NUREG/CR-5906
SAND92-1547

Decision Making Under Uncertainty: An Investigation Into the Application of Formal Decision-Making Methods to Safety Issue Decisions

Prepared by
Michael P. Bohn

Sandia National Laboratories

Prepared for
U.S. Nuclear Regulatory Commission

9301080160 921231
PDR NUREG
CR-5906 R PDR

AVAILABILITY NOTICE

Availability of Reference Materials Cited in NRC Publications

Most documents cited in NRC publications will be available from one of the following sources:

1. The NRC Public Document Room, 2120 L Street, NW., Lower Level, Washington, DC 20555

2. The Superintendent of Documents, U.S. Government Printing Office, P.O. Box 37082, Washington, DC 20013-7082

3. The National Technical Information Service, Springfield, VA 22161

Although the listing that follows represents the majority of documents cited in NRC publications, it is not intended to be exhaustive.

Referenced documents available for inspection and copying for a fee from the NRC Public Document Room include NRC correspondence and internal NRC memoranda; NRC bulletins, circulars, information notices, inspection and investigation notices; licensee event reports; vendor reports and correspondence; Commission papers; and applicant and licensee documents and correspondence.

The following documents in the NUREG series are available for purchase from the GPO Sales Program: formal NRC staff and contractor reports, NRC-sponsored conference proceedings, international agreement reports, grant publications, and NRC booklets and brochures. Also available are regulatory guides, NRC regulations in the *Code of Federal Regulations*, and *Nuclear Regulatory Commission Issuances*.

Documents available from the National Technical Information Service include NUREG-series reports and technical reports prepared by other Federal agencies and reports prepared by the Atomic Energy Commission, forerunner agency to the Nuclear Regulatory Commission.

Documents available from public and special technical libraries include all open literature items, such as books, journal articles, and transactions. *Federal Register* notices, Federal and State legislation, and congressional reports can usually be obtained from these libraries.

Documents such as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings are available for purchase from the organization sponsoring the publication cited.

Single copies of NRC draft reports are available free, to the extent of supply, upon written request to the Office of Administration, Distribution and Mail Services Section, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at the NRC Library, 7920 Norfolk Avenue, Bethesda, Maryland, for use by the public. Codes and standards are usually copyrighted and may be purchased from the originating organization or, if they are American National Standards, from the American National Standards Institute, 1430 Broadway, New York, NY 10018.

DISCLAIMER NOTICE

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, or any of their employees, makes any warranty, expressed or implied, or assumes any legal liability of responsibility for any third party's use, or the results of such use, of any information, apparatus, product or process disclosed in this report, or represents that its use by such third party would not infringe privately owned rights.

NUREG/CR-5906
SAND92-1547

Decision Making Under Uncertainty: An Investigation Into the Application of Formal Decision-Making Methods to Safety Issue Decisions

Prepared by
Michael P. Bohn

Sandia National Laboratories

Prepared for
U.S. Nuclear Regulatory Commission

9301080160 921231
PDR NUREG
CR-5906 R PDR

AVAILABILITY NOTICE

Availability of Reference Materials Cited in NRC Publications

Most documents cited in NRC publications will be available from one of the following sources:

1. The NRC Public Document Room, 2120 L Street, NW., Lower Level, Washington, DC 20555

2. The Superintendent of Documents, U.S. Government Printing Office, P.O. Box 37082, Washington, DC 20013-7082

3. The National Technical Information Service, Springfield, VA 22161

Although the listing that follows represents the majority of documents cited in NRC publications, it is not intended to be exhaustive.

Referenced documents available for inspection and copying for a fee from the NRC Public Document Room include NRC correspondence and internal NRC memoranda; NRC bulletins, circulars, information notices, inspection and investigation notices; licensee event reports; vendor reports and correspondence; Commission papers; and applicant and licensee documents and correspondence.

The following documents in the NUREG series are available for purchase from the GPO Sales Program: formal NRC staff and contractor reports, NRC-sponsored conference proceedings, international agreement reports, grant publications, and NRC booklets and brochures. Also available are regulatory guides, NRC regulations in the *Code of Federal Regulations*, and *Nuclear Regulatory Commission Issuances*.

Documents available from the National Technical Information Service include NUREG-series reports and technical reports prepared by other Federal agencies and reports prepared by the Atomic Energy Commission, forerunner agency to the Nuclear Regulatory Commission.

Documents available from public and special technical libraries include all open literature items, such as books, journal articles, and transactions. *Federal Register* notices, Federal and State legislation, and congressional reports can usually be obtained from these libraries.

Documents such as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings are available for purchase from the organization sponsoring the publication cited.

Single copies of NRC draft reports are available free, to the extent of supply, upon written request to the Office of Administration, Distribution and Mail Services Section, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at the NRC Library, 7920 Norfolk Avenue, Bethesda, Maryland, for use by the public. Codes and standards are usually copyrighted and may be purchased from the originating organization or, if they are American National Standards, from the American National Standards Institute, 1430 Broadway, New York, NY 10018.

DISCLAIMER NOTICE

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, or any of their employees, makes any warranty, expressed or implied, or assumes any legal liability of responsibility for any third party's use, or the results of such use, of any information, apparatus, product or process disclosed in this report, or represents that its use by such third party would not infringe privately owned rights.

NUREG/CR-5906
SAND92-1547
RV

Decision Making Under Uncertainty: An Investigation Into the Application of Formal Decision-Making Methods to Safety Issue Decisions

Manuscript Completed: November 1992
Date Published: December 1992

Prepared by
Michael P. Bohn

Sandia National Laboratories
Albuquerque, NM 87185

Prepared for
Division of Safety Issues Resolution
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555
NRC FIN L1334

ABSTRACT

As part of the NRC-sponsored program to study the implications of Generic Issue 57, "Effects of Fire Protection System Actuation on Safety-Related Equipment," a subtask was performed to evaluate the applicability of formal decision analysis methods to generic issues cost/benefit-type decisions and to apply these methods to the GI-57 results. In this report, the numerical results obtained from the analysis of three plants (two PWRs and one BWR) as developed in the technical resolution program for GI-57 were studied. For each plant, these results included a calculation of the person-REM averted due to various accident scenarios and various proposed modifications to mitigate the accident scenarios identified. These results were recomputed to break out the benefit in terms of contributions due to random event scenarios, fire event scenarios, and seismic event scenarios. Furthermore, the benefits associated with risk (in terms of person-REM) averted from earthquakes at three different seismic ground motion levels were separately considered. (This was done to allow differentiation between potential retrofits that affect the risk arising primarily from earthquake levels below the design basis earthquake from those potential retrofits that affect risk resulting primarily from earthquakes well beyond the design basis.)

Given this data, formal decision methodologies involving decision trees, value functions, and utility functions were applied to this basic data. It is shown that the formal decision methodology can be applied at several different levels. Examples are given (based on assumed NRC decision-maker preferences) in which the decision between several retrofits is changed from that resulting from a simple cost/benefit-ratio criterion by virtue of the decision-makinger's expressed (and assumed) preferences. The examples given in this report demonstrate that the application of formal decision-making methodology to the NRC decision-making process can potentially make such decisions more consistent, transparent and referenceable as well as allowing for explicit inclusion of the decision-maker's preferences. In addition, it allows for explicit inclusion of uncertainties associated with modelling issues (which are usually treated in terms of sensitivity studies in a probabilistic risk assessment) directly in the decision-making process. Recommendations for further work in this area are included.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF FIGURES (continued)

LIST OF TABLES

EXECUTIVE SUMMARY

As part of the NRC-sponsored program to study the implications of Generic Issue 57, "Effects of Fire Protection System Actuation on Safety-Related Equipment," a subtask was performed to evaluate the applicability of formal decision analysis methods to generic issues cost/benefit type decisions and to apply these methods to the GI-57 results. Following a literature review, it was determined that the concept of "utility" of a consequence as first developed by von Neumann and Morgenstern, and later extended to "multivariate utilities" in the work of Keeney and others was potentially applicable to the US NRC's decision-making role associated with determining backfits arising from generic issue studies.

To provide concrete examples of the application of the formal decision-making methods, the numerical results obtained from the analysis of three plants (two PWRs and one BWR) developed in the technical resolution program for GI-57 were studied. For each plant, these results included a calculation of the person-REM averted due to various accident scenarios and various proposed modifications to mitigate the accident scenarios identified. These results were recomputed to break out the benefit in terms of contributions due to random event scenarios, fire event scenarios, and seismic event scenarios. Further, the benefits associated with risk averted (in terms of person-REM) from earthquakes at three different seismic ground motion levels were separately considered. This was done to allow differentiation between potential retrofits which affect the risk arising primarily from earthquake levels below the design basis earthquake from those potential retrofits which affect risk resulting primarily from earthquakes well beyond the design basis.

In the initial sections of the report, the use of a cost benefit analysis approach to identifying viable (i.e., cost-effective) retrofits by comparing the cost benefit ratio (CBR) to a numerical criterion of \$1000/person-REM was reviewed. It was shown that, using the current definition of the cost benefit ratio, a different decision would be reached if one computes the CBR based on a mean value of the cost (in the numerator) divided by the mean value of the benefit (in the denominator) than would result from an exact calculation of the mean value of the cost benefit ratio itself. This follows from the fact that the uncertainty in the denominator (i.e., in the benefit in terms of risk averted) is orders of magnitude greater than the uncertainty in the numerator (i.e., in the cost). A simple fix is demonstrated to remove this inconsistency in the application of the cost benefit ratio.

Following this, a review of formal decision-making methodologies is presented, and four levels of decision making under uncertainty are considered:

1) Decisions based on a single attribute with no uncertainty

2) Decisions based on multiple attributes with no uncertainty

3) Decisions based on a single attribute with uncertainty

4) Decisions based on multiple attributes with uncertainty

Applications of viewing various potential actions as points in a multidimensional space of consequence attributes are shown, and the use of the classical concepts of "value functions" and "dominance" are illustrated. Finally, a brief review of the concept of utility of a consequence attribute is given, and the extension to a multivariate utility vector is indicated. A proposal to use the mean value of utility divided by cost (of each action) to make retrofit decisions (as an extension of cost benefit analyses) is made. Limitations (in terms of the arbitrary scaling usually permitted in assessing utilities) implied by the use of this normalized utility are identified.

Examples are presented to illustrate the application of these formal decision methodologies (involving decision trees, value functions, and utility functions) to the four different cases, using the plant evaluation data developed for the GI-57 technical resolution. In particular, examples are given (based on <u>assumed</u> NRC decision-maker preferences) in which the decision between several retrofits is changed from that resulting from a simple cost/benefit ratio criterion by virtue of the decision-makinger's (assumed) preferences.

The examples demonstrate that the application of formal decision-making methodology to the US NRC decision-making process can potentially make such decisions more consistent, transparent and referenceable as well as allowing for explicit inclusion of the decision-maker's preferences. In addition, it allows for explicit inclusion of uncertainties associated with modelling issues (which are usually treated in terms of sensitivity studies in a probabilistic risk assessment) directly in the decision-making process.

Implicit in the methods proposed herein is the requirement that key US NRC decision-makers can assess their individual preferences in terms of utility functions, and that consensus amongst them can be achieved so that an "approved" set of utility functions is available within the US NRC for use by individual decision-makers. This may be a difficult task. Alternative approaches (to the formal wager querying approach presented herein) to assessing individual utility functions have been suggested by reviewers of this report (especially the use of the Analytical Hierarchy Process) and further work in this area is indicated. In addition, further work evaluating the use of other consequence attributes (e.g., risk averted from different containment failure modes, etc.) and use of Bayesian updating in the decision analysis process (to assess the value of delaying decisions until new data becomes available) is identified.

1.0 INTRODUCTION AND OVERVIEW

1.1 Background

As part of the NRC-sponsored program to study the implications of Generic Issue 57, "Effects of Fire Protection System Actuation on Safety-Related Equipment," a subtask was added to consider the effects of uncertainty on the decision-making process for unresolved safety issues and to evaluate the potential applicability of formal decision-making procedures to safety issues retrofit decisions. The specific goals of this subtask were:

a. To evaluate the applicability of formal decision analysis methods to generic issues cost/benefit type decisions and to include the impacts of uncertainties in these methods and,

b. Apply these methods to GI-57 results.

In part, this subtask was in response to a letter (July 19, 1991) from the Advisory Committee on Reactor Safeguards (ACRS) to the USNRC Chairman. In this letter, the ACRS advocated

"...a deeper and more deliberate integration of (PRA) methodology into the NRC activities," and

"...a decision-making algorithm that prescribes a confidence level for (a) decision and uses both bottom line probability and uncertainty to achieve this "

In this report, both the issues of uncertainty in PRA results and the use of PRA results in decision making will be addressed.

One can envision a number of current issues and areas where more refined decision-making methods could provide substantial benefits to the USNRC in the process of assessing the viability of retrofits resulting from analyses of the unresolved safety issues. In the current decision-making process, a calculation of the cost benefit ratio (CBR) is made for each proposed modification (retrofit) and this CBR is compared with the criterion of $1K/person-REM averted. If the cost benefit ratio is less than this criterion, then the retrofit is considered viable.

Several issues can be raised with this procedure. First of all, if uncertainties are propagated through the calculation of the CBR, then it is to be expected that the <u>probability</u> of the CBR being less than the $1K/person-REM will be significantly different for those retrofits which primarily affect internal events scenarios as contrasted to those retrofits which primarily affect external events scenarios. (This follows from the fact that probability density functions associated with core damage frequencies due to external events scenarios are commonly found to have wider dispersion and to be more skewed than those

associated with internal events scenarios.) The CBR non-exceedance probabilities may even be significantly different for different retrofit options based on different accident sequences within the realm of internal events. Thus, in comparing point estimate values of the CBR, one is comparing different probabilities of exceedance of the cost benefit ratio between different retrofits. This is illustrated in Figure 1, where Figure 1(a) shows an approximately symmetric distribution on risk due to internal (for which the mean corresponds typically to a 50-60% probability of non-exceedance) while Figure 1(b) shows the very skewed distribution typical of seismic events (for which the mean often corresponds to a 80-90% probability of non-exceedance).

A second issue arises when one considers a number of different retrofits and the situation occurs in which the corresponding distributions of the cost benefit ratios of several viable fixes overlap significantly as shown in Figure 2. In this case, one may not be able to choose between the fixes due to the overlapping of the distributions. However, one fix may be more significant from the NRC viewpoint due to a number of other preferences or beliefs. Formal decision-making procedures would allow NRC to inject their preferences and differentiate between these viable retrofits.

More generally, it is possible that, in the NRC's view, certain retrofits are more or less desirable and that these preferences should be included in the decision-making process. For example, even though certain seismic and random retrofits might yield the same (point estimate) cost benefit ratio, an NRC decision-maker might feel that seismic events are "less likely" to be the cause of the next core damage event in the US commercial nuclear industry and feel, rather, that the next such event is more likely to be due to a random or human error. It is possible to use formal decision-making methods to build these preferences into the ranking of the various retrofits.

Another issue of likely interest to the NRC is that, given a number of retrofits for a specific plant, it may be possible to group these retrofits into sets (all of which could be accomplished at the next plant outage) in such a way as to obtain an optimal cost benefit achievement. This consideration was done on an ad hoc basis in the program for the resolution of USI A-45 (Reference 1). However, there currently exists no formal procedures to combine and rank groups of retrofits to achieve optimality.

An important consideration in retrofit recommendations resulting from studies on unresolved safety issues is the consideration of the need for obtaining additional data. In particular, the cost of additional data (whether obtained experimentally or through analysis) should (and can be) explicitly included in the cost benefit decision. If it turns out that the proposed additional experimentation or analysis (at whatever cost) will not impact the decision choice, then there is clearly no value in performing the experimentation or analysis. However, if it is

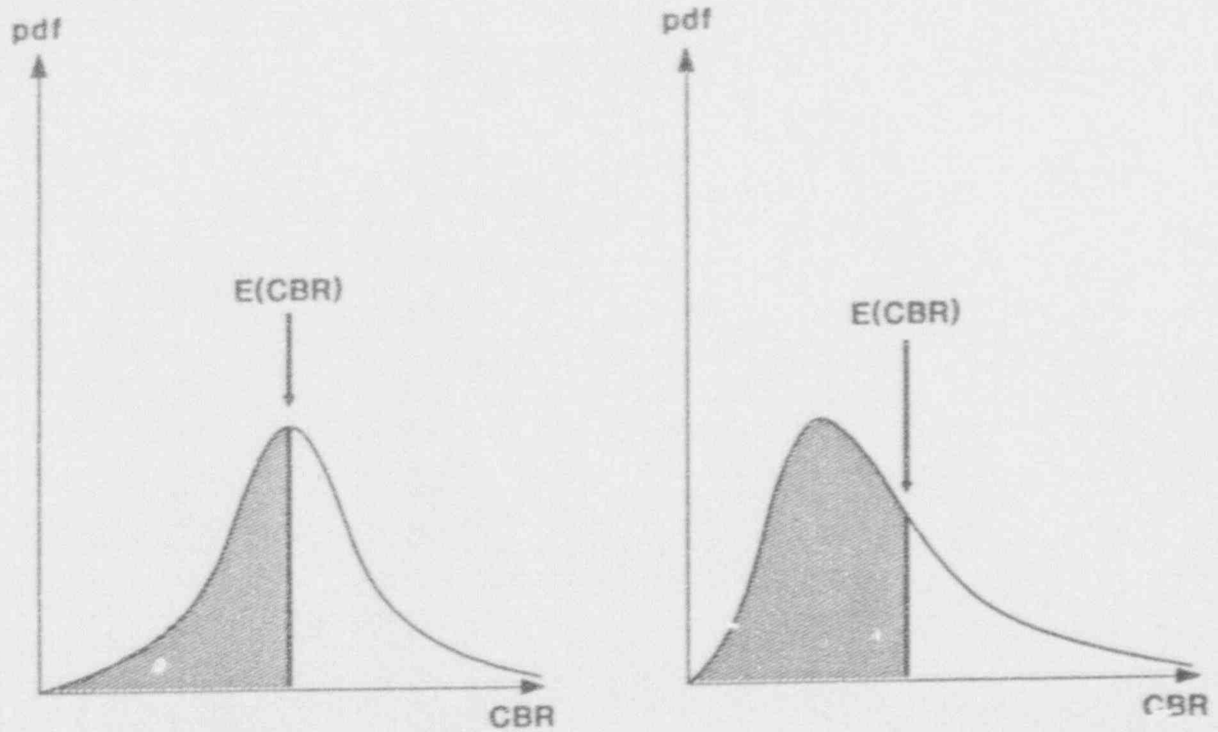

Figure 1 Schematic of the Mean of a Symmetric (a) and Skewed (b) Probability Density Function (pdf) of Cost Benefit Ratio (CBR)

