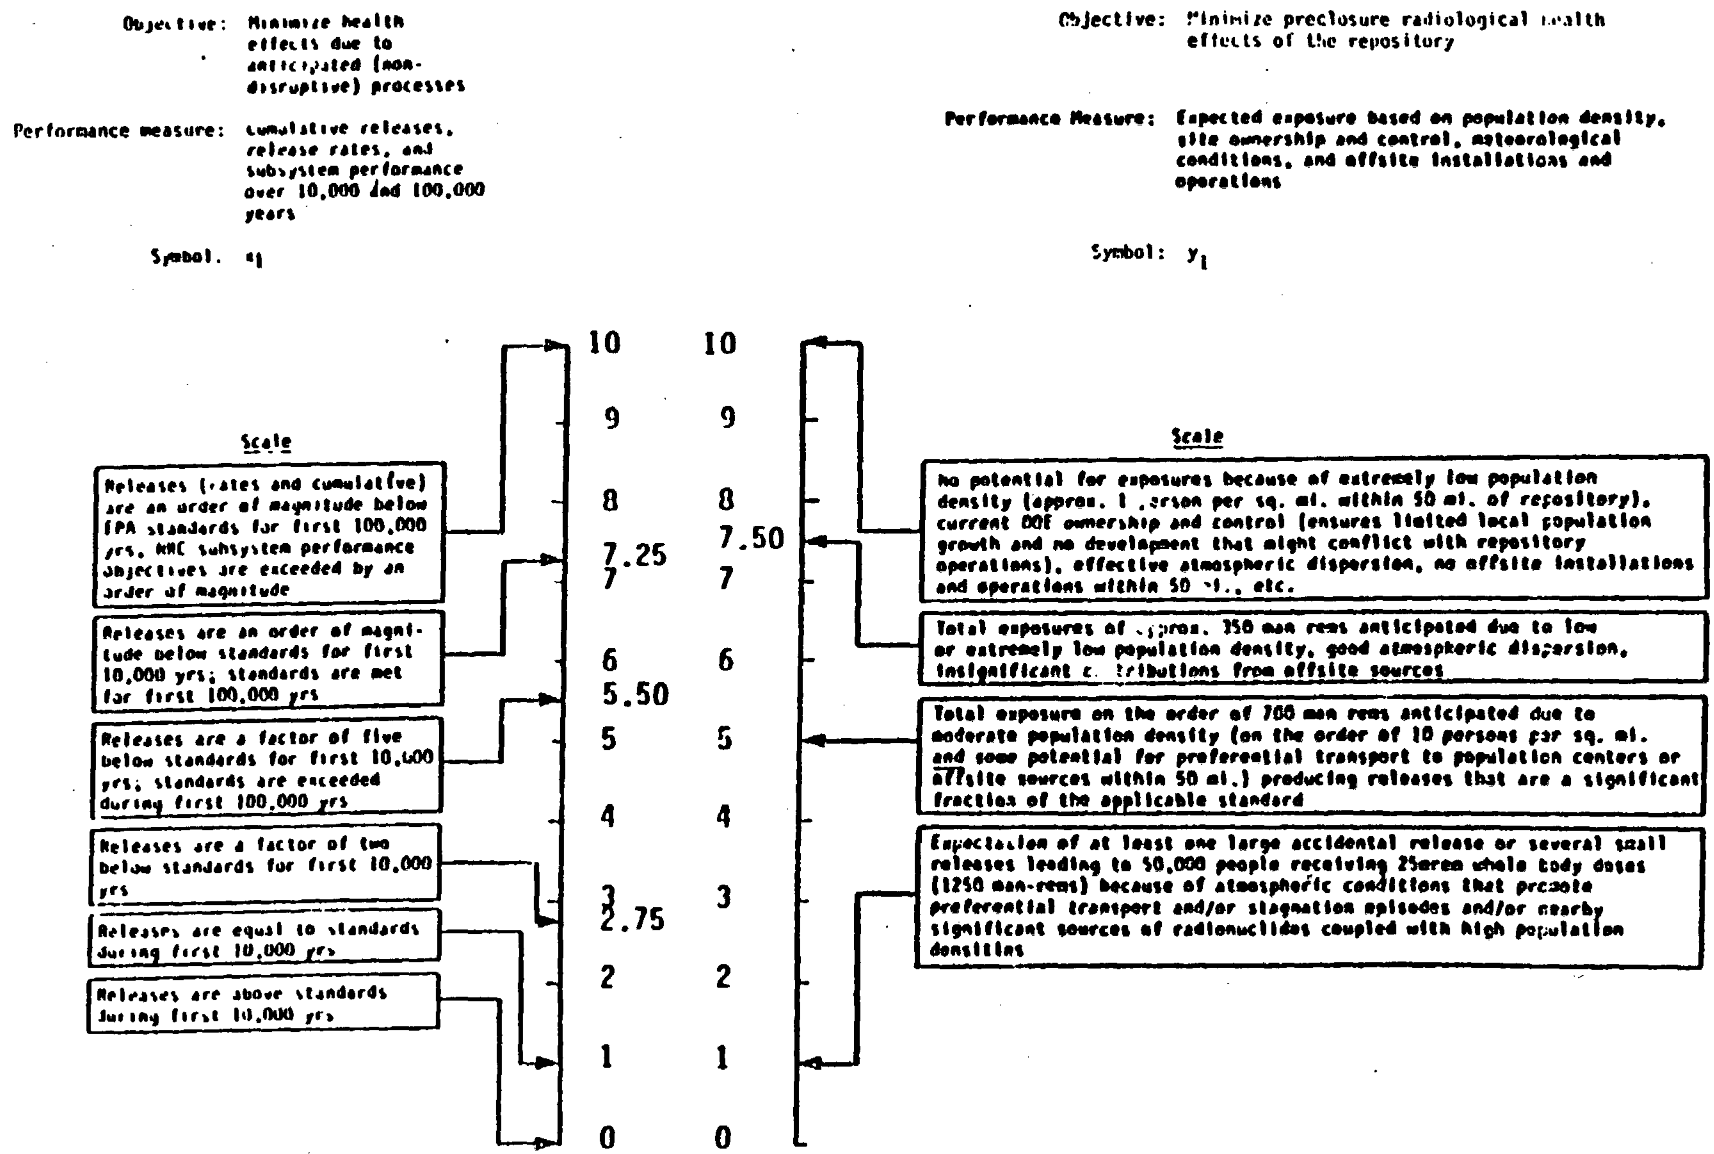

Figure 9. Sample performance measures.

first yields equal chances for the favorable outcome (Z_1 low, Z_2 low) and the unfavorable outcome (Z_1 high, Z_2 high). The second yields equal chances for the mixed outcomes (Z_1 low, Z_1 high) and (Z_1 high, Z_1 low). Note that both lotteries have an equal (namely, 0.5) chance at either (Z_1 low, Z_1 high) and that both also have an equal 0.5 chance at (Z_2 low, Z_2 high). Both lotteries are therefore said to have the same marginal probability distributions.

If Z_1 and Z_2 are additive independent, then one must be indifferent between the first lottery and the second.

Assuming additive independence among all performance measures, it is possible to express a site's postclosure utility, denoted U_{post}, by an additive equation. For example, if there are three postclosure-performance measures, x_1, x_2, and x_3, then

$$U_{post} = w_1U_1(x_1) + w_2U_2(x_2) + w_3U_3(x_3) \quad (1)$$

where w_1, w_2, and w_3 are weights (scaling factors) and U_1, U_2, and U_3 are single-attribute utility functions defined over the respective performance measures x_1, x_2, and x_3. Similarly, if there are 11 preclosure-performance measures, then the preclosure utility of a site can be computed from an additive equation of the form

$$U_{pre} = k_1V_1(y_1) + k_2V_2(y_2) + k_3V_3(y_3) + k_4V_4(y_4) + k_5V_5(y_5) + k_6V_6(y_6) + k_7V_7(y_7) + k_8V_8(y_8) \quad (2)$$

where k_1 through k_8 are weights and V_1 through V_8 are single-attribute utility functions defined over the preclosure-performance measures y_1 through y_8, respectively. The overall utility is then given by

$$U_{overall} = k_{post}\,U_{post} + k_{pre}\,U_{pre} \quad (3)$$

Step 4. Assess Single-Attribute Utility Functions

Performance measures are important proxies for determining how well a site meets a particular objective. However, by themselves, these measures do not quantify performance against a particular objective. For example, it does not follow that an objective is 90 percent met just because the level of performance is 90 percent of its maximum value (i.e., the site is assigned a score of 9). Depending on the objective, it might be, for example, that most of the intent of the objective is met when the performance measure reaches only 20 percent of its maximum possible value (i.e., achieves a score of 2). Therefore, a scale is needed to represent the relative desirability of achieving different scores for the performance measures. The concept of a single-attribute (marginal) utility function provides such a scale. As noted in Section II, an extensive literature has been developed on the meaning and uses of utility functions.

Simply stated, a utility function is a mathematical expression for the subjective tradeoffs that are inherent in any judgment that one site is better than another for a repository. Logically, the values that are represented in a utility function should be those of the decisionmaker--in this case the DOE. The DOE will incorporate as appropriate the values of others in the value structure. For example, public comments on the weighting allocations among guideline sets and groups presented in Chapter 7 of the draft EAs will be considered. Methods for accomplishing this integration are discussed by Keeney and Raiffa (1976) and Keeney (1980).

Marginal utility functions that reflect the preferences of an individual can be derived by assessing a few points on the function corresponding to various values of the performance measure and then fitting a smooth curve. Using techniques recommended in the decision-analysis literature, decision analysts experienced in utility assessment are constructing utility functions in interviews with DOE management, staff, and consultants. For example, a technique being used to assess the single-attribute utility function U_1 is the midpoint method (Changkong and Haimes, 1983). This procedure involves successively identifying levels of performance whose utilities (desirabilities) seem to be halfway between already established utilities. To illustrate, consider a utility function for measuring performance with respect to nondisruptive processes (x_1). Arbitrary utilities of 0 and 100 may be assigned to performance levels for x_1 of 1 and 10, respectively. Various intermediate performance levels are then selected until a level, denoted x', is found such that it is judged to be equally desirable to change a site whose performance level is $x_1 = 1$ to the level $x_1 = x'$ as it would be to change a site whose performance level is $x_1 = x'$ to the level $x_1 = 10$. The resulting performance level is called "the midpoint" because the utility function evaluated at this point is midway between the utilities of the other two outcome levels that were considered. This same process is repeated to find other midpoints (e.g., the midpoint between $x_1 = 1$ and $x_1 = x'$) until enough are identified to permit fitting a smooth curve. A sample utility curve for U_1 is shown in Figure 10. For comparison, a sample utility curve for U_2 is shown in Figure 11.

Step 5: Assess Scaling Factors or Weights

The constants $w_1, w_2, \ldots, k_1, k_2, \ldots, k_{pre}$, and k_{post} in Equations 1, 2, and 3 represent scaling factors or weights designed to account for the relative value of trading off performance on one performance measure for another. The scaling factor assigned to a given performance measure defines the increment of overall utility associated with increasing that measure's performance outcome from a score of 1 to a score of 10. Clearly, the scaling factor must depend on the definitions of "1" and "10," which, as described in step 2, must be consistent with the siting guidelines. In other words, the scaling factors must be consistent with the definitions established for performance-measure scores.

As outlined in Section II, the method generally recommended for establishing scaling factors that reflect preferences is to fix all but two of the performance measures and then to allow these two to vary, in order to find combinations that the policymaker finds equally preferable. In this case, the multiattribute utilities will be equal by definition, and therefore it is possible to generate equations in which the weights are unknowns. The solution of these equations then yields the values for the weights.