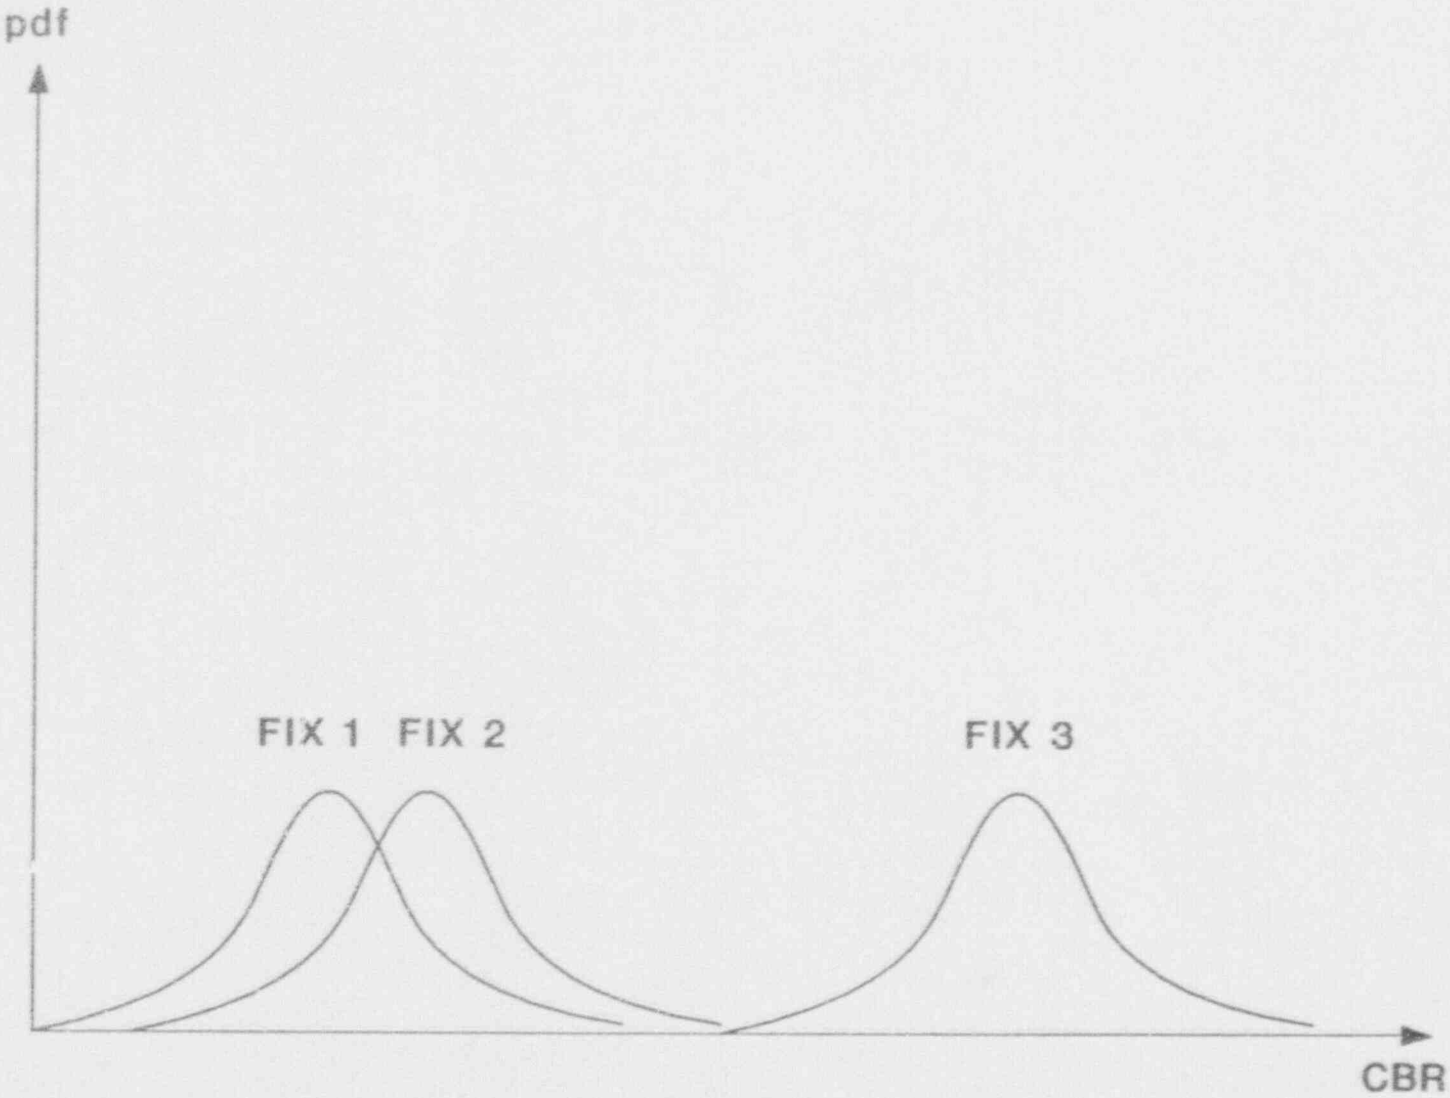

Figure 2 Schematic Illustrating Overlapping Cost Benefit (CBR) Probability Density Functions (pdf)

likely that the additional data could potentially affect the decision, then the NRC would have a clear basis for proceeding in that direction prior to making the decision. There are numerous examples of these considerations in the civil engineering realm where the design engineer must deal with uncertainties as to soil properties, etc. (See, for example, Reference 2).

Finally, the issue of timely implementation of retrofits is one of importance. Given expensive retrofits or retrofits requiring significant downtime, it is often the case that the power plant owner will propose to phase in the retrofits over the next several years. If this plan were to be accepted by NRC, it might be desirable to accumulate data during that time period and to determine the impact of the additional years of plant specific data on the decision. It may turn out, for example, that the additional years of data could change the cost benefit decision. As an example, in the area of risk due to fires, it is possible that the additional years of data could reflect enhanced housekeeping procedures implemented at the plant and that this demonstration of superior housekeeping could be used to reconsider the need for certain retrofits.

1.2 Results Available from GI-57 Studies

The NRC-sponsored program studying the implications of Generic Issue 57 provides an ideal opportunity for studying the potential applications of formal decision analysis methods to NRC decision making. In this program, three nuclear power plants were analyzed at a similar level of detail (including detailed plant walkdowns). The methods used to analyze these three plants were based on the plant logic models developed in the Unresolved Safety Issue TAP A-45 and NUREG-1150 programs performed at Sandia Labs (Refs. 1 and 3). Accident scenarios from a wide range of root causes leading to core damage due to the effects of actuations of fire protection systems were developed and analyzed. The root causes leading to the scenarios are shown in Table 1. Of these root causes, five are fire related, four are random related, three are seismic related, and one reflects a seismic fire interaction root cause. The accident sequences leading to core damage were then quantified and mapped onto containment failure models and release categories so that the end result is a calculation of person-REM exposure at the plant boundary calculated for a remaining plant life of twenty years.

In addition, physical plant modifications (retrofits) which would reduce or eliminate the accident sequence scenarios were identified. After a final plant visit, during which the feasibility of these modifications was evaluated, a detailed analysis of the costs of the various modifications was made, including on-site averted costs (OSACS).

Thus, for each plant considered in the GI-57 program, a complete analysis of accident scenarios and resulting risk as well as associated modifications and costs was performed. Furthermore, uncertainty analyses

Table 1

Potential Root Cause Scenarios Resulting
From Fire Protection System (FPS) Actuation

1. Fire in an adjacent zone causing FPS actuation
2. Fire-induced FPS actuation (due to fire in an adjacent zone) preventing recovery actions for random failures
3. Fire-induced FPS actuation (due to fire in an adjacent zone)
4. FPS actuation caused by human error
5. FPS actuation caused by steam pipe break
6. FPS actuation caused by hardware failures of FPS components
7. Seismic FPS actuations resulting from dust-triggered smoke detector activation
8. FPS actuations caused by seismic relay chatter
9. FPS actuations resulting from seismic-caused mechanical failures of FPS
10. Fires external to plant causing FPS actuation
11. Fire-induced FPS actuation due to a fire in the same zone
12. Seismic-fire interaction leading to FPS diversion
13. FPS actuation due to unknown causes

were performed on all these calculations so that probability distributions of all the output results are available. These results form the basic data which is available as input to a decision-making methodology having the goal of identifying which actions (retrofits or modifications) represent the optimal choice, given the NRC decision maker's goals, values and beliefs.

As mentioned, three plants were studied in detail in the GI-57 technical resolution program:

A General Electric BWR
A Westinghouse PWR
A Babcock and Wilcox PWR

The basic data developed for these three plants are shown on Tables 2, 3, and 4. For each plant, several proposed plant modifications are listed. Each proposed modification is associated with a total person-REM averted and the cost of the modification is also shown. (The values shown on these tables are mean values as computed from the actual distributions resulting from the uncertainty analyses.) The modifications identified for each of the three plants are shown on Tables 5, 6, and 7.

Note that in the original GI-57 technical program, only the benefit in terms of total person-REM averted was computed and reported for each modification. However, for this investigation of formal decision-making methodologies, the total benefit (in terms of person-REM averted) was further broken out into sources about which it is anticipated that the NRC decision makers would have strong beliefs, preferences and prejudices. To accomplish this, the uncertainty analyses were repeated, and the person-REM averted were recomputed and broken out into contributions due to internal event scenarios, fire related scenarios, and seismic scenarios. Furthermore, the seismic contribution was broken out into the contributions due to earthquakes occurring in three peak ground acceleration (pga) levels, corresponding (approximately) to

$$\text{OBE} < \text{pga} < 2\ \text{SSE} \quad (\text{SEIS } 1)$$
$$2\ \text{SSE} < \text{pga} < 6\ \text{SSE} \quad (\text{SEIS } 2)$$
$$\text{pga} > 6\ \text{SSE} \quad (\text{SEIS } 3)$$

where the OBE and SSE are the operating basis earthquake and the safe shutdown earthquake for the site, respectively.

The mean seismic hazard curves for these three plant sites are shown on Figure 3. Each curve gives the annual frequency of exceeding any particular value of peak ground acceleration at the site for which it was derived. From these curves, the annual frequency of exceeding the OBE or SSE can be determined.

While this breakdown choice was arbitrary, past experience with NRC personnel tends to indicate that these three earthquake levels are viewed

Table 2

GE Plant - Benefits (person-REM averted) and Costs

	MOD 1	MOD 2	MOD 3	MOD 4
INTERNAL	0	1.1	7.9	0
FIRE	0	7.5	10.5	0
SEIS 1	63.4	2.3	21.8	20.1
SEIS 2	5.1	0.1	34.8	35.7
SEIS 3	0	0	17.4	17.9
TOTAL	68.5 M-R	10.0 M-R	92.4 M-R	73.7 M-R
COST* (W/OSAC)	$66K	$69K	$172K	$101K
COST/BENEFIT ($K/Person-REM)	0.9	7.1	1.8	1.4

*Costs include on-site averted costs (OSAC) as given in Refs. 5, 6, and 7.

Table 3

B & W Plant - Benefits (person-REM averted) and Costs

	MOD 1 & 10	MOD 3	MOD 5 & 11	MOD 7
INTERNAL	0	0.9	0	0
FIRE	0	0.1	0	0
SEIS 1	3.1	0	6.8	6.4
SEIS 2	1.7	0	53.2	52.1
SEIS 3	0	0	17.0	16.9
TOTAL	4.8 M-R	1.0 M-R	77.0 M-R	75.5 M-R
COST*(W/OSAC)	$14K	$250K	$40K	$15K
$\frac{\text{COST}}{\text{BENEFIT}}\left[\frac{\$K}{\text{Person-REM}}\right]$	2.9	250.0	0.5	0.2

*Costs include on-site averted costs (OSAC) as given in Refs. 5, 6, and 7.

Table 4

W Plant - Benefits (person-REM averted) and Costs

	MOD 1 & 10	MOD 3	MOD 11
INTERNAL	0	0	0
FIRE	0	0	0
SEIS 1	0	0	0
SEIS 2	0.3	0.1	1.0
SEIS 3	0.1	0.2	1.8
TOTAL	0.4 M-R	0.3 M-R	2.8 M-R
COST*(W/OSAC)	$14K	$200K	$86K
$\frac{\text{COST}}{\text{BENEFIT}} \left(\frac{\$K}{\text{Person-REM}}\right)$	35.0	666.6	30.7

*Costs include on-site averted costs (OSAC) as given in Refs. 5, 6, and 7.

Table 5

General Electric BWR - Plant Modifications Analyzed

- Modification 1: Upgrade the FPS Controller with Seismically Qualified Printed Circuit Boards (Eliminates Root Cause 8)
- Modification 2: Replace Smoke Detector Actuated FPS with a Heat Detector Actuated FPS (Eliminates Root Cause 7)
- Modification 3: Reroute Safety-Related Cables (Reduce all root cause core damage frequency contribution by 0.1)
- Modification 4: Seismically Qualify CO2 Tank, Outlet Piping and Battery Rack (Eliminates Root Cause 12)

Table 6

Babcock and Wilcox PWR - Plant Modifications Analyzed

- Modification 1: Upgrade the FPS Controller with Seismically Qualified Printed Circuit Boards (Eliminates Root Cause 8)
- Modification 3: Reroute Safety-Related Cables (Reduce all root cause core damage frequency contribution by 0.1)
- Modification 5: Seismically Qualify a Battery Rack (Eliminates Root Cause 12)
- Modification 7: Provide Fire Wraps for Safety-Related Cable (Eliminates Root Cause 12)
- Modification 10: Replace Low Fragility Control Relays with Hardened Relays (Eliminates Root Cause 8)
- Modification 11: Seismically Anchor Electrical Cabinets (Eliminates Root Cause 12)

Table 7

Westinghouse PWR - Plant Modifications Analyzed

- Modification 1: Upgrade the FPS Controller with Seismically Qualified Printed Circuit Boards (Eliminates Root Cause 8)
- Modification 3: Reroute Safety-Related Cables (Reduce all root cause core damage frequency contribution by 0.1)
- Modification 10: Replace Low Fragility Control Relays with Hardened Relays (Eliminates Root Cause 8)
- Modification 11: Seismically Anchor Electrical Cabinets (Eliminates Root Cause 12)

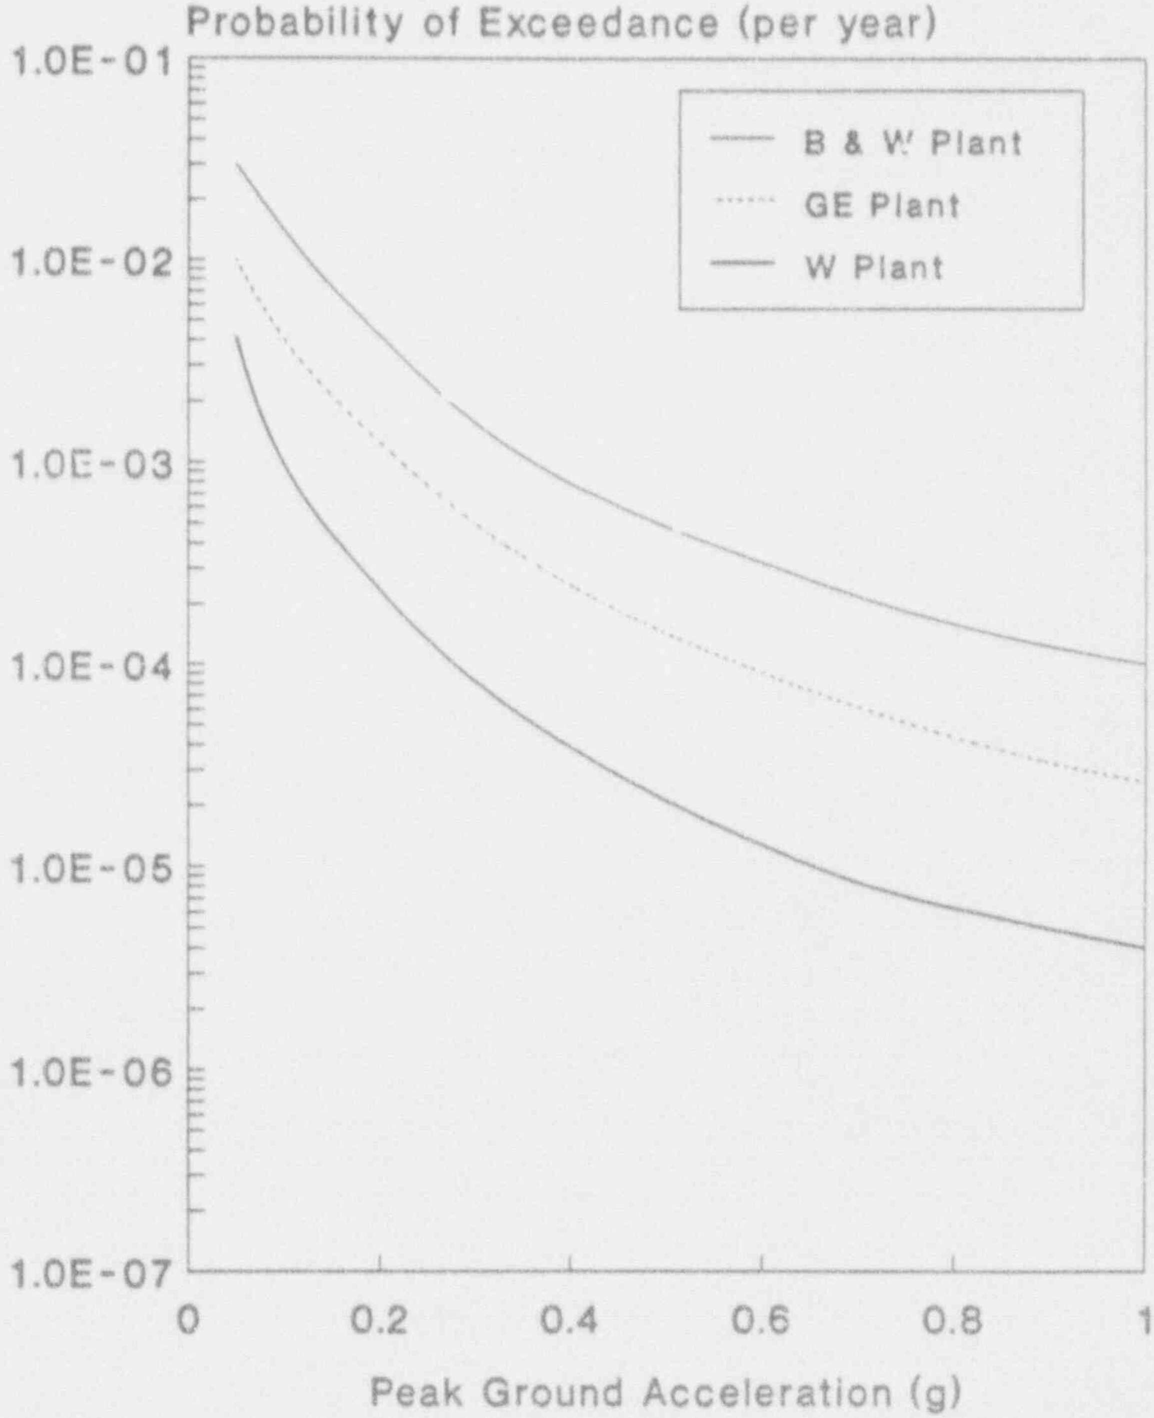

Figure 3 Mean Hazard Curves for the GI-57 Plants

substantially differently as to their likelihood of causing a future core damage accident. (These breakdowns will be used later in the report in examples of potential decision-making approaches.) Of course, the total sum of the mean person-REM averted given for each modification must sum to the total values computed and reported in the GI-57 individual plant reports (References 4, 5, and 6).

Finally, on the bottom of each of Tables 2, 3, and 4, the cost benefit ratio for each modification (as computed from the mean of the person-REM averted and the mean of the cost) is shown. According to the current NRC guidelines, a cost benefit ratio less than $1K/person-REM indicates a potentially viable retrofit. The data in Tables 2, 3, and 4 constitute the basic data which will be used in this report to study potential applications of formal decision-making methodologies to the NRC decision-making process, with emphasis on their application to retrofit decisions resulting from evaluation of Unresolved Safety Issues.

1.3 The Cost Benefit Ratio and Its Inverse

In the current value impact methodology, once the annual total person-REM of released dose has been calculated, then a cost benefit ratio is computed based on the cost of each proposed retrofit divided by the person-REMs which are averted by the retrofit. Usually, this calculation is performed both with and without onsite averted costs (OSAC). This cost benefit ratio is then used as a basis for identifying viable fixes. In the current procedure, according to the Handbook of Value Impact Assessment, if the cost benefit ratio (CBR) is less than $1.0K/person-REM averted, then the proposed modification is felt to be viable. If the CBR is greater than 1.0, then the retrofit is not taken to be viable.

Implicit in this definition is the assumption that the cost benefit ratio so computed is either a best estimate value or a mean point estimate. (The mean point estimate is the mean of the cost divided by the mean of the benefit.) The reason for this statement will be shown below. The implication, however, is that the cost benefit ratio as currently defined is <u>not</u> the appropriate decision variable to be used for cost benefit decisions when a full uncertainty analysis is performed in calculating the benefit. Rather, it will be found that the <u>inverse</u> of the cost benefit ratio is the quantity which should be utilized if a full uncertainty analysis is to be performed as part of the cost benefit decision.

To see this, consider the general form of the cost benefit ratio as used in a point estimate context today. The numerator of this ratio is the cost of the retrofit and the uncertainty in that cost. This uncertainty is usually relatively small and, for GI-57, was taken as plus or minus 12% of the best estimate value. The denominator, however, incorporates significantly more uncertainty and represents the contribution of each

accident sequence to the total increment in person-REMs averted. The general form of the denominator is

$$\Sigma \ CDF_i * CSF_i * ST_i, \qquad i=1,n$$

In this expression the CDF_i are the annual frequencies of core damage associated with the different containment event tree branches. The terms CSF_i are the split fractions which allocate the contributions of the branches to the n different release categories. Finally, the ST_i are the source terms associated with each release category.

The split fractions and the source terms are usually assumed have relatively little uncertainty in them, and are often represented as either point values or by simple triangular distributions. By contrast, the core damage frequency terms (CDF_i) typically encompass a large range of uncertainty, and have cumulative distribution functions which are very nearly lognormal in form. Thus, it can be seen that the behavior of the cost benefit ratio is approximately the behavior of the inverse of a lognormal random variable.