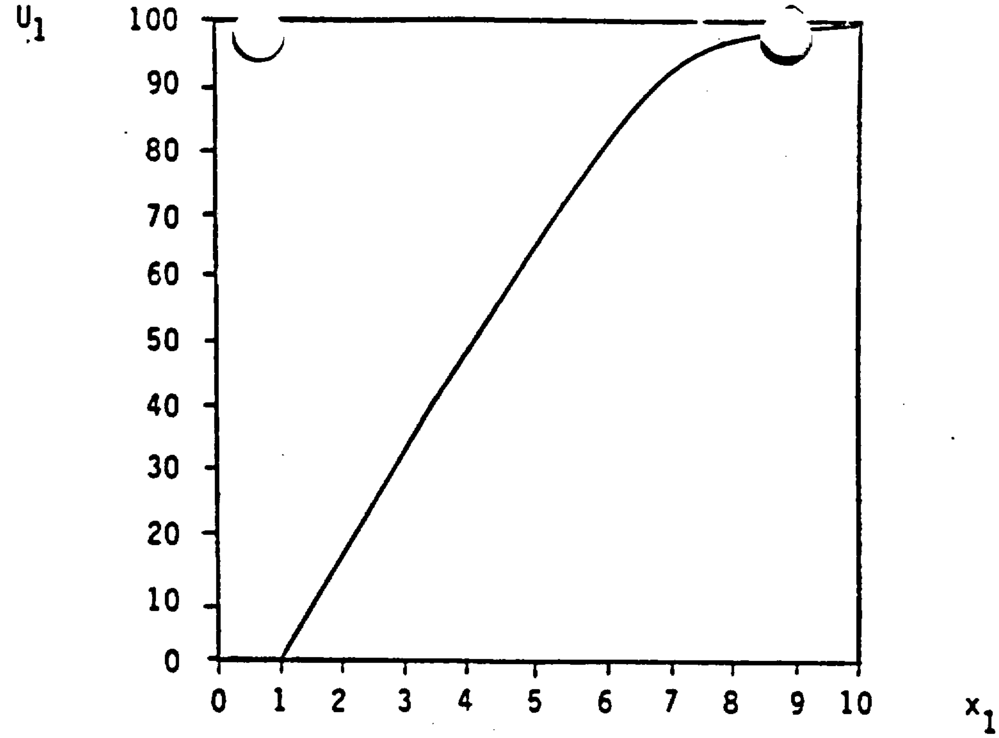

Figure 10. Sample single-attribute utility curve for postclosure performance with respect to nondisruptive events.

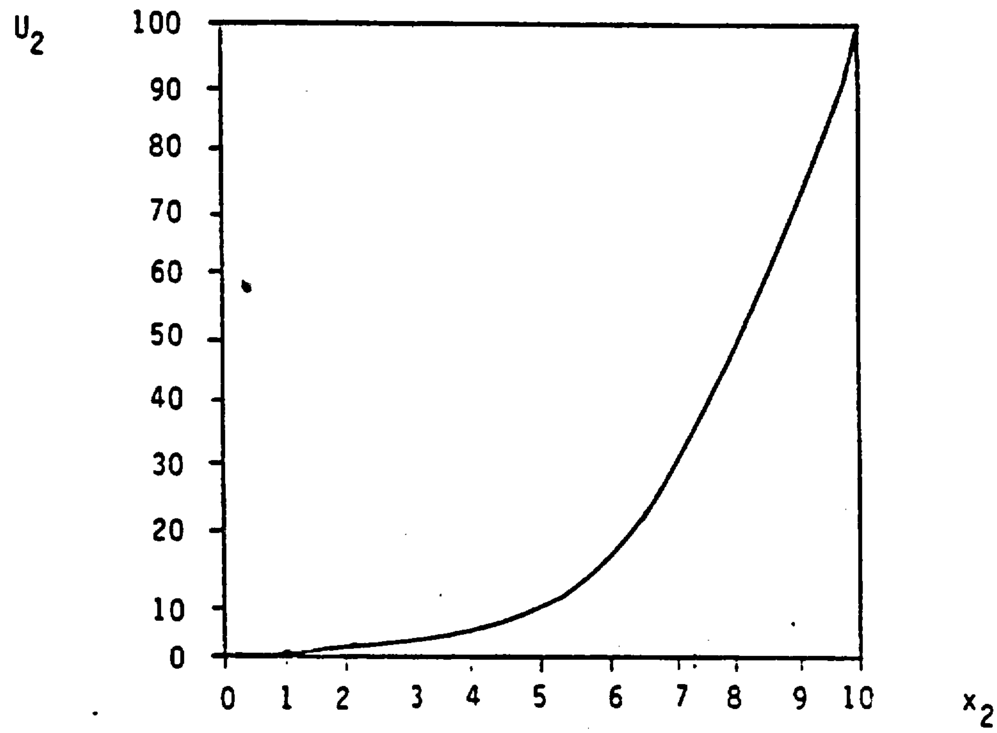

Figure 11. Single-attribute utility curve for performance under disruptive geologic events and processes.

To illustrate the methods, consider preferences for trading off performance between performance measures x_1 and x_2. As shown in the example below, different radionuclide-release scenarios may be considered until two are found that are regarded as equally undesirable:

	Site A		Site B	
x_1	10	(Releases from nondisruptive processes are 10 times lower than the standard during the first 100,000 years)	1	(Releases from nondisruptive processes are equal to the standard during the first 10,000 years)
x_2	1	(Releases from disruptive geologic events are 10 times higher than the allowable releases for the first 10,000 years)	4	(Releases from disruptive geologic events are three times higher than the allowable releases for the first 10,000 years)
x_3	1		1	

From Equation 1 and Figures 10 and 11 (and the fact that utilities are defined to equal 0 and 100 for scores of 1 and 10, respectively), the postclosure utility of site A is

$$U^{A}_{post} = w_1U_1(10) + w_2U_2(1) + w_3U_3(1) = 100w_1$$

Similarly, the postclosure utility of site B is

$$U^{B}_{post} = w_1U_1(1) + w_2U_2(4) + w_3U_3(1) = 5w_2$$

Because indifference between point A and point B implies equal utility,

$$100w_1 = 5w_2$$

To obtain additional relationships among the weights, other tradeoffs among various levels of performance measures must be considered.