The inverse of a lognormal random variable is a function for which the mean of the function is nowhere near the function of the mean. In fact, using the well-known properties of the lognormal distribution, it is easy to show that the true mean of the inverse of a lognormal random variable is, in fact, equal to the mean point estimate multiplied by $\exp(\beta^2)$ as shown in Table 8. Thus, for example, if we were dealing with a cost-benefit ratio that consisted only of a constant cost divided by a lognormal random variable for which the log standard deviation β was equal to 2, then the true mean of the distribution of the cost benefit ratio would, in fact, be equal to approximately 50 times the point estimate mean. (By contrast, of course, the true mean of the ratio of a lognormal random variable benefit divided by a constant cost is exactly equal to the mean of the benefit distribution divided by the constant cost - and is thus exactly equal to the mean point estimate.) The implications of this are that if one chooses to propagate uncertainties through the entire cost benefit calculations and then compare some ratio with the value of $1000 per person-REM, it is <u>not</u> appropriate to use the form of the cost-benefit ratio as currently defined. Rather, it is appropriate to use the inverse of the cost-benefit ratio.

To illustrate this with actual results obtained in the GI-57 program, consider the four modifications obtained in the GE Plant analysis (Reference 4). The first column of Table 9 presents the mean point estimate (MPE) values of the cost benefit ratio as presented in the earlier GI-57 reports, as well as its inverse. The second column of Table 9 presents the results of a full uncertainty analysis in which all uncertainties are propagated through the analysis to obtain a distribution on the cost-benefit ratio (as currently defined) for each

Table 8

Mean Point Estimate and True Mean of the Function 1/X, Where X is a Lognormal Random Variable

---

Consider the function $Y = 1/X$, where X is a lognormal random variable characterized by its median M_x and its log standard deviation β_x (i.e., β_x is the standard deviation of the logarithm of X). The probability density function f(X) of X is given by

$$f(X) = \frac{1}{\beta_x X \sqrt{2\pi}} \exp\left[\frac{-1}{2\beta_x^2} (\ln X - \ln M_x)^2\right]$$

By the properties of a lognormal random variable, the median and the log standard deviation of the random variable $Y = 1/X$ are given by

$$M_y = 1/M_x \quad \text{and} \quad \beta_y = \beta_x$$

and the true mean of Y is given by the expression

$$\mu_y = M_y \exp(\beta_y^2/2)$$

$$= \frac{1}{M_x} \exp(\beta_x^2/2) \qquad (1)$$

But, for the lognormal random variable X, the true mean and the mean point estimate are both given by

$$\mu_x = M_x \exp(\beta_x^2/2)$$

and thus, the mean point estimate of Y is

$$\frac{1}{\mu_x} = \frac{1}{M_x \exp(\beta_x^2/2)} \qquad (2)$$

Comparing equations (1) and (2) shows that

$$\mu_y = \frac{1}{\mu_x} \exp(\beta_x^2)$$

and thus the true mean of $Y = 1/X$ is a factor of $\exp(\beta^2)$ times the point estimate mean of Y (which is $1/\mu_x$) when the random variable X is lognormal.

Table 9

Comparison of Mean Point Estimates (MPE) and True Means
For the GI-57 General Electric Plant Modifications
Cost Benefit Ratios (CBRs)

Modification 1		CBR*		1/CBR	
	MEAN =	3.60E+01		MEAN =	9.42E-01
MPE(CBR) = 1.00	95% =	1.26E+02		95% =	3.58E+00
	85%	4.24E+01		85% =	1.18E+00
MPE(1/CBR) = 0.922	50% =	5.96E+00		50% =	1.67E-01
	15% =	7.82E-01		15% =	2.31E-02
	5% =	2.60E-01		5% =	7.06E-03
	M_{CBR} =	6.02E+00		$M_{1/CBR}$	1.66E-01
	b_{CBR} =	1.91E+00		$b_{1/CBR}$	1.91E+00
Modification 2		CBR		1/CBR	
	Mean =	2.92E+01		MEAN =	1.46E-01
MPE(CBR) = 8.05	95% =	9.56E+01		95% =	4.27E-01
	85% =	4.67E+01		85% =	2.24E-01
MPE(1/CBR) = 0.124	50% =	1.51E+01		50% =	6.50E-02
	15% =	4.24E+00		15% =	2.12E-02
	5% =	2.26E+00		5% +	9.75E-03
	M_{CBR} =	1.50E+01		$M_{1/CBR}$ =	6.65E-02
	b_{CBR} =	1.17E+00		$b_{1/CBR}$ =	1.17E+00
Modification 3		CBR		1/CBR	
	MEAN =	1.42E+01		MEAN =	3.88E-01
MPE(CBR) = 2.54	95% =	5.13E+01		95% =	1.56E+00
	85% =	2.28E+01		85% =	6.19E-01
MPE(1/CBR) = 0.394	50% =	5.76E+00		50% =	1.72E-01
	15% =	1.57E+00		15% =	4.30E-02
	5% =	5.96E+00		5% =	1.90E-02
	M_{CBR} =	5.61E+01		$M_{1/CBR}$ =	1.68E-01
	b_{CBR} =	1.32E+00		$b_{1/CBR}$ =	1.32E+00
Modification 4		CBR		1/CBR	
	MEAN =	4.46E+02		MEAN =	7.94E-01
MPE(CBR) = 1.93	95% =	1.60E+03		95% =	1.77E+00
	85% =	4.19E+02		85% =	3.91E-0[illegible]
MPE(1/CBR) = 0.517	50% =	3.06E+01		50% +	3.14E-02
	15% =	2.46E+00		15% =	2.27E-03
	5% =	5.50E+0[illegible]		5% =	5.43E-04
	M_{CBR} =	3.1[illegible]		$M_{1/CBR}$ =	3.19E-02
	b_{CBR} =	2.[illegible]		$b_{1/CBR}$ =	2.44E+00

*Units of CBR in this table are $K/person-REM.

of the modifications. The third column presents the results of a full uncertainty analysis as applied to the inverse of the (same) cost benefit ratio. For each modification, columns 2 and 3 present the exact mean of the distributions of CBR and 1/CBR, several percentiles of the distributions, and the best fitting lognormal distribution parameters: the medians (M_{CBR}, $M_{1/CBR}$) and the log standard deviations (b_{CBR}, $b_{1/CBR}$). The first observation from this table is that the true mean values of the cost benefit ratio are from 10 to 200 times larger than the mean point estimate values. This is to be expected given on the arguments based on the properties of a log normal random variable presented above. Further, using the best fit lognormal distribution median (M_{CBR}) and log standard deviation (b_{CBR}) listed on Table 9, it is easily shown that the true mean of the cost benefit ratio distribution is indeed approximately equal to $\exp(b^2)$ times the mean point estimate value. This implies that the denominator in the cost benefit ratio, even though not exactly lognormal, is behaving very nearly like a lognormal function.

Examining the third column shows that the true mean of the distribution of the inverse of the cost benefit ratio is very nearly equal to the mean point estimate of the inverse of the cost benefit ratio which, again is to be expected based on the arguments presented above. All of this indicates that, if one is interested in understanding the implications of using a point estimate calculation of the cost benefit ratio and comparing it with the criteria of a $1,000 per person-REM, then it is only appropriate to perform the uncertainty analysis on the inverse of the cost benefit ratio rather than the cost benefit ratio itself. Thus, in all the studies and results to be presented later in this report addressing various aspects of decision making under uncertainty, it is the inverse of the cost benefit ratio and generalizations thereof that will be analyzed.

This inverse cost benefit ratio will be denoted as the BPB ratio (standing for Bang per Buck). A BPB ratio greater than 1.0 person-REM/$K denotes a viable retrofit (corresponding to a CBR less than 1.0 $K/person-REM). The BPB ratios for the GE and B & W plants are shown on Tables 10 and 11. (The Westinghouse plant will not be considered since its cost/benefit ratios indicate that none of the modifications are even remotely cost-effective.)

1.4 Formal Decision-Making Methodologies

As discussed above, the use of formal decision-making methods in ranking retrofits resulting from studies on unresolved safety issues has great potential. In general, a formal decision-making methodology includes the following steps.

1. Identification of actions (that is, backfits or retrofits)

2. Identification of one or more consequence attributes which characterize the results of each of these actions (e.g., exposure to the public, cost, etc.)

Table 10

GE Plant - BPB Ratios (person-REM/$K)

	MOD 1	MOD 2	MOD 3	MOD 4
INTERNAL	0	0	0.05	0
FIRE	0	0.11	0.06	0
SEIS 1	0.96	0.03	0.13	0.20
SEIS 2	0.08	0	0.20	0.35
SEIS 3	0	0	0.10	0.18
TOTAL	1.04	0.14	0.54	0.73

Table 11

B & W Plant - BPB Ratios (person-REM/$K)

	MOD 10	MOD 5	MOD 7
INTERNAL	0	0	0
FIRE	0	0	0
SEIS 1	0.24	0.17	0.43
SEIS 2	0.13	1.33	3.47
SEIS 3	0	0.43	1.13
TOTAL	0.37	1.93	5.03

3. Input of the decision maker's preferences as to the relative desirability of the consequence attributes. These are expressed mathematically in terms of value functions or utility functions. These latter functions are also termed the "outrage factor" and reflect the decision makers aversion to certain outcome attributes.

4. Calculation of the consequences of each of these actions incorporating all significant uncertainties in the process.

5. Ranking the actions relative to the utilities of the consequence attributes (which express the decision makers preferences) and then making the decision.

All formal decision-making methodologies would incorporate the above steps.

An additional step may be included if additional data is sought or additional exrerIments performed. In this case, one can use a Bayesian updating procedure to alter certain of the assumed probabilities in the calculation of consequences and thus determine the impact on the final decision.

Formal decision-making methods can be applied in the four distinct situations shown on Figure 4 (as taken from Reference 10). The simplest case is that in which we associate a single attribute to the consequence of each action and there is no uncertainty in calculating this attribute. An example of this is the use of a best estimate cost benefit ratio. In this case, the decision is made simply by comparing the numerical ranking of the attributes.

The next more complicated situation involves a single attribute associated with each action, but in this case we recognize uncertainties associated with the calculation of this attribute. This, for example, would be the case when we base the decision on the true mean value of the cost benefit ratio or some percentile in the distribution of the cost benefit ratio and then compare that percentile with the numerical criterion.

The third level would involve situations in which each consequence would be characterized by a number of different attributes (i.e., a vector of attributes) but for which there was no uncertainty in the calculation of these attributes. An example of this might be a case in which one would consider the increments in person-REMs associated with different release categories which result from each particular retrofit. Another example would be a case in which, for seismic-related retrofits, the consequences could be expressed in terms of the person-REMs associated with earthquakes of different levels (say less than the OBE, between the OBE and the SSE, and for earthquakes greater than the SSE). In this case,

SINGLE ATTRIBUTE NO UNCERTAINTY	MULTIPLE ATTRIBUTES NO UNCERTAINTY
SINGLE ATTRIBUTE WITH UNCERTAINTY	MULTIPLE ATTRIBUTES WITH UNCERTAINTY

Figure 4. Formal Decision Methods Can Be Applied At Four Levels

the decision maker may want to make his decision based on the consequences of earthquakes in the lower categories and place little significance on the predictions of earthquakes greater than the SSE.

Finally, the most general case involves the situation in which for each action, there is are multiple attributes (an attribute vector) and, further, there is uncertainty involved in the calculation of these multiple attributes. This represents the most general case and is undoubtedly the case that has most applicability to the USNRC's retrofit decisions. Solution of decision problems for this last case involves concepts and procedures developed for each of the three simpler cases. Thus, in the following, each of these four levels of decision making will be addressed explicitly, with examples and illustrations of their potential usefulness.

2.0 DECISIONS WITHOUT UNCERTAINTY

Before going on to consider the general problem of decisions involving uncertain outcomes we will first discuss decisions between several actions having either a single attribute or multiple attributes associated with each action but for which there is no uncertainty in the attributes resulting from each action. This is done, first, to illustrate the decision tree symbolism which will be used in later discussions and, second and more importantly, to show that in certain circumstances it is possible to make meaningful decisions based on topological concepts <u>without</u> direct consideration of the uncertainties involved.

2.1 Single Attribute Consequences

This is the simplest decision problem as illustrated in Figure 5. As shown we have a decision between actions a' and a", each of which with certainty lead to a consequence C' and C" respectively. That the outcomes are certain is indicated on this figure by the fact that the probabilities of the outcomes given the actions (shown as p' and p") are both equal to unity. In this case, of course, the decision analysis is straightforward and is made based on comparison of the numerical values of the scalar consequences. An example of this simple case is the use of the mean point estimate of the cost/benefit ratio as the criterion for choosing between two retrofit modifications.

2.2 Multiple Attribute Consequences

The general decision model here is also a branching tree of different potential actions, and associated with each action is a consequence vector which is known with no uncertainty, as shown in Figure 6. Thus, in this case, we must make a decision by comparing the different consequence <u>vectors</u> corresponding to the different actions.

In this case, it is useful to think of each consequence vector as defining a point in multidimensional space as illustrated on Figure 7. Thus to each action A_i there corresponds a point in the multidimensional space of consequences. Each axis coincides with one of the components of the consequence vector.

2.3 Dominance

The important concept in this case is that of dominance. To illustrate this concept, let us first stipulate that increasing values of each component of the consequence vector are increasingly preferred. Then, the point $X(A_1)$ <u>dominates</u> the point $X(A_2)$ when