As mentioned previously, in the case of preclosure, the scaling factors are partially constrained by the requirements of the siting guidelines. The guidelines specify that the order of importance for the three preclosure-guideline groups, from greater to lesser importance, is (1) preclosure radiological safety; (2) environment, socioeconomics, and transportation; and (3) ease and cost of siting, construction, operation, and closure. Suppose the correspondence between performance measures and preclosure-guideline groups were as follows:

Guideline group	Performance measure
Preclosure radiological safety (repository)	Radiological safety of repository operation (y_1)
Environment, socioeconomics, and transportation	Radiological safety of waste transportation (y_3)
	Nonradiological safety of waste transportation (y_4)
	Performance with respect to the natural environment (y_5)
	Performance with respect to socioeconomics (y_6)
	Performance with respect to transportation costs (y_8)
Ease and cost of siting, construction, operation, and closure	Nonradiological safety of repository workers (y_2)
	Performance with respect to repository costs (y_7)

The relative-importance stipulation in the guidelines is interpreted as requiring that the total weight given to the utility of performance for measures associated with preclosure radiological safety must be greater than the total weight given to the utility of performance for measures associated with the environment, socioeconomics, and transportation. Similarly, the total weight given to the utility of performance for the environment, socioeconomics, and transportation must be greater than the total weight given to the utility associated with the ease and cost of siting, construction, operation, and closure. Thus,

$$k_1 > k_3 + k_4 + k_5 + k_6 + k_8 > k_2 + k_7 \qquad (4)$$

The approach for generating the scaling factors consists of deriving tentative values, using methods similar to that described above, and then checking whether those values satisfy the above equation. In all cases, the tradeoff judgments are being provided by DOE management and staff most familiar with repository-siting objectives and are chosen, wherever possible, so as to be consistent with tradeoffs established by other social decisions. To the extent that judgmental value tradeoffs produce scaling factors that violate Equation 4, these tradeoffs are adjusted until consistency with Equation 4 is obtained.

Step 6: Assign Site Performance Scores, Compute Utilities, and Perform Sensitivity Analysis

After the development of single-attribute utility functions and nominal scaling factors, Equations 1, 2, and 3 are applied to compute preclosure, postclosure, and overall utilities for each site. Sensitivity studies are then undertaken to identify critical numerical assumptions and the sensitivity of the overall utilities to these assumptions.

The information contained in the final EAs is being used to summarize the expected performance of each site by estimating appropriate values for the performance measures established in step 2. In the absence of complete models for simulating site performance, performance-measure scores are being obtained as judgments provided by panels of experts. The scores assigned by each panel must be consistent with the definition of the performance-measure scales and must logically account for all characteristics of the site represented in the associated influence diagram. If there is substantial uncertainty about the value of a performance measure for a given site, alternative scores may be specified with associated probabilities.

For an example of how utilities are being computed, consider the evaluation of overall postclosure utilities. Given the example used throughout this section, and assuming that independence is verified in step 3, the multiattribute utility theory suggests that a measure of postclosure performance that takes into account nondisruptive geologic processes, disruptive geologic events, and human interference can be obtained by using Equation 1 to calculate the expected utility. Mathematically, the calculation of expected utility can be expressed as

$$E(U_{post}) = \int [w_1U_1(x_1) + w_2U_2(x_2) + w_3U_3(x_3)]\ dP \qquad (5)$$

where the symbols $\int$ and dP denote the process of computing all possible performance outcomes, computing the resulting utility values, weighting these values by their probabilities, and taking the resulting weighted average.