$$X_i(A_1) \geq X_i(A_2) \text{ for all } i$$

and

$$X_i(A_1) > X_i(A_2) \text{ for some } i.$$

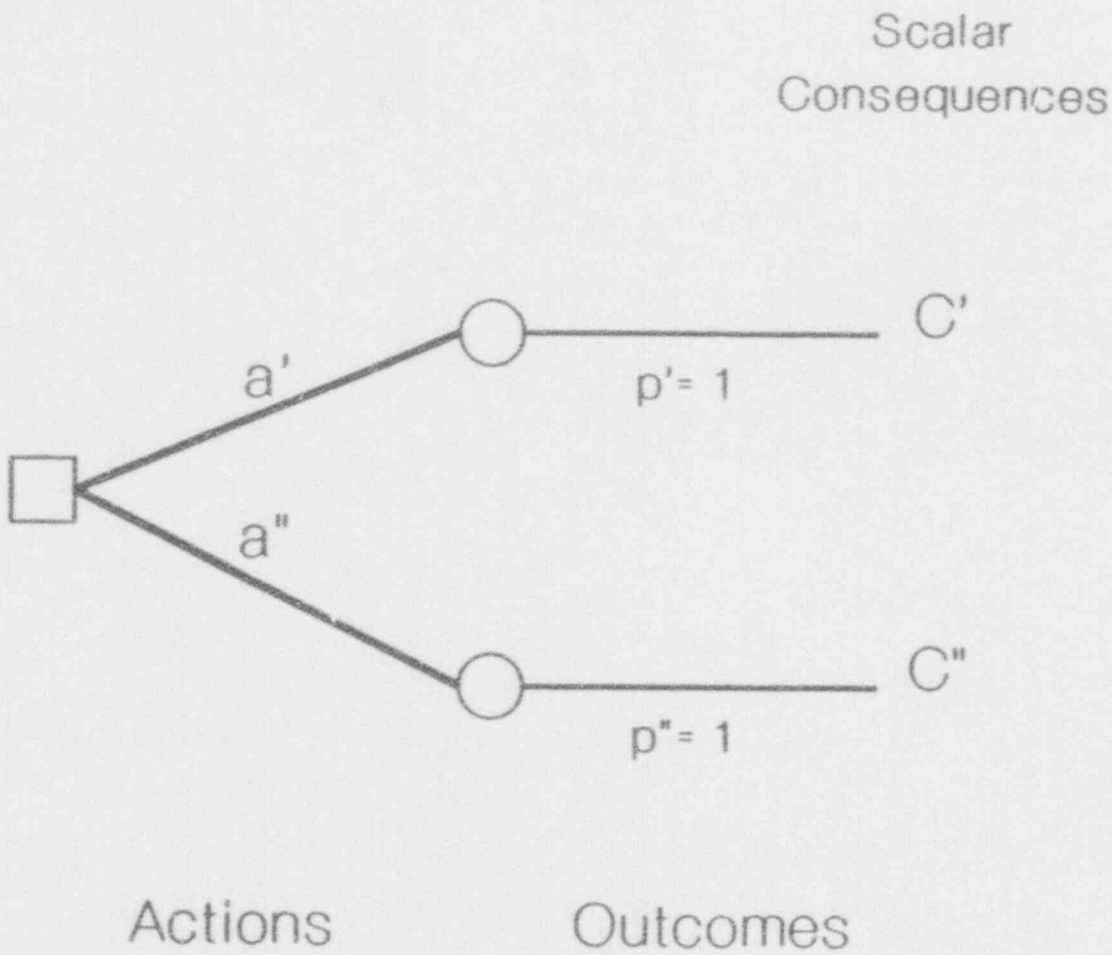

Figure 5 Decision Tree for Single Attribute Consequence without Uncertainty

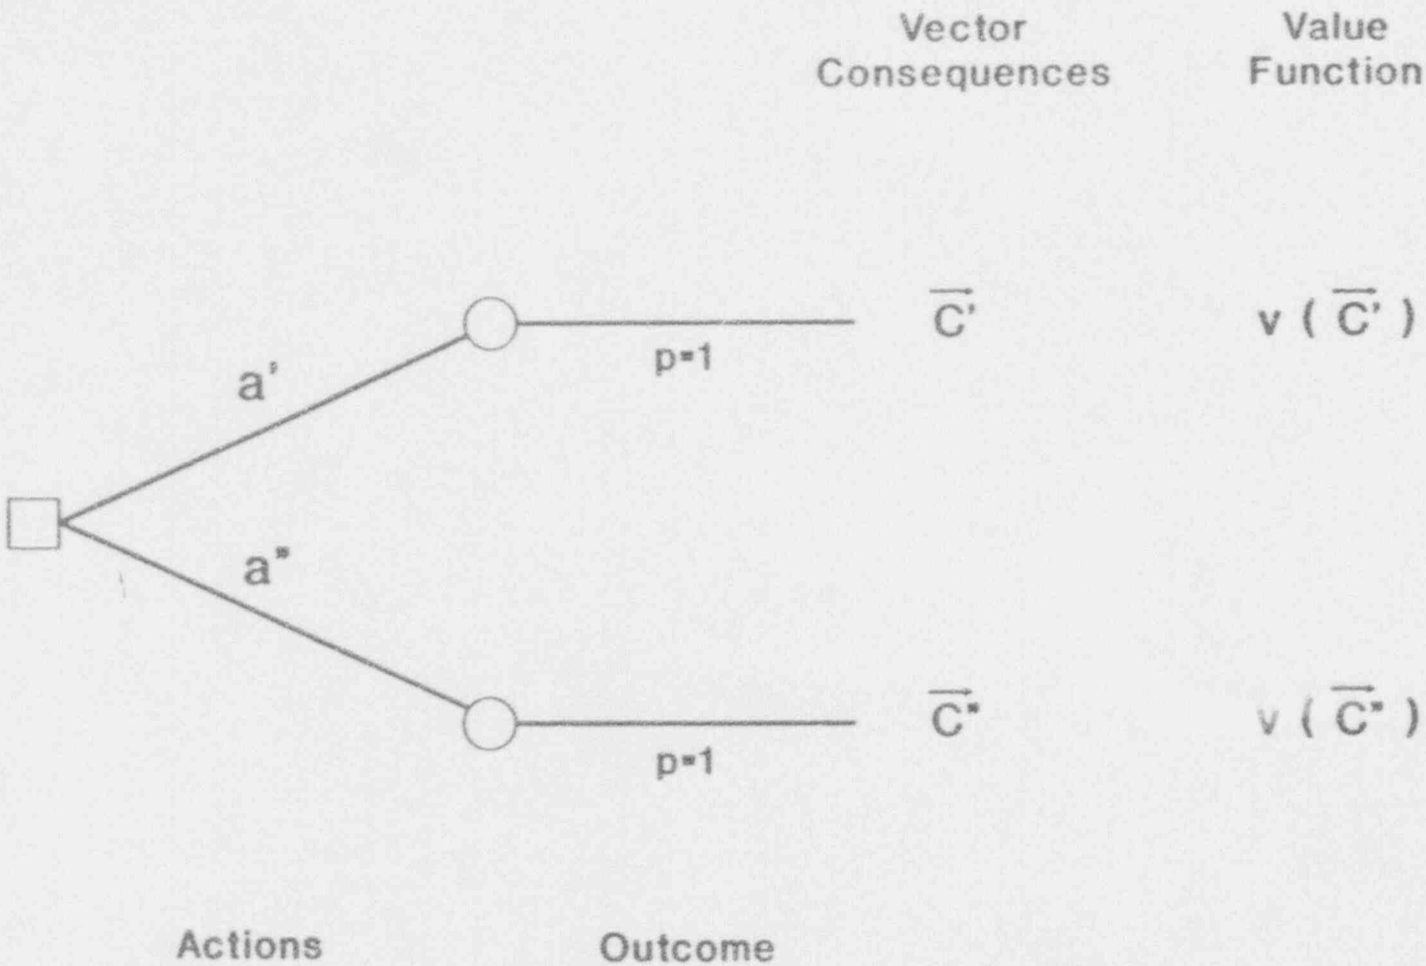

Figure 6 Decision Tree for Multiple Attribute Consequence without Uncertainty

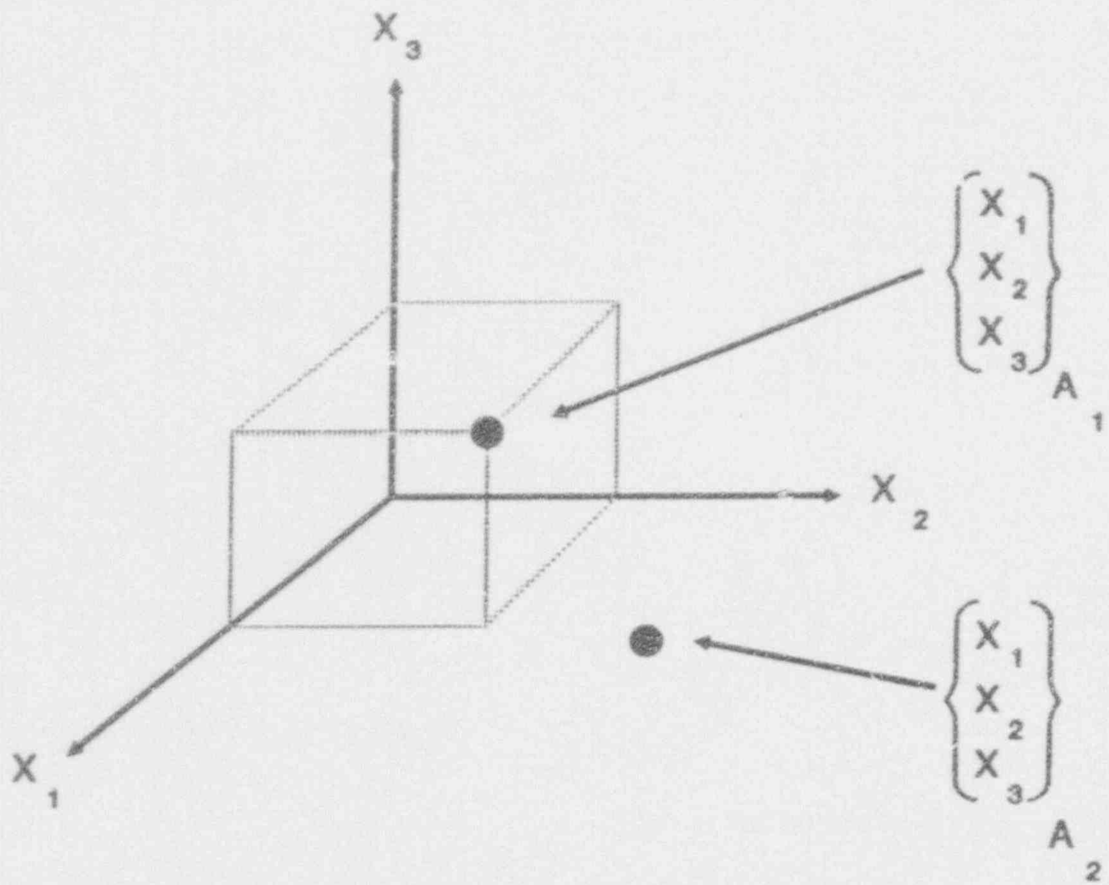

Figure 7 Consequence Vectors as Points in N-Space

That is, all the consequence vector attributes arising from action A_1 are equal to or greater than the corresponding attributes of action A_2, and at least one attribute of A_1 is greater than the corresponding attribute of action A_2. This is illustrated in two dimensions in Figure 8 where the point X_1 is above and to the right of point X_2, which shows point X_1 indeed dominates point X_2. Thus, assuming that increasing values of each attribute are to be preferred, one sees that all points "below and to the left of a given point" are dominated by that point. Hence in Figure 8, point X_2 is dominated by all points in the non-cross-hatched region.

It is important to note that, in looking for dominance, we never compare the different attributes so that the scales of the different axes may be entirely different. Furthermore, the scales may be quantitative scales or subjective scales (often denoted as derived scales) for which different levels of an attribute are qualitatively ranked. This is so since, in looking for dominance, we compare consequences only on a component by component basis. Finally, when one act has been shown to dominate all other acts, then it will always be preferred no matter what reasonable preferences are expressed concerning the utility of these various attributes in the consequence vector.

2.3.1 Example of Dominance

As an example of dominance, consider the data for the B & W plant expressed in terms of their BPB ratios. This data is repeated in Table 12 where three modifications (Mod 10, Mod 5, and Mod 7) are shown along with their corresponding consequences in terms of person-REM/$K due to the five sources of internal events, fires, and three seismic levels. As can be seen, Mod 7 dominates both Mod 5 and Mod 10 on a component by component basis. Thus Mod 7 will always be preferred no matter what reasonable preferences are expressed concerning the utility of these five attributes.

2.4 Value Functions

Continuing with the discussion of the concept of the consequence vector of each action being a point in the multidimensional space, consider the schematic shown in Figure 9. (This, of course, is a two dimensional example.) The solid dots indicate points which are dominant points whereas the open circles indicate points which are dominated. The set of consequence points not dominated is called the "efficient frontier" of the set of all actions being considered. This is also known as the "Pareto Optimal Set". Thus the points on the efficient frontier represent points which are not dominated and which are, in each case, in a certain sense optimal.

In general, the efficient frontier may be either discrete (as shown) or may be continuous. Furthermore, it is not necessarily convex. However, if it is convex, then cardinal as well as ordinal notions will apply.

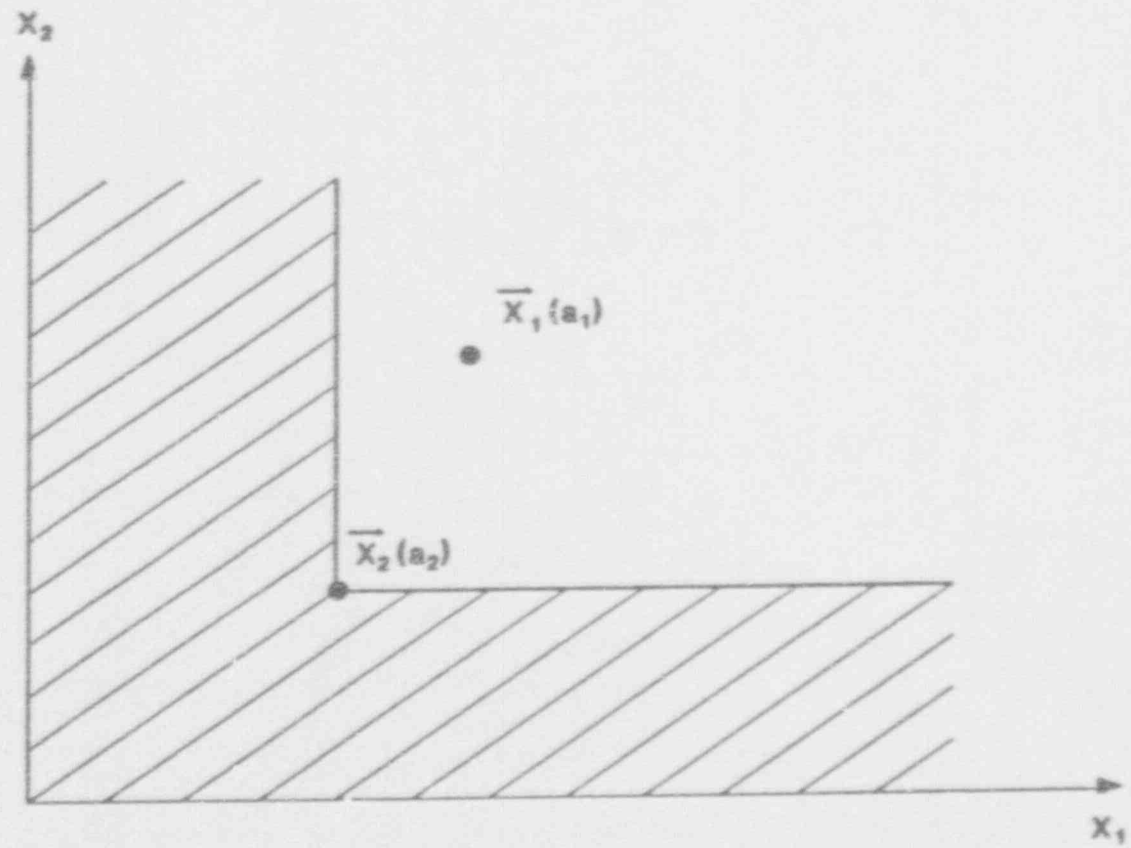

Figure 8 Illustration of Dominance in 2 space

Table 12

Example Of Dominance - B & W Plant BPB Ratios (person-REM/$K)

	MOD 10	MOD 5	MOD 7
INTERNAL	0	0	0
FIRE	0	0	0
SEIS 1	0.24	0.17	0.43
SEIS 2	0.13	1.33	3.47
SEIS 3	0	0.43	1.13

Figure 9 Illustration of the "Efficient Frontier"

The situation is thus that the decision maker wants to choose one act out of the set of acts, each of whose consequence vector defines a point which lies on or near the efficient frontier so that his particular preferences as to the consequences are satisfied.

In general, it should be noted that for a discrete number of acts with a discrete number of consequences associated with each act, the optimal set may be determined directly by computer search for the dominated points. However it may turn out that there are no dominated points. In this case, the decision maker must consider the entire set of acts in further expressing his preferences. Finally, for actions representing a continuous range of choices and hence a continuous set of consequence points, the efficient frontier may be found exactly by using linear programming techniques as given in Reference 7. These techniques allow the decision maker to explore the efficient frontier and choose one or more actions that result in consequences which he prefers. Decisions based on this process, however, are highly judgmental and not amenable to documentation and referenceability and hence are not commonly used. Alternate methods for expressing preferences in a quantitative fashion are described below.

Let us assume that the decision maker can define a set of curves in 2 space or surfaces in n space such that he is <u>indifferent</u> to actions which result in points which lie on the same curve (surface) as shown on Figure 10. Thus each of these curves is a so-called "indifference curve" or surface. Furthermore, we assume that an index can be associated with each curve such that an increasing index corresponds to the decision maker's increasing preference as to the action outcomes. Once such a family of curves or surfaces is defined to specify the decision maker's preferences, then the solution to the decision problem is to find that point on the efficient frontier which coincides with or is tangent to the indifference curve that is most preferred by the decision maker. This is illustrated in Figure 11. In this figure, consequence X_0 is the best consequence of all those consequences in the complete set of acts (defined as the region R) given the preference structure illustrated.

Analytically, the preference curves or surfaces are defined by algebraic expressions denoted as value functions. By definition, a function $v(x)$ which associates a real number v to each point x in the consequence space is said to be a value function provided

a) $v(x') = v(x'')$ implies that the consequence x' is equally preferred to consequence x", and

b) $v(x') > v(x'')$ implies that the consequence x' is more preferred than the consequence x".

Thus if a value function $v(x)$ is known, then the decision problem reduces to that of a maximization problem. That is, to find that act a

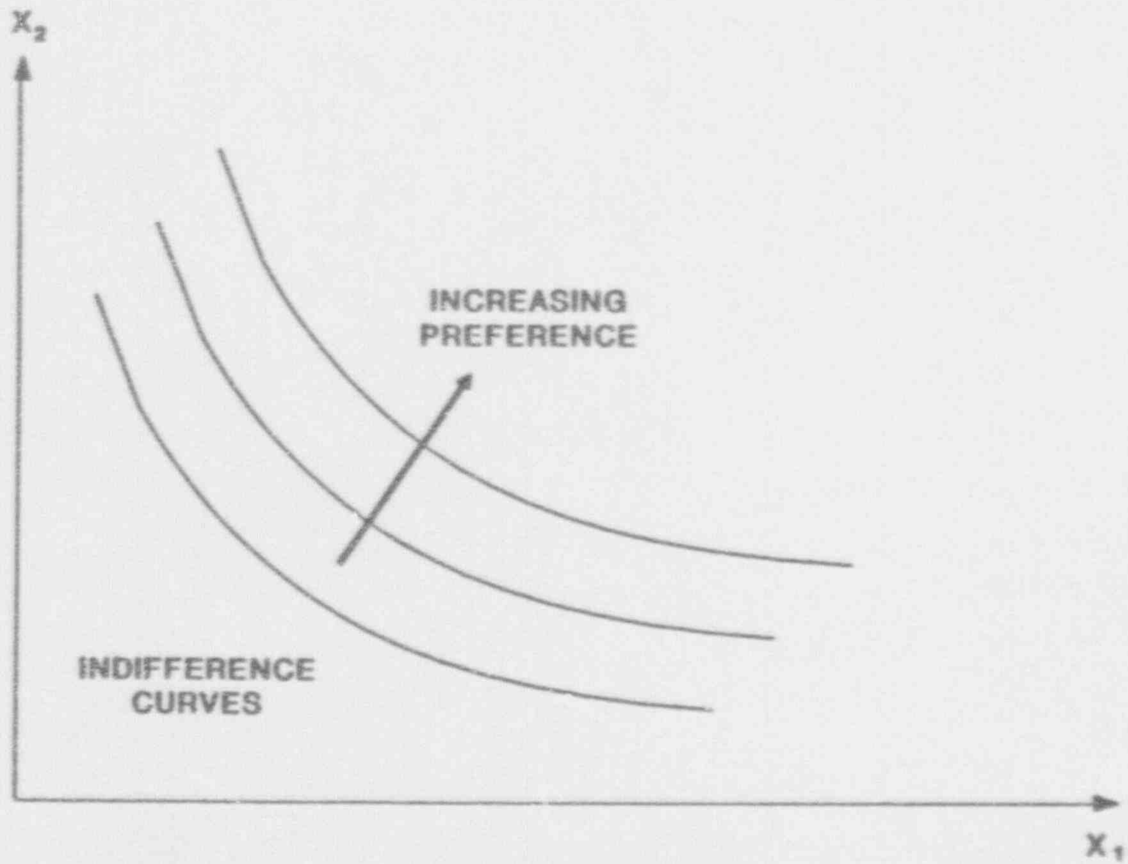

Figure 10 Example of Preference Curves

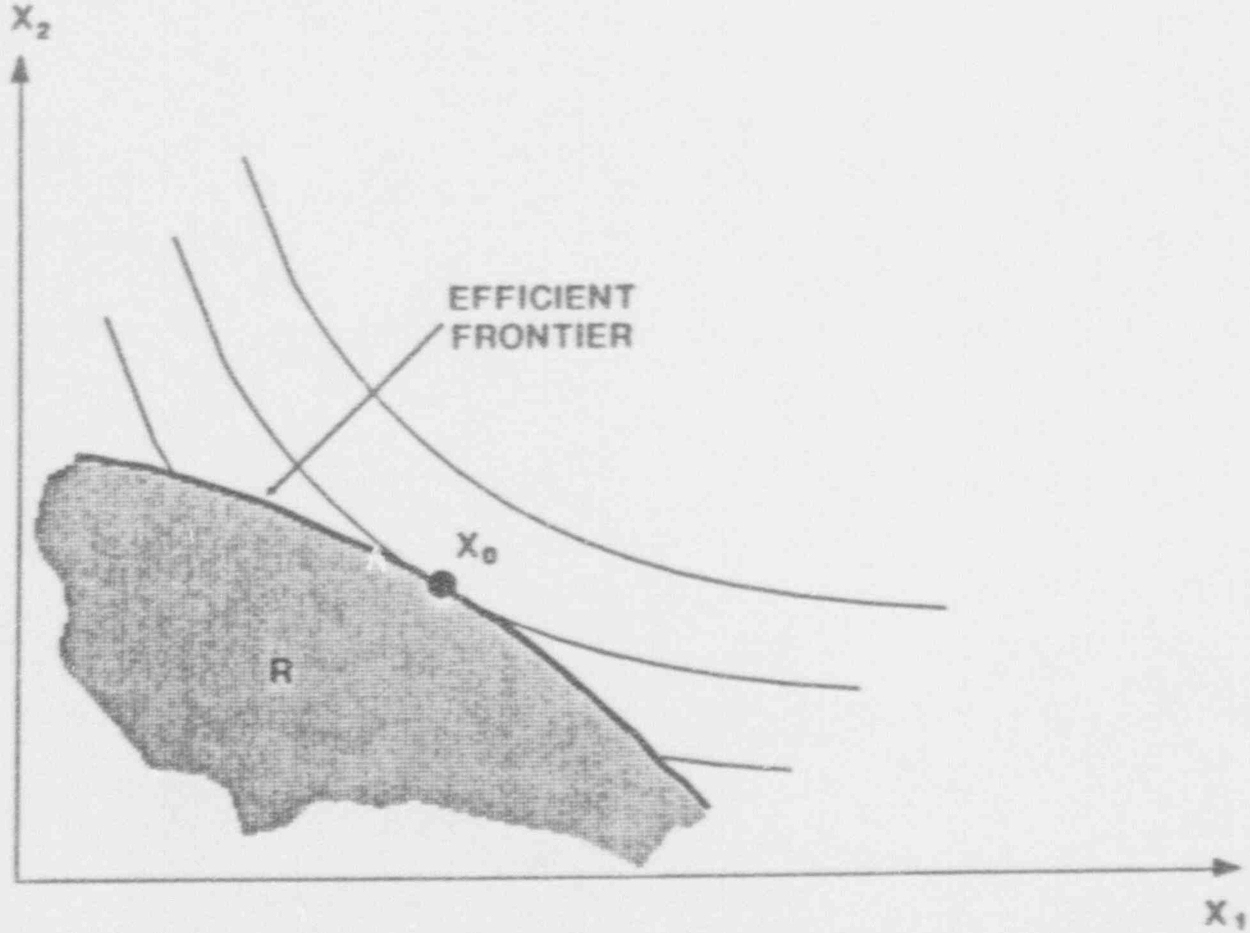

Figure 11 Optimal Point on the Efficient Frontier given a Family of Preference Curves

in a set of all acts being considered which maximizes v[x(a)]. In general, various linear and nonlinear programming techniques can be applied to this maximization problem when cast in a standard form. Furthermore, when a decision problem is cast in this format, there is a close relationship between these problems and the Theory of Games and Strategy which can also be solved using linear programming techniques.

Value functions are commonly used in making decisions today. For example, the cost/benefit ratio as currently used by the USNRC can be viewed as a value function. To illustrate this in 2 space, consider consequence X_1 to be the benefit due to internal event scenarios and consequence X_2 to be the benefit due to external events scenarios (that is, due to seismic, fire etc.) Thus, in this case, the cost/benefit ratio is defined as

$$CBR = Cost/(X_1 + X_2)$$

and the preference curves are the family of straight lines given by

$$X_1 + X_2 = constant$$

shown on Figure 12.

In a similar fashion the various load combination rules used for combining structural loads due to pressure, temperature, and seismic motion, etc. are also value functions. As will be seen later, a multivariate utility function is, in fact, a special case of a value function which allows for the input of the decision maker's preferences given a decision between actions whose outcomes are uncertain.

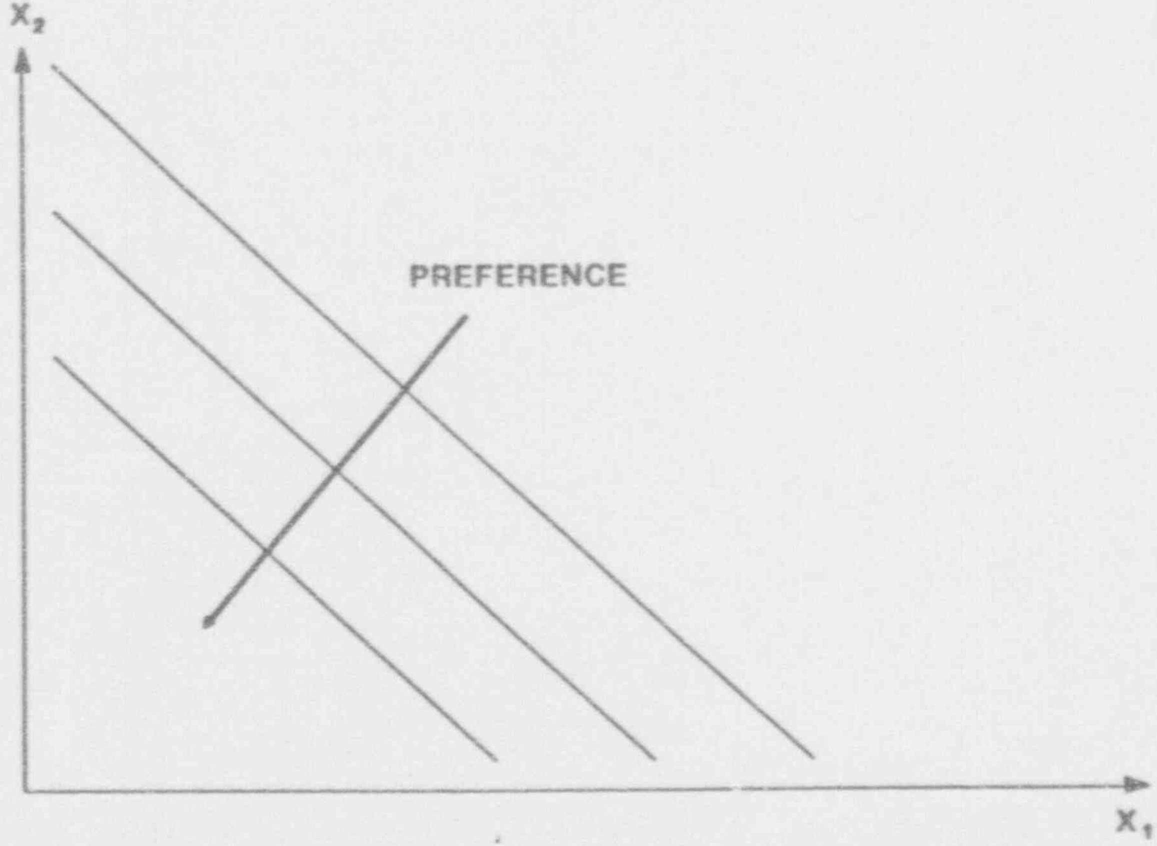

Figure 12 Preference Curves Implied by the Cost/Benefit Ratio

3.0 DECISIONS WITH UNCERTAINTY

3.1 The Utility Concept

It is when making a decision between several actions whose outcomes are uncertain that we must introduce new concepts to assist us in this process. The concept that has proved most useful in the commercial business sector in addressing this situation is that of the utility of a consequence. Much of the literature on the concept of utility has been expressed in terms of monetary value (say, dollars), and much of the original work was presented in an economic format. The basic concept behind the association of a numerical value of "utility" with each uncertain outcome associated with an action is that the value of the outcome to a decision maker is represented by characteristics that go significantly beyond the simple numerical (say, monetary) value of the outcome. That is, the value of the outcome to the decision maker depends on his current status, his perceived needs, and his value system, and the representation of these characteristics cannot be captured in a single monetary value.

Before going any further, an example to clarify this basic concept is presented. Consider two college students each having a total of $15 in his pockets. The first student has made plans to go bowling that evening, and he knows that $15 is an adequate sum of money to enjoy an evening of bowling with refreshments. The evening of bowling is his top priority. Both these students have a friend down the hall who is addicted to small games of chance and is constantly encouraging them to wager with him. For example, he approaches the student with the bowling plans with a wager that could double his money at very reasonable odds. The decision the student must make is whether or not to take the wager. Even though the odds might be quite reasonable, it is highly likely that the would-be bowler will decline to make the wager for the simple reason that the $15 he has in his pocket is adequate for his highest priority needs, and the winning of the wager to increase his net worth to $30 provides very little added value to him, given that bowling is his top priority. Furthermore, the chance that he will lose his entire fortune of $15 and hence not be able to bowl at all far outweighs the added value he would receive by winning the extra $15. Thus, in this case, the would-be bowler is highly likely to decline the wager.

By contrast, consider his roommate who also has $15 in his pocket. In this case, the roommate has made a date that weekend with a charming lady. However, he knows that to wine and dine this lady in suitable fashion will require considerably more than the $15 he has at his disposal. In fact, he estimates that it will take at least $30 to provide for a memorable evening with the lady in question. Thus, in this case, this student is highly likely to accept a wager with a chance of increasing his net worth to $30, as the $15 he has in his pocket is relatively valueless to him. Hence, he is likely to seek out their gambling friend and make a wager with the hope of increasing his net

worth. Furthermore, he is more than likely to accept the wager at less than even odds as his needs prescribe a significant increase in funds in order to meet his top priority objective.

Thus it can be seen that the added value of $15 that could accrue to each of these students has a significantly different value to each of them. In the case of the would-be bowler, he is more likely to be averse to taking any risks that might result in his being unable to pursue his bowling evening. In the case of his roommate, however, he is highly likely to accept a wager at any (even unreasonable) odds with the hopes of meeting his objectives. Thus the utility of that same $15 is entirely different to each of these students.

It is important to note that the odds of the wager being offered to the two students do play a role in whether or not they will take the wager. For example, in the case of the would-be bowler, if the wager was such that he was highly certain of winning, he would more than likely take the wager. Further, for his roommate, there are limits to the poor odds that he would accept even given his resource needs for his potential romantic evening. Thus whether or not each of these students would accept the wager depends upon the odds of the wager being offered. However, in the former case, much higher odds would have to be offered before he would choose to gamble than in the case of his roommate. As we will see, it is possible to query these two students and essentially determine the odds that would have to be offered to them before they would consider gambling away their current resources. In fact, a relationship between the odds being offered and the potential payoff of the wager can be developed for each of these two students which would prescribe whether or not they would choose to accept the wager. This functional relationship is known as a utility function and is the mechanism by which a consistent decision can be made between a set of actions in which the outcome of each action is uncertain (i.e., a wager with known payoffs).

The concept of the utility function was proposed and developed in the 1930s by von Neumann and Morgenstern in the seminal text, Theory of Games and Economic Behavior (Reference 8). Speaking of this work, Luce and Raiffa (in Reference 9) state:

> "only a very few scientific volumes as mathematical as this one have aroused as much interest and general admiration. Yet we know that much of the material has lain dormant in the literature for two decades. Presumably the recent war was an important contributing factor to the later rapid development of the theory. During that period considerable activity developed in .....such topics as logistics, submarine search, air defense, etc."

Discussions of the concept of utility and derivations based on different sets of assumptions are given in References 9 through 15.

Before dealing in more detail with the concept of the development of the utility function representing a decision maker's values and preferences, the basic distinction concerning the concept of the utility function must be emphasized. This point is that, if the consequences of an action are uncertain, we can compute the expected value of the consequences resulting from each action using the usual theories of statistics (and knowing the odds of each outcome and the payoff associated with each outcome), assuming that the game is played many times. By contrast, by developing a utility function for each of the outcomes associated with the uncertain consequences, we can make a decision as to whether or not we should even play the game and further, this decision can be made whether or not the game is only to be played once! (That is, utility theory provides a guide to decision making even when a one-time decision must be made.)

Consider now the general model of the decision-making process as applied where a decision must be made between two actions a' and a" as shown in Figure 13. As shown on this figure, as a consequence of action a' there are a number of possible outcomes, C'_i associated with the individual probabilities p'_i. Thus, this is a general model in which we have uncertainty involving a single scalar consequence, and our goal is to map each of these consequence values onto a single scalar utility value that reflects our goals, values and preferences. It is important to note that the probabilities associated with the outcomes of an action (the p') may be derivable from the laws of statistics and the knowledge of the particular uncertainty source or they may reflect the decision maker's belief in the likelihood of the various consequences and thus be primarily judgmental.

Given this general model, it can be shown that the appropriate choice between action a' and action a" is made by choosing that action that has the larger expected value of the scalar utility. That is, we choose a' over a" if

$$\Sigma\, p'_i u'_i \;>\; \Sigma\, p''_i u''_i$$

This follows from the construction of the utility function itself as described below.

In order to develop a utility function for a scalar attribute, let us first assume that we have three values of the attribute A, B, and C and furthermore assume that we prefer A over B over C (and also A over C) according to our value system. Further we assume that our preferences imply

$$u(A) > u(B) > u(C)$$

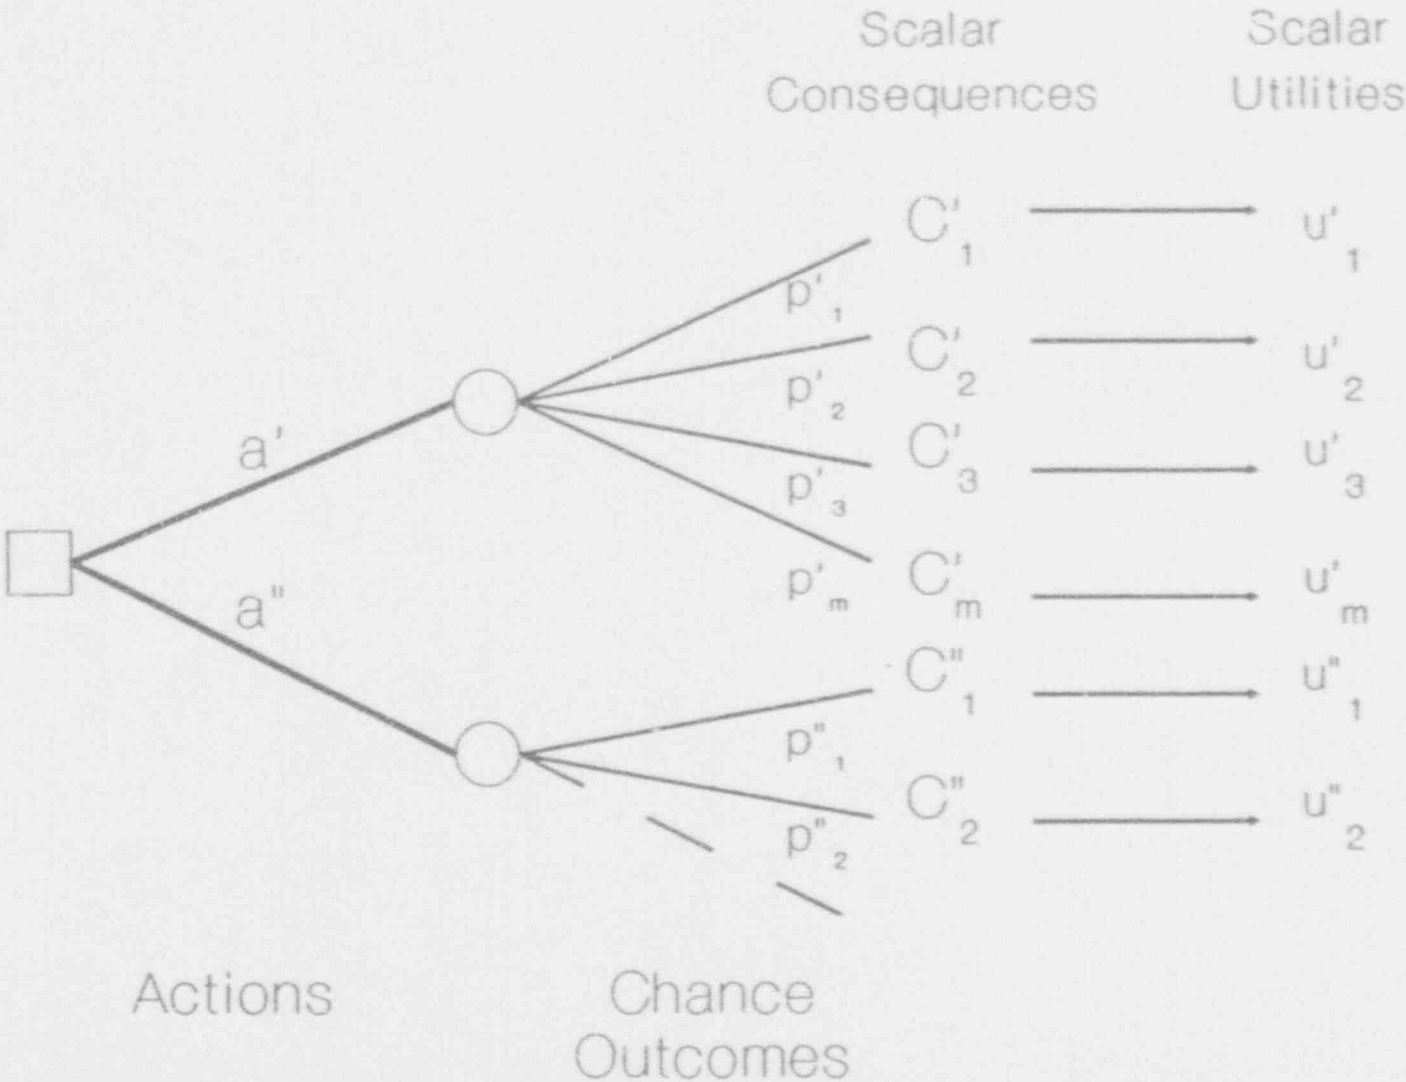

Figure 13 General Decision Tree with Uncertain Consequences of Each Action

That is, increasing values of utility are indicative of the more preferred actions or values. In general it can be shown that we may arbitrarily choose any two of the three values. Then, by querying the decision maker, we can find the probability p for which the decision maker is indifferent between two options, namely,

a) Receiving value B for sure, or

b) Receiving A with probability p or receiving C with probability (1-p) as shown by the decision tree in Figure 14.

Typically, the values of A and C are the most valuable and least valuable outcomes and are chosen to represent a reasonable physical range for the scalar attribute whose utility we are attempting to derive. And, as mentioned above, the values of u(A) and u(C) are arbitrary. B is some intermediate outcome value whose utility (to the decision maker) we seek to ascertain. Thus the choice between action a' which essentially consists of receiving value B (or keeping value B) and action a" (which is to accept the wager with the probability of winning p and the probability of losing 1 - p) hinges very much on the value of p (that is, the odds of the wager). Clearly there are values of p sufficiently high such that the decision maker is sure to take the gamble, and similarly there are values of p for which the odds of losing are so high that he is very unlikely to take the gamble and will instead choose action a' which involves no risk. There is some intermediate value of p - denoted the <u>indifference probability p*</u> - for which the decision maker is indifferent between sitting pat and receiving value B or taking the wager. Once the numerical value of the indifference probability p* has been identified, then the value of the utility of B is, by definition,

$$u(B) = p^* u(A) + (1 - p^*) u(C)$$

That is, the utility of B is the <u>expected value of the utility</u> of the wager of action a". This process is repeated for a number of intermediate values of B and for each intermediate value of B chosen, a new indifference probability p* is determined (by querying the decision-maker) and then the utility of each new value of B is computed as above. From these, a continuous curve of the utility of B can be constructed as shown on Figure 15. As a result, as shown by von Neumann and Morgenstern, the choice between any two actions with uncertain outcomes is made based on selecting that action which has the largest expected value of utility.

A number of aspects of utility must be emphasized. First, a reasonable range must be assigned to the scalar consequence whose utility is being evaluated. It is no use to query a decision maker over a range of the scalar attribute that is inconceivable to him. Second, in generating the

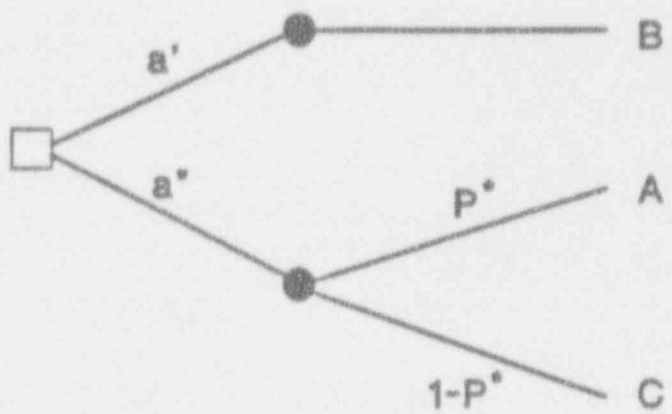

Figure 14 Decision Tree Used to Define u(B)

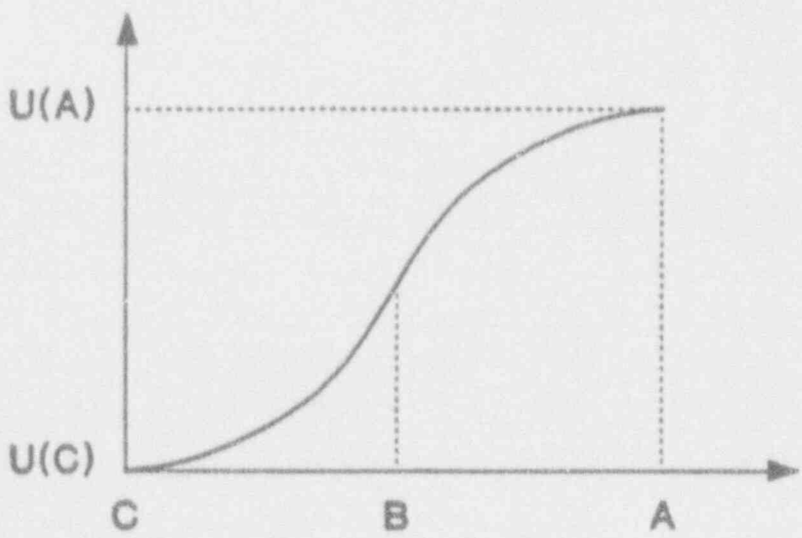

Figure 15 Sample Utility Functions

utility function, any two values of the utility function can be selected arbitrarily. That is, we may assume the utility function varies between 0 and 1 or -500 and +500 or any other particular range we choose. It is the intermediate values developed by the process described above that are important. Thirdly, it can be shown that a linear transformation of any given utility function is also itself a utility function. The implication of this is that if, after querying the decision maker, it is determined that his utility of a scalar consequence is linear with respect to that consequence, then the numerical value of the consequence itself represents the utility to the decision maker.

The shape of the utility function thus captures the decision maker's preferences, values, and goals and is the basis for making decisions between risky alternatives. To illustrate the effect of different shapes of utility functions, let us assume that $u(A) = 1$ and $u(C) = 0$ where it is recalled that the range of the utility function is arbitrary. Then the utility of any intermediate value B is given by

$$u(B) = p^* u(A) + (1 - p^*) u(C)$$

$$= p^*$$

Thus by the choosing the (arbitrary) limits of the utility function to be 0 and 1, we have forced the utility of B to equal the indifference probability itself. Consider the two utility curves shown in Figure 16. One is labeled a "risk averter" curve, and one is labeled a "risk taker" curve. To see why this is so, consider a point B' where the decision maker's bankroll (B') is relatively small. At this point, the ordinate on the curve (which gives the indifference probability) is relatively high for the risk averter. In other words, he wants relatively good odds before he is indifferent between betting and not betting. By contrast, the ordinate for the risk taker curve is relatively small. In other words, the risk taker is very likely to bet even though the odds are very much unfavorable to him.