To simplify the application of Equation 5, it might be assumed that there is no significant uncertainty in the specification of the performance outcome x_1 for a site. Furthermore, uncertainty in the specification of performance outcomes x_2 and x_3 might be assumed to be due only to uncertainty in the occurrence of disruptive geologic events and human interference. The occurrences of disruptive geologic events and human interference might be assumed to be probabilistically independent. With these assumptions, Equation 5 can be expressed as

$$E(U_{post}) = w_1U_1(x_1) + w_2\int_{S_2} U_2[x_2(S_2)]p_2(S_2)dS_2 + w_3\int_{S_3} U_3[<_3(S_3)]p_3(S_3)dS_3 \qquad (6)$$

where $x_2(S_2)$ represents the performance outcome with respect to disruptive geologic events given a disruptive-event scenario S_2; $x_3(S_3)$ is the performance outcome with respect to human interference given a human-interference scenario S_3; $p_2(S_2)$ is a probability density function describing the likelihood of various disruptive-event scenarios; and $p_3(S_3)$ is a probability density function describing the likelihood of various human-interference scenarios.

Similarly, the expected utility of preclosure, assuming 8 preclosure performance measures, as in the previous examples, would be given by

$$E(U_{pre}) = \int [k_1V_1(y_1) + k_2V_2(y_2) + k_3V_3(y_3) + k_4V_4(y_4) + k_5V_5(y_5) + k_6V_6(y_6) + k_7V_7(y_7) + k_8V_8(y_8)]dP \quad (7)$$

If there is no significant uncertainty in the assignment of performance scores, Equation 2 could be used directly to compute preclosure utilities.

The single-attribute utility scores and associated probabilities assessed for each siting objective are being aggregated to obtain an overall expected utility and associated probability distribution on utility summarizing overall site attractiveness.

The output of this final step for each site will be a point estimate if there is little uncertainty about the performance-measure scores that represent the ultimate attractiveness (total utility) of the site. Alternatively, the final results could be presented as probability distributions, which would permit both the expected values and the uncertainty in the values to be compared among sites.

Sensitivity studies will be performed to explore the effect of changing assumptions and differences of opinion. For example, significant differences in the utility functions assessed by different individuals can be organized, and a sensitivity analysis can be used to determine the extent to which such differences alter the relative evaluation of sites.

Different weights representing a range of different views will be developed. In particular, a range of postclosure versus preclosure weights, consistent with an assumption that postclosure be assigned greater importance than preclosure, will be considered. In addition, the weighting relationship among the three preclosure-guideline groups will be varied, again consistent with the siting guidelines (see the discussion of step 5). The significance of these differing opinions will be investigated through sensitivity analyses. An important advantage of the decision-aiding methodology is that extensive sensitivity analyses representing differing value judgments can be developed quickly and inexpensively. This ability to answer many "what if" questions decreases the likelihood that inappropriate values will be used in the decision process and increases the likelihood that the most advantageous group of sites will be identified and recommended for characterization.

REFERENCES

Changkong, V., and Y. Y. Haimes, 1983. Multiobjective Decision Making: Theory and Methodology, North-Holland, New York.

Edwards, W., and J. R. Newman, 1982. Multiattribute Evaluation, Sage University, Beverly Hills, California.

Hobbs, B. F., 1982. Analytical Multiobjective Decision Methods for Power Plant Siting: A Review of Theory and Applications, Brookhaven National Laboratory, Upton, New York.

Holloway, C., 1979. Decision Making under Uncertainty: Models and Choices, Prentice-Hall, New York.

Howard, R. A., and J. E. Matheson, 1980. "Influence Diagrams," Decision Analysis Department, SRI International, Menlo Park, California.

Keeney, R. L., 1980. Siting Energy Facilities, Academic Press, New York.

Keeney, R. L., and H. Raiffa, 1976. Decisions with Multiple Objectives, Wiley, New York.

Merkhofer, M. W., 1983. "A Comparative Evaluation of Quantitative Decision-Making Approaches," SRI International, Menlo Park, California.

Merkhofer, M. W., in press. Decision Science and Social Risk Management: A Comparison of Cost-Benefit Analysis, Decision Analysis, and Other Decision-Aiding Approaches, D. Reidel, Dordrecht, the Netherlands.

Owen, D., 1978. "The Use of Influence Diagrams in Structuring Complex Decision Problems," in Proceedings of the Second Lawrence Symposium on Systems and Decision Sciences (reproduced in Applied Decision Analysis, by Derek Bunn, McGraw-Hill, 1984).

Spetzler, C. S., and C.-A. S. Stael von Holstein, 1975. "Probability Encoding in Decision Analysis," Management Science, Vol. 22, pp. 340-358.