Now look at a second point (B") where the bankroll is relatively large. In this case, the risk averter requires highly favorable odds (almost a sure thing) before making the decision to bet rather than standing pat with his bankroll (B"). By contrast, the risk taker will still be willing to gamble at any reasonable odds. So in effect, the risk taker's view of money is to "go for broke," while the risk averter's view is to "hang on to what you've got,"

Note that, as shown in this figure, the risk averter's curve is concave downward. It can be shown that the condition of risk aversion implies that his utility curve will always be concave. By contrast a risk taker's curve will always be concave upward. However, it is not necessary that the utility curve be always concave or convex, but it may

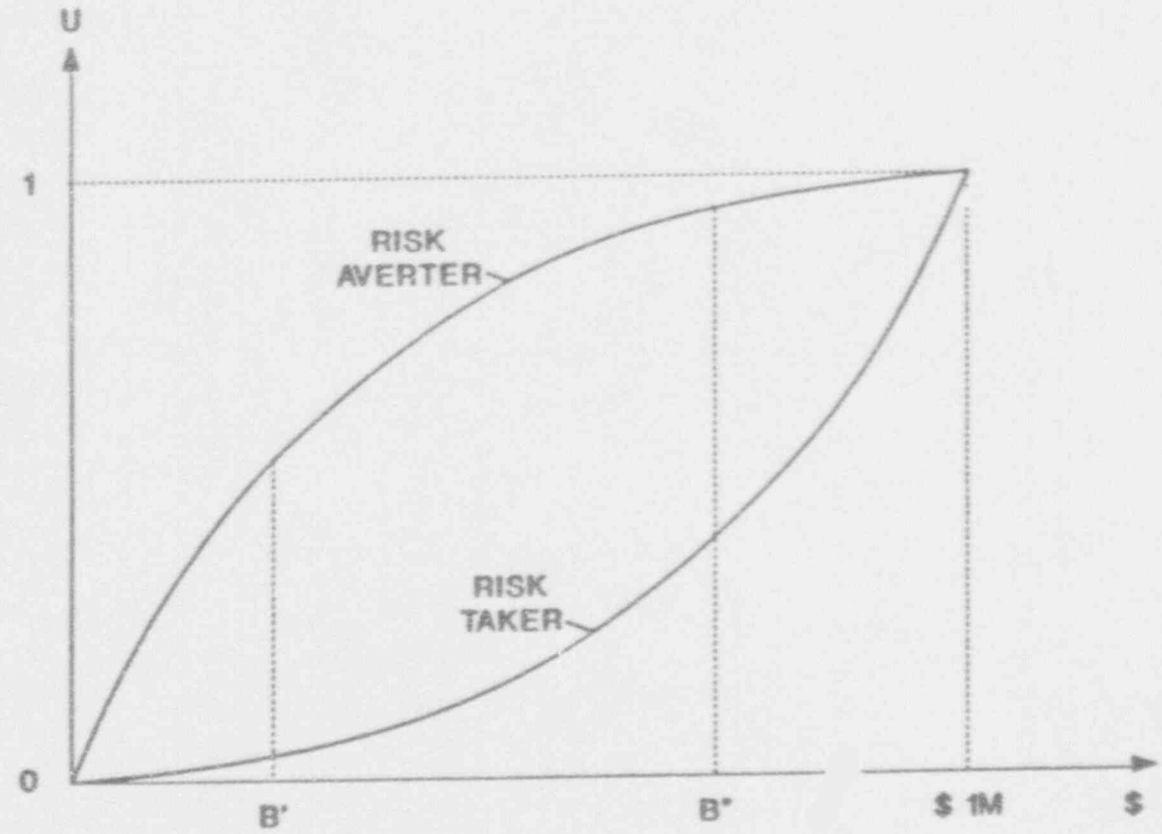

Figure 16 Example of Risk Averter and Risk Taker Utility Functions

vary over different anges of the consequence attribute. This can only be determined by the results of the querying process. Finally, it should be noted that it is possible that the utility curve is linear. This case is denoted as the "risk neutral" case and, as mentioned earlier, the implications are that the decision maker would be willing to accept any wager provided the odds of the wager and the payoff were fair.

In summary, the above interpretation, cached in a gambling format, shows that the decision between making a wager and not making a wager depends only on the decision maker's utility curve. The utility curve is developed based on his values, goals, and preferences, and is totally independent of the nature of the wager (odds, number of options, payoffs, etc.,). Thus, once the utility curve has been established for a given decision maker, then the utility curve forms the basis for his decision as to whether or not to bet - depending only on the odds of the wager (i.e., on the probabilities of the uncertain consequences associated with each action).

The above interpretation was presenced in terms of a decision between gambling and not gambling. From the viewpoint of the NRC, the consequences of importance have to do with the risks posed to the public by commercial nuclear power plants, and, with respect to the resolution of generic issues, the decision to be made is whether or not a certain proposed retrofit (resulting in a given number of person-REMs averted) is cost effective. And decision makers within the NRC must assess the relative value of these retrofits given that the retrofits result in consequences which are uncertain. Development of the utility function on person-REMs averted would allow for the use of utility functions and formal decision-making methodology in assessing the relative value of different proposed retrofits. For example, an NRC regulator might view incremental achievements in person-REMs averted differently when they arise from different sources, i.e., from core damage events resulting from random accident scenarios, from fire-induced accident scenarios, or from earthquake-induced accident scenarios. Thus, it might be true that the NRC regulator could view only large increments in seismic risk averted as being significant. By contrast, he might view any increments in person-REM averted arising from accident scenarios due to fires as being equally significant, and he might very well view even small increments in averted risk due to random or internal events accident scenarios as highly significant. In effect, if these views were to be established by interrogation of the decision maker, it would imply that he is a risk averter in regards to internal events, risk neutral in regards to fire events and a risk taker in regards to earthquake-induced accidents, and his utility functions for these three attributes would be as shown schematically in Figure 17. In effect, the decision maker is implying that for a given amount of money to be spent in making retrofits, he is willing to "bet on" his belief that the next core damage accident in the US commercial nuclear power industry will arise from random or human error related events, and thus he puts his largest "value" on those retrofits which avert person-REMs due to internal events sources.

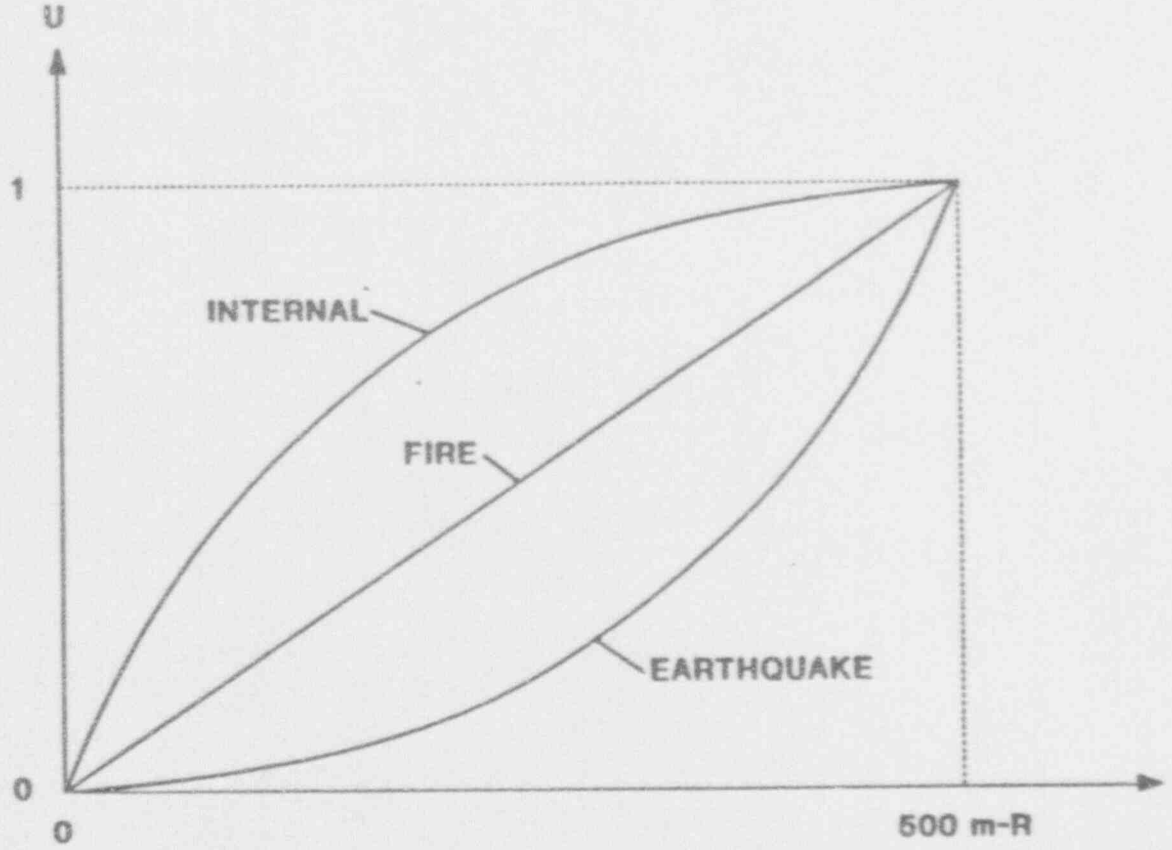

Figure 17 Schematic of Different Utility Functions For Internal, Fire and Seismic Events.

Similarly, he is implying that retrofits primarily affecting person-REMs averted due to fire-induced accidents sequences as being more important than person-REMs averted due to earthquake-induced accident sequences.

3.1.1 Extension of Utility Analysis To Incorporate Costs

In keeping with NRC's current thinking on using a cost/benefit ratio as a measure of viability of a retrofit, it is possible to extend the utility analysis to include costs directly. This could be done in several ways.

a) Consider each outcome as a multivariate consequence, with cost being one of the attributes, or,

b) Extend the concept of a utility-based decision-making criterion to allow for normalization by cost, i.e., we will choose actions a' over a" if

$$\frac{\Sigma p'_i C'_i}{\text{Cost}(a')} > \frac{\Sigma p"_i C"_i}{\text{Cost}(a")}$$

which, in effect, normalizes the expected value of the utility of each uncertain consequence vector C_i by the costs associated with that action.

This second approach is followed in the examples in this report, as it most closely ties in with current NRC thinking on assessing the viability of retrofits. Further, this criterion reduces to the usual cost-benefit ratio when the decision-maker is risk neutral about the uncertain consequences (i.e., his utility curve for the C_i is linear). However, the second approach is a new mode of application of the utility concept of decision making, and as shown below, implies a restriction on the families of utility functions to be generated. This restriction, however, poses no lack of generality.

In order for the normalized utility to be a valid criterion for decision, it must be indifferent to the scale selected for the utility function. This can only be true if we restrict the choice of utility functions to the family of functions for which $u(C) = 0$ and $u(A) > 0$. This implies that, in querying the decision maker as to his utility function on the consequence, one is free to choose u(A) but not u(C). Then any resulting decision (i.e., ranking of the actions) will be independent (not changed) by the choice of u(A). This follows since, due to requirement that $u(C) = 0$, a different choice of u(A) gives a positive scaling:

$$\bar{u} = ku$$

and if, for the initial choice of scaling,

$$\frac{\Sigma p'_i u'_i}{C'} > \frac{\Sigma p''_i u''_i}{C''}$$

then this implies (for any other scaling)

$$\frac{\Sigma p'_i \bar{u}'_i}{C'} > \frac{\Sigma p'_i \bar{u}''_i}{C''}$$

provided $k > 0$. Thus no matter what the scale of the utility function initially chosen, the same decision between two alternatives will be selected.

The restriction that $u(C) = 0$ is physically reasonable in the context of cost/benefit type decisions, for which a negative benefit (at whatever cost) is of no interest to the decision maker, and clearly a value of zero benefit is of no utility. Thus the requirement that $u(C) = 0$ is practical and physically meaningful.

3.2 Typical Forms and Querying Process

As described above, the shape of the utility function indicates whether or not the decision maker is risk prone, risk averse, or risk neutral (at least locally). A full discussion of the different shapes of utility functions is presented in Reference 10. In general, the exact shape of the utility function can be determined by a sufficient amount of querying of the decision maker. However, there exists a family of utility functions denoted as the exponential-linear functions which, by varying a single parameter, can model varying degrees of risk proneness or averseness. In fact, according to Reference 10,

> "experience has shown that (such) fine tuning is rarely required for the single attribute utility functions when they are part of a multi-attribute formulation. It will almost always suffice to use a single parameter utility function."

The linear and exponential functions that can be used in this context are shown below:

$$U(x) = a + b(-e^{-cx})$$

$$U(x) = a + b(cx)$$

$$U(x) = a + b(e^{cx})$$

where the constants a and b are used to select the scale of the utility function and c is a positive constant (assuming that increasing values of the attribute have increasing utility). The value of c is determined by querying the decision maker and is a measure of his degree of risk proneness or risk aversion. This family of exponential and linear utility functions is shown in Figure 18. To determine the shape of a utility function by querying the decision maker, we assess the utility of the attribute (in the decision maker's view) for several points on the utility curve and then adjust the parameter c to provide an appropriate fit.

As an example, consider developing a utility function for an attribute x which varies between 0 and 100. We would first ask the decision maker whether he would prefer a 50-50 wager having a payoff of either 0 or 100 as compared to receiving a sure payoff of the midpoint of the range, namely 50. If the decision maker prefers the numerical value of 50 over the wager, this would indicate that he may be risk averse. The same question might be posed for other subranges of the attribute x, for example, the range 25-50 or 50-100. The responses of the decision maker would then determine whether he is risk averse, risk neutral or risk prone. (Or perhaps risk averse for some range of values of the attribute and risk prone in other portions.) However, assuming that he is one or the other over the entire range of the attribute (which is the usual case) then one can select a value of the constant c to specify his utility preferences. For example, we might pick two values of the attribute, say 25 and 75, and ask the decision maker for what intermediate value x^* of the attribute would he be indifferent to receiving the value x^* or accepting a 50-50 wager with the payoff of either the value 25 (if he loses) or 75 (if he wins). Assume that the decision maker indicates that for values of x less than 40, he would be more likely to accept the wager, whereas for values of x greater than 40, he would be more likely to accept the sure payoff of 40. Thus his indifference probability is $p^* = 0.5$ for $x = 40$. Using the defining equation of the utility function, one thus has a single equation

$$U(40) = 0.5\ U(75) + 0.5\ U(25)$$

which is used to determine the value of the constant c in the exponential equation for a risk adverse utility.

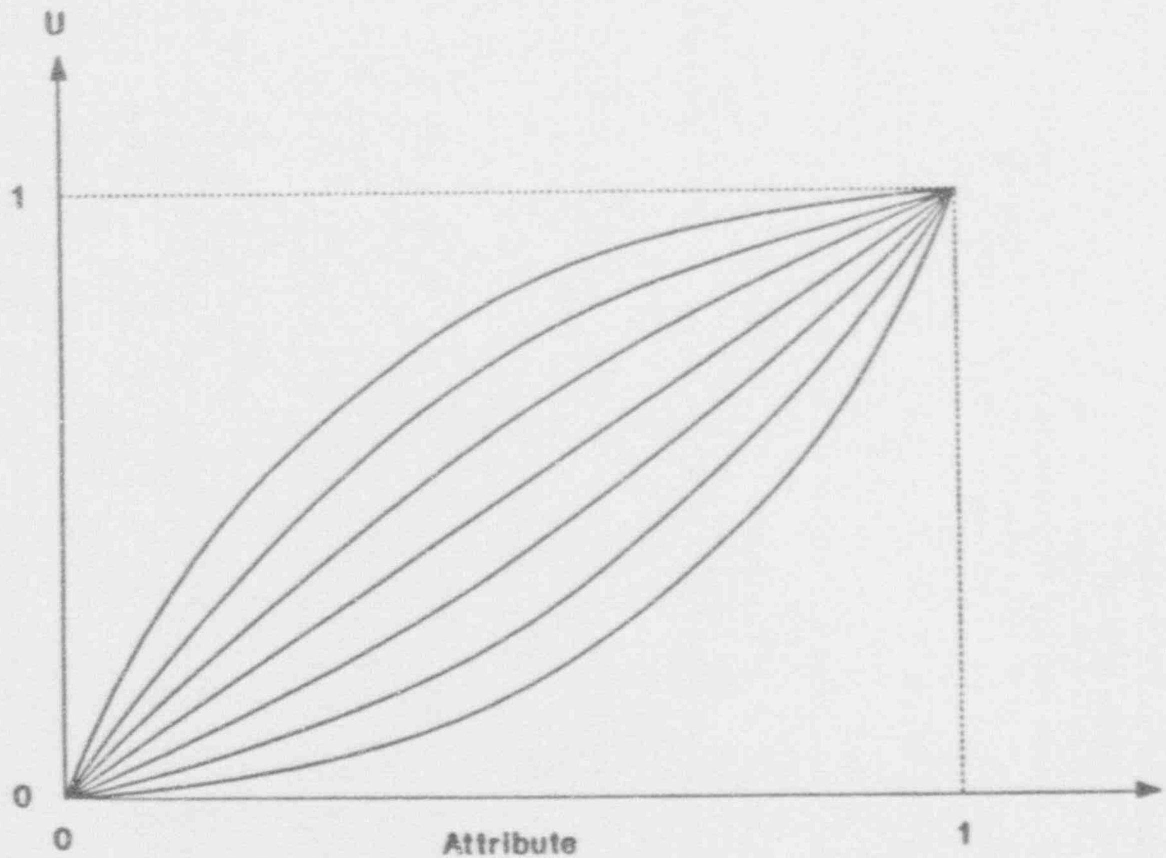

Figure 18 Exponential (and Linear) Utility Curves

It must be reiterated that, while the exponential forms of the utility function have been found to be useful in a wide variety of circumstances, the use of such forms is not essential. If the decision maker was found to be risk averse in one range of the attribute and risk prone in another range of the attribute, the precise form of the utility function could be determined by extensive questioning of the decision maker.

In addition, in all the above discussion, it has been assumed that increasing values of the attribute are preferred, and hence we are dealing with a monotonically increasing utility function. It is a simple extension to consider decreasing values of an attribute (for example, costs) as being preferred. In fact, in this case, the constant c for the exponential family of utility functions is always negative, but is determined in the same fashion as described above.

3.3 Application to the GE Plant

Referring back to the GE plant data given in Section 1, it is shown on Table 2, that of the four modifications proposed for the GE plant, only modifications 1 and 4 are potentially cost effective. Note further that the benefit in terms of person-REM averted for modifications 1 and 4 is due only to the seismic root causes. In this example, we will illustrate how decision analysis can be used to make a meaningful distinction between these two modifications.

First, let us assume that the responsible decision maker has been queried as to the utility of averted seismic person-REM, and that the utility function derived from this querying process corresponds to that shown in Figure 19.

The decision analysis tree for this example is shown in Figure 20. In this case we have two potential actions (Modification 1 and Modification 2 with corresponding costs of $66K and $101K, respectively). Note that the consequence of each of these actions has three uncertain outcomes: an earthquake at seismic level 1, at seismic level 2, or at seismic level 3. Thus the uncertainty considered in this example is the uncertainty in the size of the earthquake which is likely to affect the site. The probabilities of obtaining earthquakes in these three levels were derived from the seismic hazard curve for the site in question, which was shown in Figure 3 in Section 1. The SSE for the GE plant is 0.12 g. From this figure, it can be determined that (if an earthquake occurs occurs at the site) there is a 77% chance of a level 1 earthquake (as defined in Section 1.2), a 19% chance of a level 2 earthquake and a 4% chance of a level 3 earthquake (where the annual frequencies are normalized to the annual frequency of the OBE - 0.06 g). These are denoted as the earthquake split fractions in Fig. 20 and in subsequent examples.

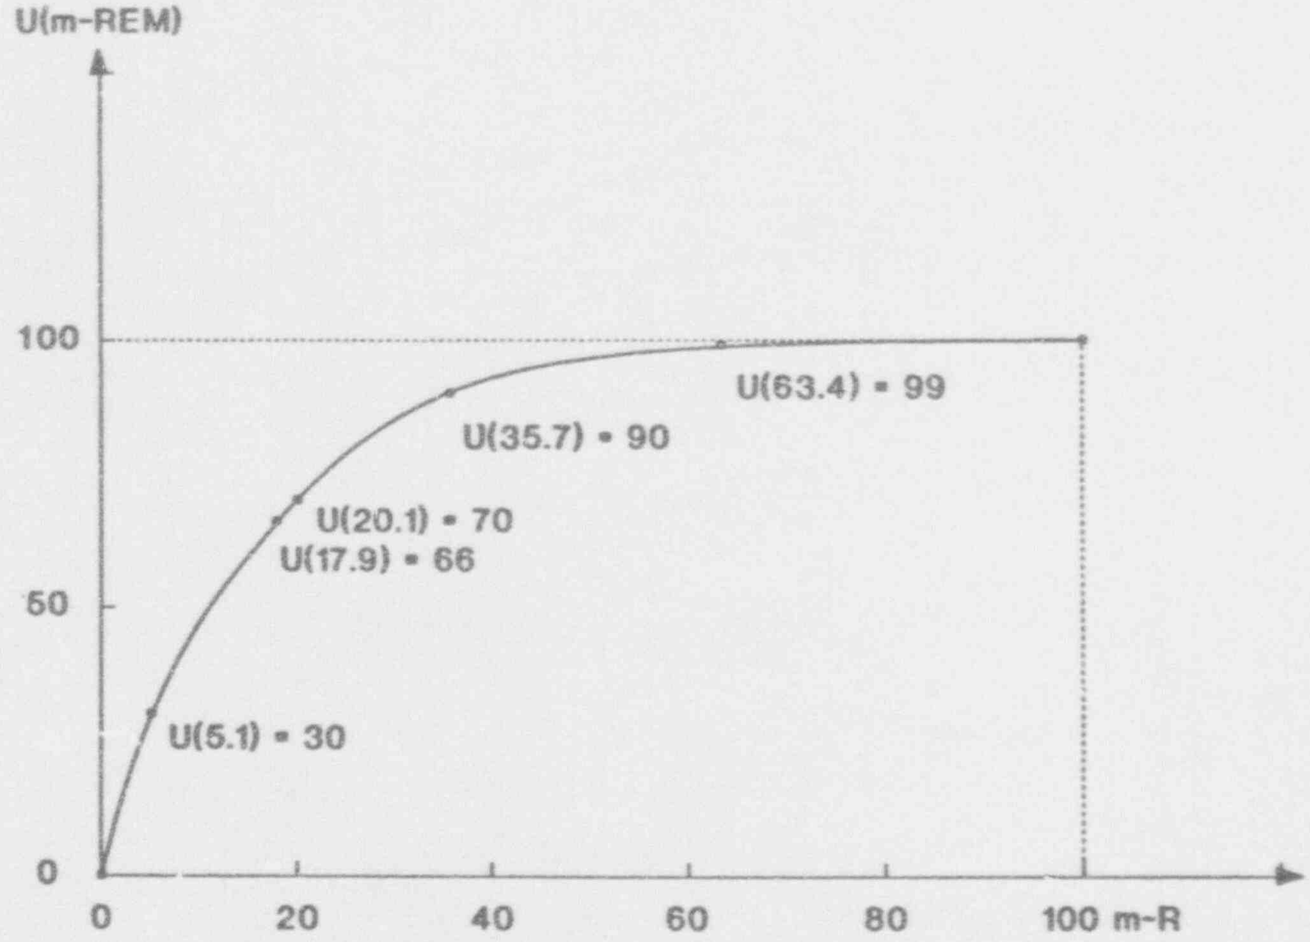

Figure 19 Assumed Utility Curve for Example of Section 3.3

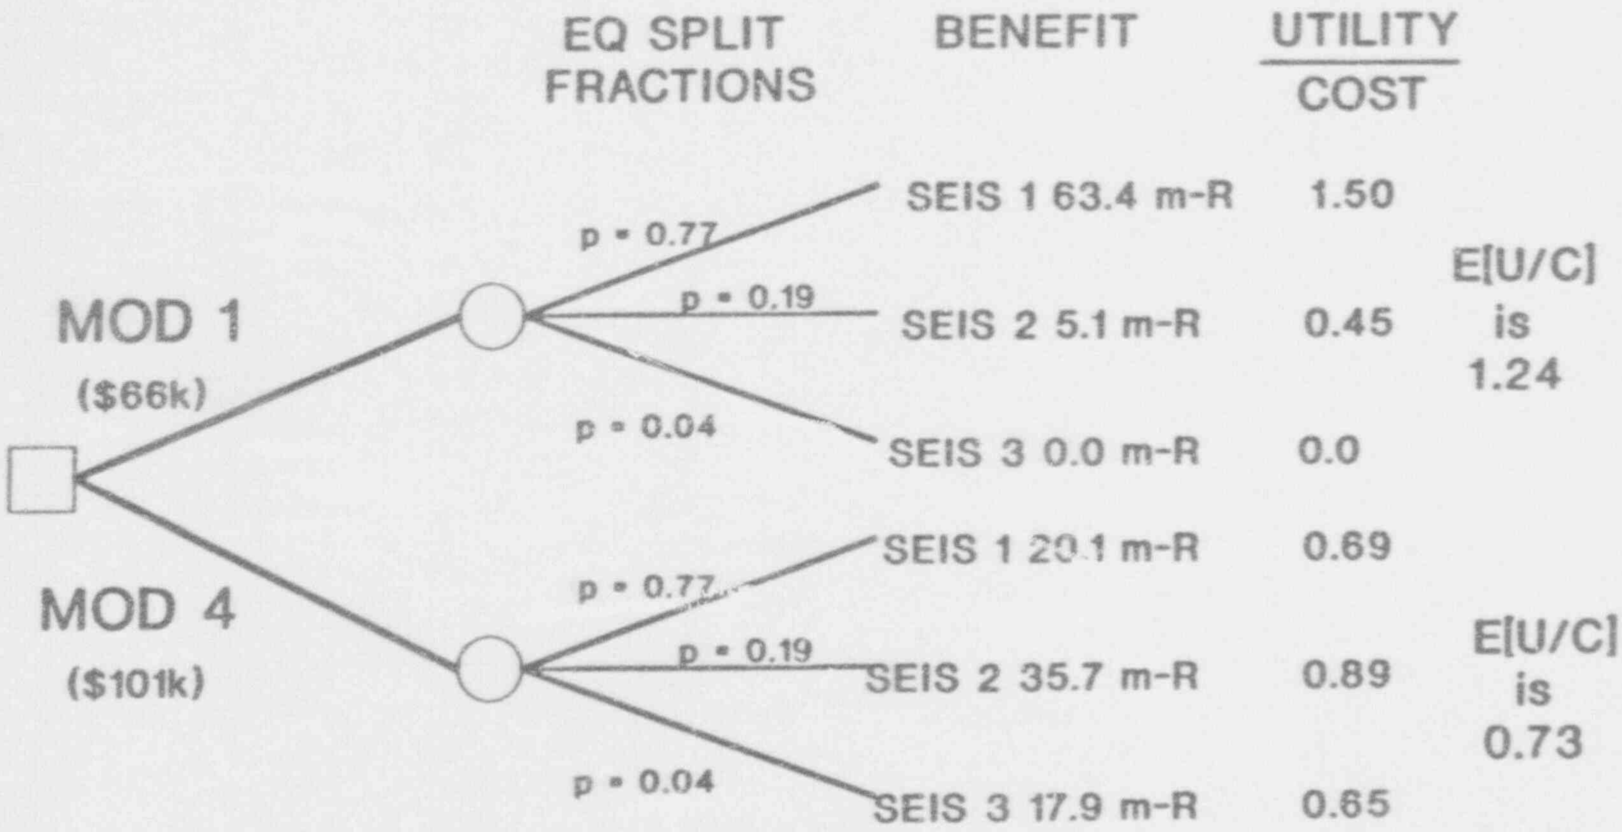

Figure 20 Decision Tree for Example in Section 3.3

The benefits in terms of person-REMs averted for each of the seismic level outcomes are also shown in the figure. For example, given that we have made the retrofit identified as Mod 1 for a cost of $66K, the person-REM averted is 63.4 person-REM given that an earthquake in seismic level 1 has occurred, and so forth. From the utility function for person-REM averted that we have assumed for this example (Fig. 19), numerical values of the utility of the various person-REM averted can be computed. (The points coinciding with these numerical values are shown on Figure 19.) Finally, at the right hand side of the decision tree is shown the utility of each outcome normalized by the cost of the modification and then the expected value of normalized utility is shown. As can be seen for Modification 1, the expected value of the utility normalized by cost of Modification 1 is 1.24. Similarly for Modification 4, the expected value of the normalized utility is 0.73. Thus, in this case, both the cost benefit analysis and the utility analysis indicate that Modification 1 is the appropriate choice to make.

To illustrate the potential use of the utility analysis, however, let us make one minor change in the example discussed above. Let us assume that Modification 4 is associated with a cost of $77K rather than $101K. Assuming a cost of $77K implies that both Modifications 1 and Modifications 4 have a BPB ratio of 1.04. The decision tree is now as shown in Figure 21. Thus, in this case, the cost/benefit analysis would be unable to distinguish between these two modifications. However, when we repeat the utility analysis, the only change comes in normalizing the computed utilities by the costs. Now, the expected value of the normalized utility for Modification 1 is 1.24 (as before) but now, the expected value of the normalized utility for Modification 4 has a value of 0.96. Thus, in this case, even though a cost/benefit analysis would indicate that both modifications were equally viable, a subsequent utility analysis specifies that Modification 1 is the appropriate choice.

This example does not imply that decisions based on utility theory are "better" than those based on a cost benefit analysis (as indeed the numbers in the example could be altered to yield identical values of mean utility for each modification, but different cost benefit ratios.) Whichever approach taken should be that which the decision-maker feels is most appropriate for the decision to be made in terms of reflecting his subjective feelings as to the underlying factors on which the decision is to be made, and his ability to separate and quantify these factors.

However, this example illustrates how the utility analysis concept can be used in conjunction with the cost/benefit concept in a meaningful fashion, in two ways. First, a utility analysis tells the decision-maker which of two or more actions he should take (given his expressed preferences as encoded in his utility functions). It does not tell him whether, indeed, <u>any</u> of the actions should be taken, or no action should

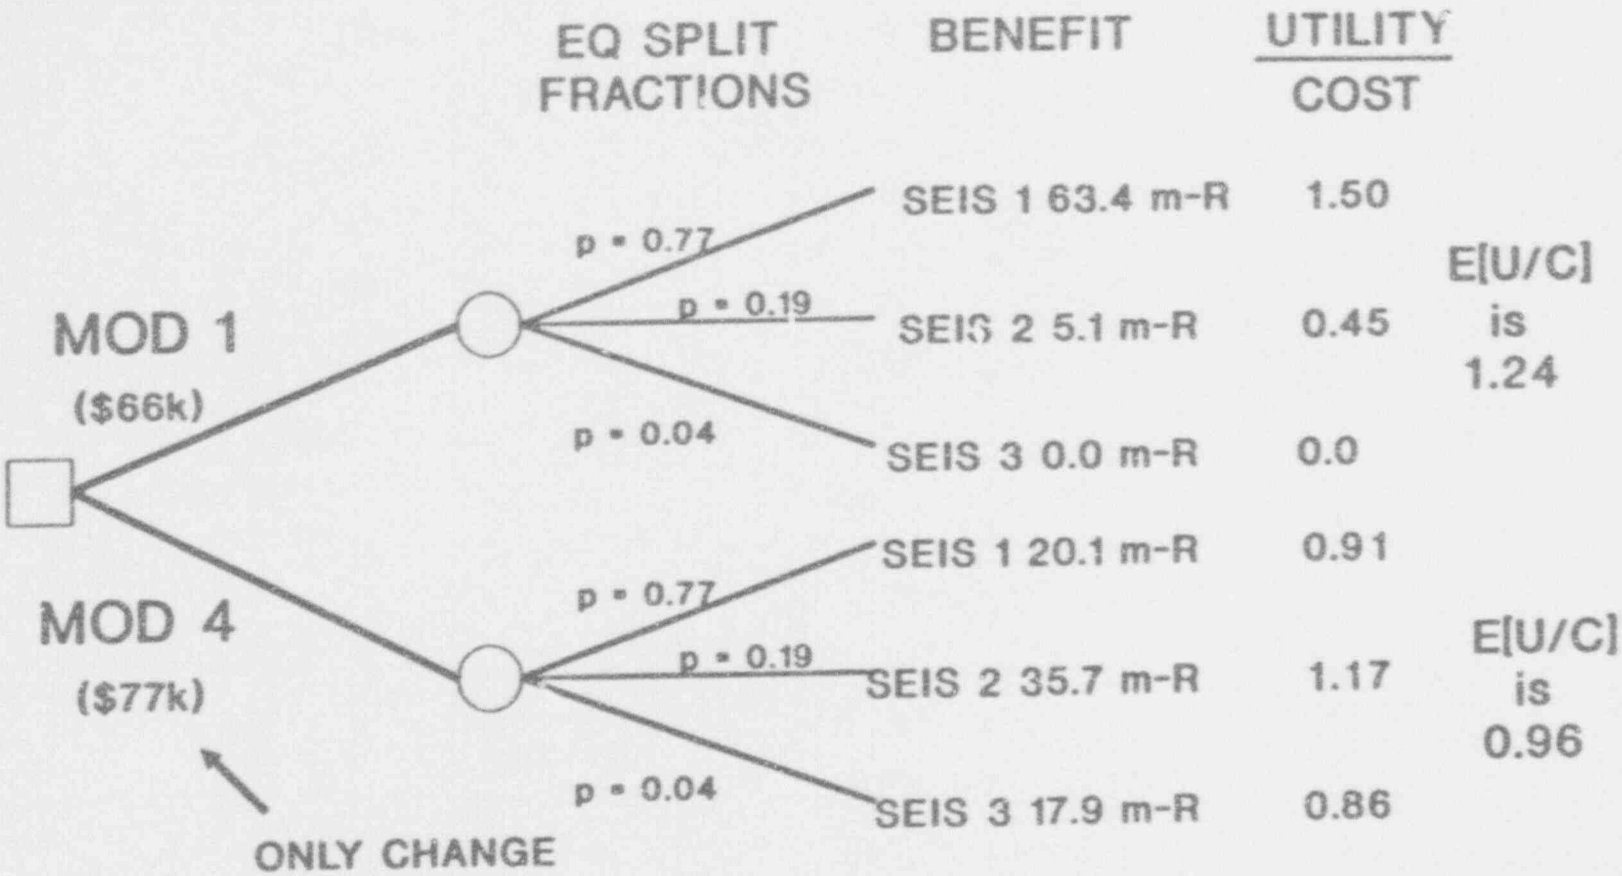

Figure 21 Modified Decision Tree for Example in Section 3.3

be taken at all. The cost benefit analysis (comparing the CBR to the criterion of $1K/person-REM) does indicate whether any action or no action should be taken.

Second, the utility analysis includes the magnitudes of the various consequences (i.e. the benefits in terms of person-REM averted) in an explicit fashion (which the cost benefit analysis does not). That is, the benefits in terms of person-REM averted are, in effect, weighted by the utility of the benefit according to the decision maker's preferences. This is then normalized by the cost of the modification. Thus, the utility analysis would differentiate between two actions having the same cost benefit ratio, but for which the magnitude of the benefit (and associated cost) was substantially greater for one action than the other. (That is, it would distinguish between an action with small cost and small benefit, and a second action having large cost and large benefit - but both having the same cost benefit ratio.)

These points suggest that use of a utility analysis in conjunction with a preliminary cost benefit analysis would be a useful approach to the US NRC's retrofit decisions. The cost benefit analysis would be used to identify those actions which would be considered viable - if any. (For example, these might be all actions with a cost benefit ratio less than 1.5 $K/person REM*.) Then, the utility analysis would be used to choose the appropriate action (from the viable set) which reflected the decision-makers preferences. And this choice would no longer necessarily match the cost benefit ratio ranking. Indeed, in examples later in the report, it will be shown that the use of such an analysis can actually change the order of ranking of various actions from that dictated by a simple cost/benefit analysis.

3.4 Multivariate Utility Functions

We now consider the most general case involving decisions between actions associated with multiple uncertain outcomes, and for which each outcome is characterized by a vector of consequences, as illustrated by the decision tree in Fig. 22. In this case, we must make use of a multivariate utility function which maps the components of each vector of consequences onto a single scaler utility value. The means of performing the mapping from the vector of consequences to a single scalar utility is the multivariate utility function discussed below.

In general, we need to construct a function of the form

$$u[\{C\}] = F[u_1(C_1),\ u_2(C_2),\ \ldots\ldots,u_m(C_m)\ ]$$

where {C} is the vector of multiple attributes, i.e.,

* The exact CBR cutoff value for identifying a viable action would have to be chosen after the adopted preference structure (i.e., utility function) was selected.

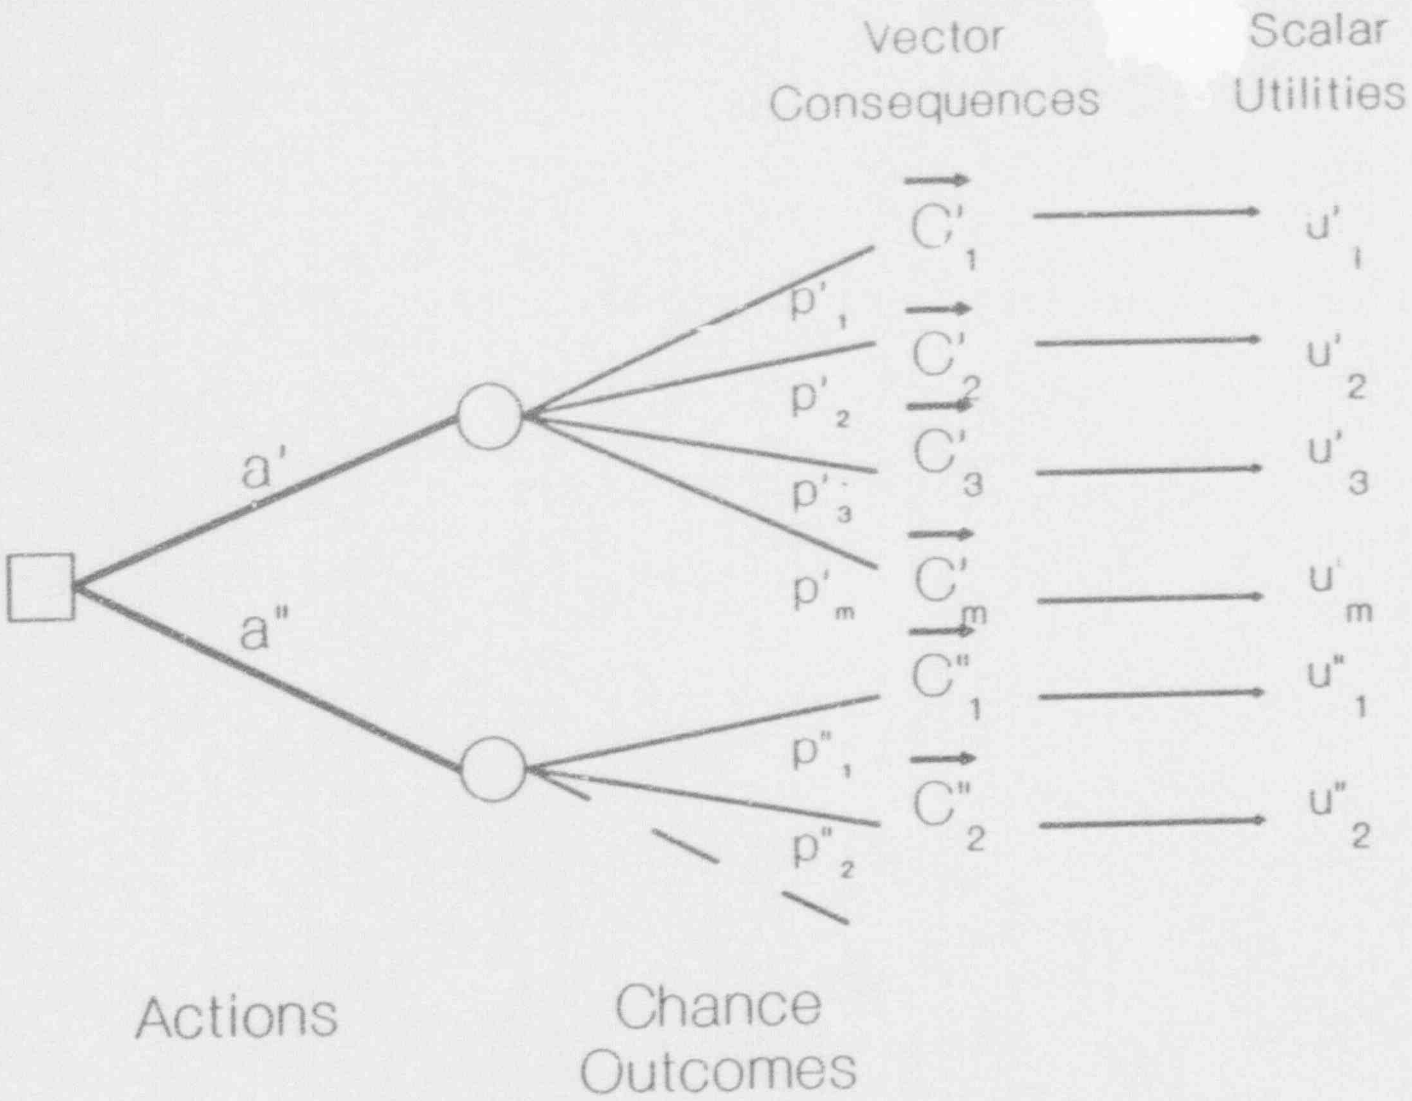

Figure 22 General Decision Tree with Uncertain Consequences of Each Action

$$(C) = (C_1, C_2, \ldots\ldots, C_m)$$

and the $u_i(C_i)$ are the individual (marginal) utility functions for each of the attributes C_i. It can be shown (see, for example, Reference 10) that the form of the multivariate utility function is determined by the concepts of

a) preferential independence,

b) weak-difference independence,

c) utility independence, and

d) additive independence.

Whether or not one or more of these independence conditions hold is determined by querying the decision makers. Depending on which combination of independence conditions can be established, the form of the multivariate utility function can be prescribed to within a number of constants to be determined based on the decision maker's relative preference for the individual attributes.

The simplest multivariate utility function is the so-called additive multivariate utility function which has the form

$$u[(C)] = k_1\, u(C_1) + k_2\, u(C_2) + \ldots\ldots$$

where the k_i are scaling constants. This form of multivariate utility function is appropriate if and only if the attributes C_i are additive independent as determined by querying the decision maker.

There are other forms of multi-attribute utility functions; for example, the multiplicative multivariate utility function which is the appropriate form when certain conditions as to both preferential independence and utility independence of the attributes can be established. A full discussion of the various forms of multi-attribute utility functions and the general techniques which may be used to establish the correct form by querying the decision maker is given in Reference 10.

In general, the k_i constants associated with a multi-attribute utility function indicate the decision maker's preferences (value tradeoffs) between various pairs of attributes. Once the form of the multi-attribute utility function has been determined (as a function of the utility functions of the individual attributes), then one develops a set

of independent equations (equal to the number of unknown constants) which are solved for the values of the unknown constants. To develop these equations, the decision makers are queried as to their perceived value of various pairs of attributes. For example, the decision maker is queried to find the values of two consequences that (in his view) are equally preferable (i.e., that he is indifferent between). Then the multi-attribute utility function - evaluated with these two consequence values - must yield the same numerical value of utility and thus one equation is available. This is repeated so as to generate a sufficient number of equations to solve for all the unknown constants. Again, the reader is referred to Reference 10 for a full discussion of the methods of querying decision makers and examples of such dialogue.

As an example of an multivariate utility function, let us consider the fivefold multi-attribute consequence vector consisting of the person-REM averted due to internal events, fire events, and seismic events for three different seismic levels. Let us hypothesize that, by querying the decision maker, it has been determined that

a) In the decision maker's view, the (marginal) utility of each of the consequence quantities (taken individually) are linear. That is, the decision maker is risk neutral as to increments of benefit accrued by making modifications which affect internal events, or seismic events, or fire events. As discussed earlier, when the utility functions are linear, then the actual value of the benefit itself may be taken as the utility function.

b) It is assumed that, in the decision maker's view, risk averted benefits accrued from internal events or fire events are equally preferred (or are equivalent in utility) but that risk averted benefits due to internal events or fire events are preferred in a ratio of 10 to 1 over seismic risk averted. This implies that, in the decision maker's view, the next core damage accident is more likely to be due to an accident involving internal (random) event scenarios or fire event scenarios than to be due to an earthquake at the power plant, and that (crudely speaking) the decision maker is likely to bet at 10 to 1 odds that such will be the case.

c) In addition, it is assumed that the decision maker prefers modifications affecting risk from the three seismic levels (Seis 1, Seis 2, and Seis 3) in the ratios of 10:5:1. Again, this can be crudely interpreted as stating that the decision maker believes that if an earthquake occurs affecting the plant, he is willing to give the above stated odds that it is more likely that the earthquake will be in the range OBE to 2 SSE than the higher earthquake levels.

Given the preferences stated above, one can equate the utilities of the various equivalent consequences described by the preference assumptions with the result that the appropriate additive multivariate utility function is given by

$$U[(C_i)] = 10 \text{ M-R(internal)} + 10 \text{ M-R(Fire)}$$
$$+ \text{M-R(Seis 1)} + 0.5 \text{ M-R(Seis 2)} + 0.1 \text{ M-4(Seis 3)}$$

(It is, of course, assumed that the necessary additive independence assumptions have been established for all pairs of the independent consequences.) And note again, that the benefits in terms of person-REM averted are the utilities in this expression, since by virtue of the assumption that all marginal utilities are linear, the benefit itself is equivalent to the utility. This multi-attribute utility function will be employed in the next two examples.

3.4.1 Application to the GE Plant

Consider now the benefit in terms of person-REM averted for the GE plant as presented in Section 1 (see Table 2). As shown on that table, the cost/benefit ratio analysis indicates that Modification 1 is preferred.

Consider now the decision tree shown in Figure 23. In this case, there is no uncertainty in the consequence vectors associated with each modification. Mapping the consequence vectors associated with each of the three modifications results in scalar utilities of Modifications 1, 3, and 4 of 0.98, 1.28, and 0.29 respectively, as shown on this figure. Thus it can be seen that, making use of the example additive multivariate utility function derived above, the decision analysis now ranks Modification 3 as being the most cost effective whereas using the cost benefit ratio approach, Modification 1 is preferred. Thus, this is an example whereby the ranking of the modifications is altered because of the preferences in terms of value tradeoffs between attributes specified by the decision maker. In effect, the decision maker's preferences have altered the order of choice in recognition of the fact that Modification 3 involves substantial contributions of person-REM averted due to internal event and fire event scenarios, whereas Modification 1 and Modification 4 result in person-REMs averted only due to seismic-induced scenarios. In effect, the decision maker is "betting" on the source of the next core damage accident and using his belief to help shape his choice of the most effective use of his resources in averting risk to the public due to future accidents. (Note also that since no uncertainty is involved, the multivariate utility function is serving as a value function as described in Section 2.4).

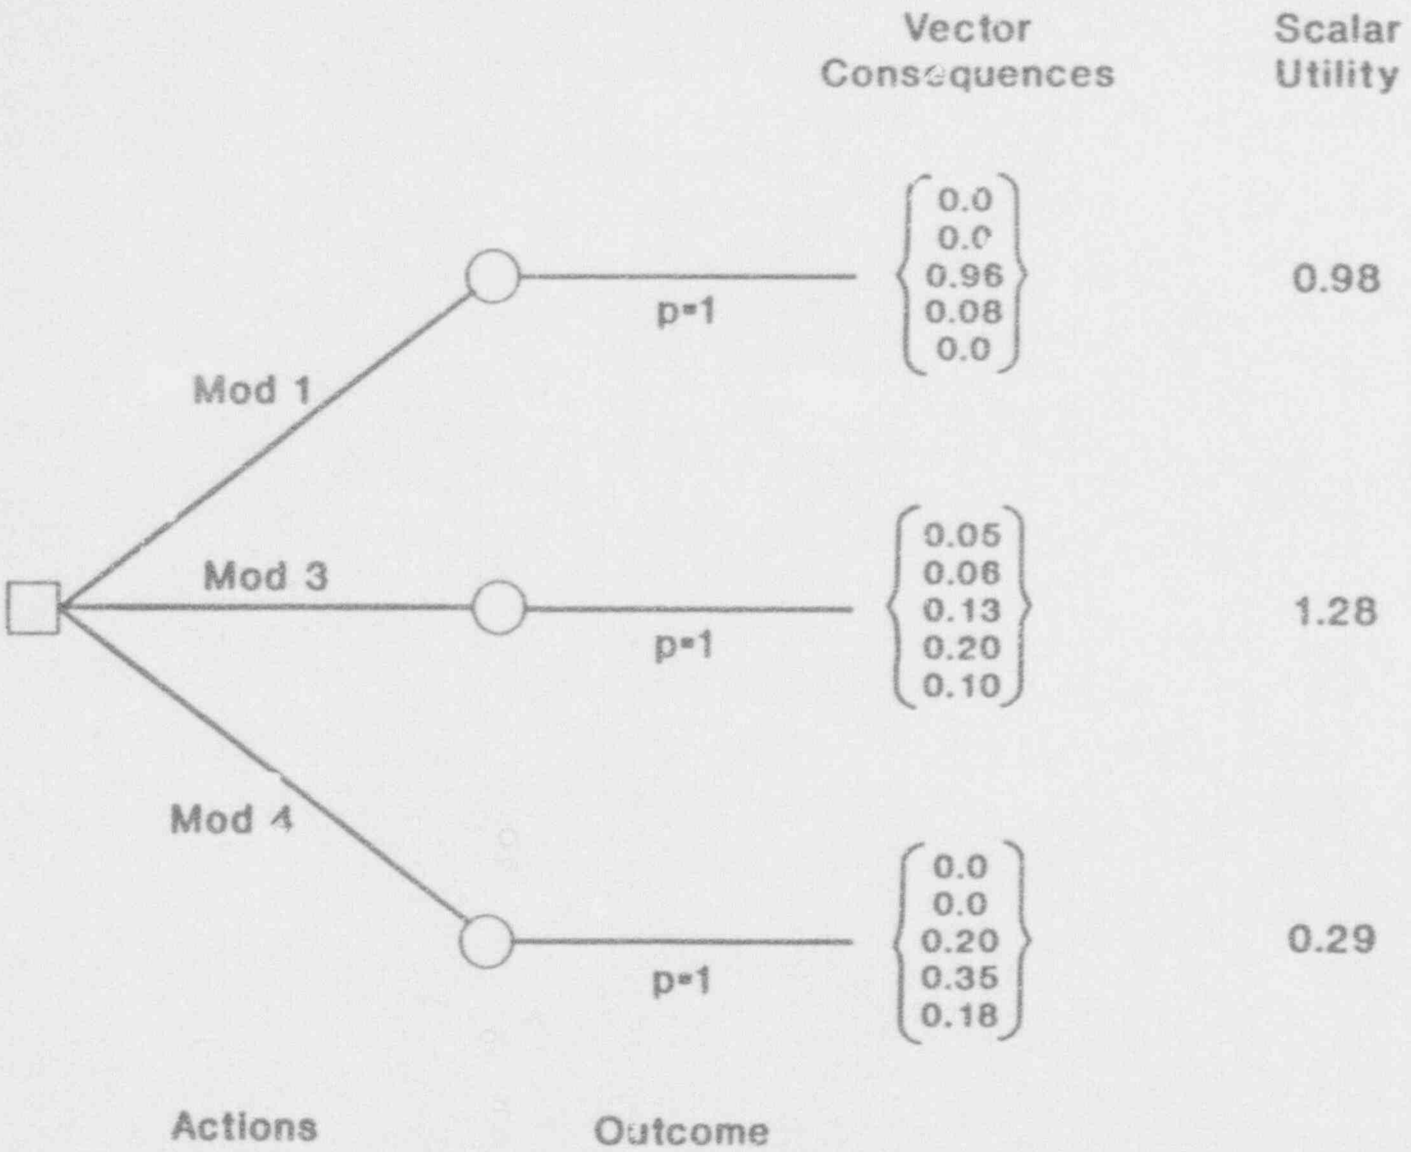

Figure 23 Decision Tree for Example in Section 3.4.1

3.4.2 Application to the B & W Plant

In this example, we will illustrate the explicit inclusion of modeling uncertainty and its impact on the decision-making process given the data for the B & W plant presented in Section 1. Note that modeling uncertainty is differentiated from random variability by virtue of the fact that random uncertainty is generally defined to be that due to the inherent randomness in the nature of a process affecting a result. Random uncertainty, for example, would be that variability in the failure level of a component which is observed in multiple tests to failure. No matter how many tests are performed, there will always be variability associated with the onset of failure. Thus random variability is essentially irreducible. By contrast, modeling uncertainty (or variability) is generally taken to be that which is due to the inability of our codes to exactly model the phenomenon being studied or the fact that, in accumulating a failure data base, one has had to combine the failure data for a number of components which are similar but not identical in design, but the conglomeration of the data was required to obtain a statistically meaningful sample. In principle, given sufficient time and money, the modeling uncertainty could be substantially reduced or eliminated. In practice, however, there is a limit to the amount of additional testing, code development, or analysis that can be performed and hence we must always make decisions in the face of (often substantial) modeling uncertainties. In fact, at the B & W plant studied in the GI-57 program, one significant modeling uncertainty issue arose.

This issue had to do with the ability of the fire growth and damage code COMPBRN to model damage to essential cables in a room whose configuration is much more complicated than the single compartment geometry for which COMPBRN is designed. Depending on the assumptions that are made in modeling this multicompartment situation, significant differences in the calculated person-REM averted result. The geometry of this situation is shown in Figure 24 where potential fire sources associated with cabinets in fire zone 95-0 were being considered. In the hall outside this fire zone there is sufficient safety-related cabling such that, if the fire spread and damaged this cabling, a core damage event would result. In Table 13 we have repeated the person-REM averted values computed for the B & W plant for the various proposed modifications as calculated for the fivefold consequence vector of internal events, fire events, and three levels of seismic events. This table is essentially that presented as Table 3 in Section 1. However an additional column denoted as Mod 5* is shown. This column presents the benefit (in terms of person-REM averted) which would accrue given that Modification 5 has been put in place, and also given that pessimistic assumptions as to the growth and spread of the seismically induced fire were made in performing the COMPBRN fire growth calculations. Thus, in this case (Mod 5*), the benefit is significantly greater than was originally shown for Mod 5 when a more optimistic calculation was made using the COMPBRN code.

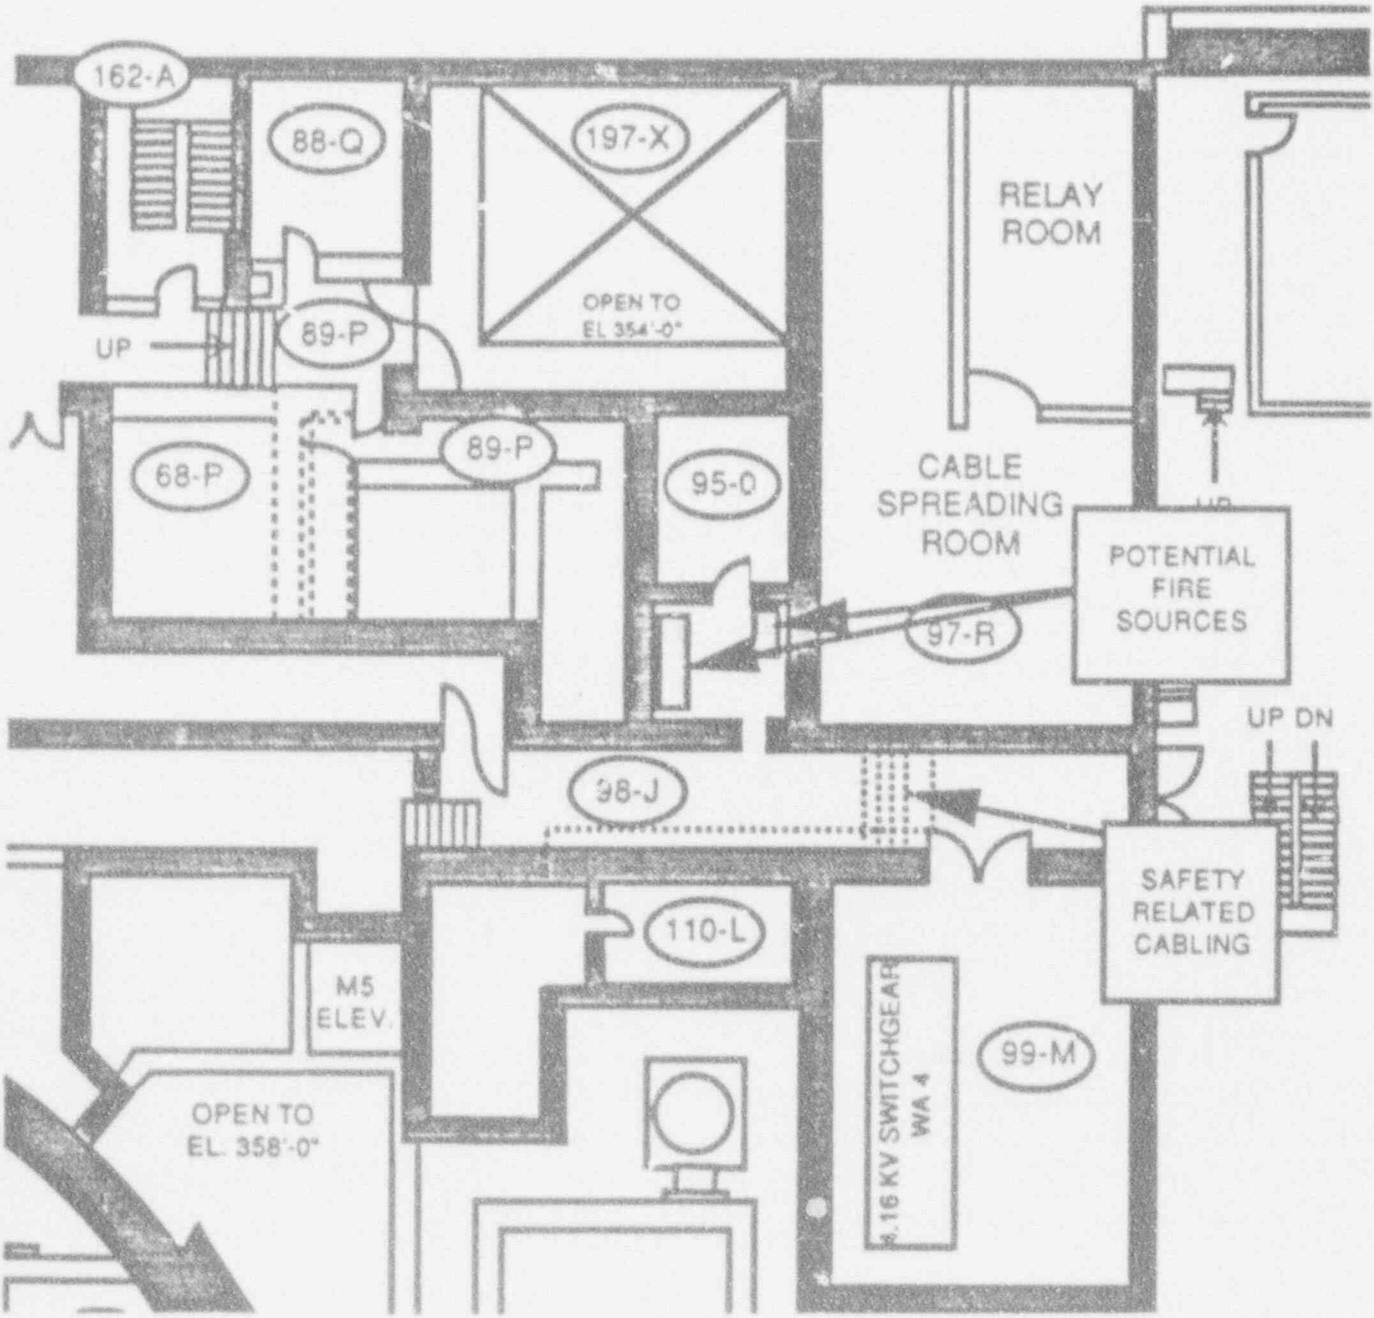

Figure 24 Room Geometry for Example in Section 3.4.2

Table 13

B & W Plant - Benefits and Costs With Uncertainty

	MOD 1 & 10	MOD 3	MOD 5 & 11	MOD 5*	MOD 7
INTERNAL	0	0.9	0	0	0
FIRE	0	0.1	0	0	0
SEIS 1	3.1	0	6.8	54.8	6.4
SEIS 2	1.7	0	53.2	315.0	52.1
SEIS 3	0	0	17.0	68.5	16.9
TOTAL	4.8 M-R	1.0 M-R	77.0 M-R	438.3 M-R	75.5 M-R

The decision tree associated with the two actions Mod 5 and Mod 7 (which are the only two viable modifications identified for the B & W plant) is shown in Figure 25. In this figure it can be seen that there is no uncertainty associated with the consequence of Mod 7. However, there are two uncertain outcomes associated with Mod 5. Judgmentally, one must assign a relative probability to the two uncertain consequences which reflects the decision maker's (or the fire phenomenology analyst's) relative confidence in the two sets of code calculations and the resulting consequences. For the sake of this example, it was assumed that the most optimistic COMPBRN code calculation had a probability of 0.67 whereas the most pessimistic calculation was associated with a probability of 0.33 as shown in this figure. Also shown in this figure are the consequence vectors associated with the outcomes of Mod 5 and Mod 7. Using the example additive multi-attribute utility function developed earlier, each consequence vector is mapped into a scalar utility and the mean of the normalized scalar utility is then computed using the probabilities of the uncertain outcomes. As shown in this figure, the expected value of the normalized utility of Mod 5 has a numerical value of 0.137, whereas the expected value of the normalized utility of Mod 7 has a value 0.124.

Thus, it can be seen that, based on a cost/benefit ratio analysis, Modification 7 would be selected over Modification 5. However, when the critical nature of the geometry of the vital area is recognized and the uncertainties in the code modeling of the damage to the critical cables is explicitly included in a multivariate utility analysis, then the preferred choice is now Modification 5. Physically, the reason for this change in the order of ranking of the modifications is the fact that, when one includes the more pessimistic COMPBRN code calculation, a greater benefit results from Modification 5 due to the fact that - with finite probability - it is possible that the more pessimistic COMPBRN code calculation is indeed the correct analysis. Thus, Modification 5 is chosen because it protects against an unlikely (but still possible) accident.

This fairly general analysis indicates the great potential usefulness of a formal decision-making methodology in the face of significant and unquantifiable uncertainty, which is often the situation faced by NRC regulators. We often work with uncertain tools and uncertain assumptions and the impact of the modeling uncertainties can often outweigh the impact of the propagation of random uncertainties in the calculation process.

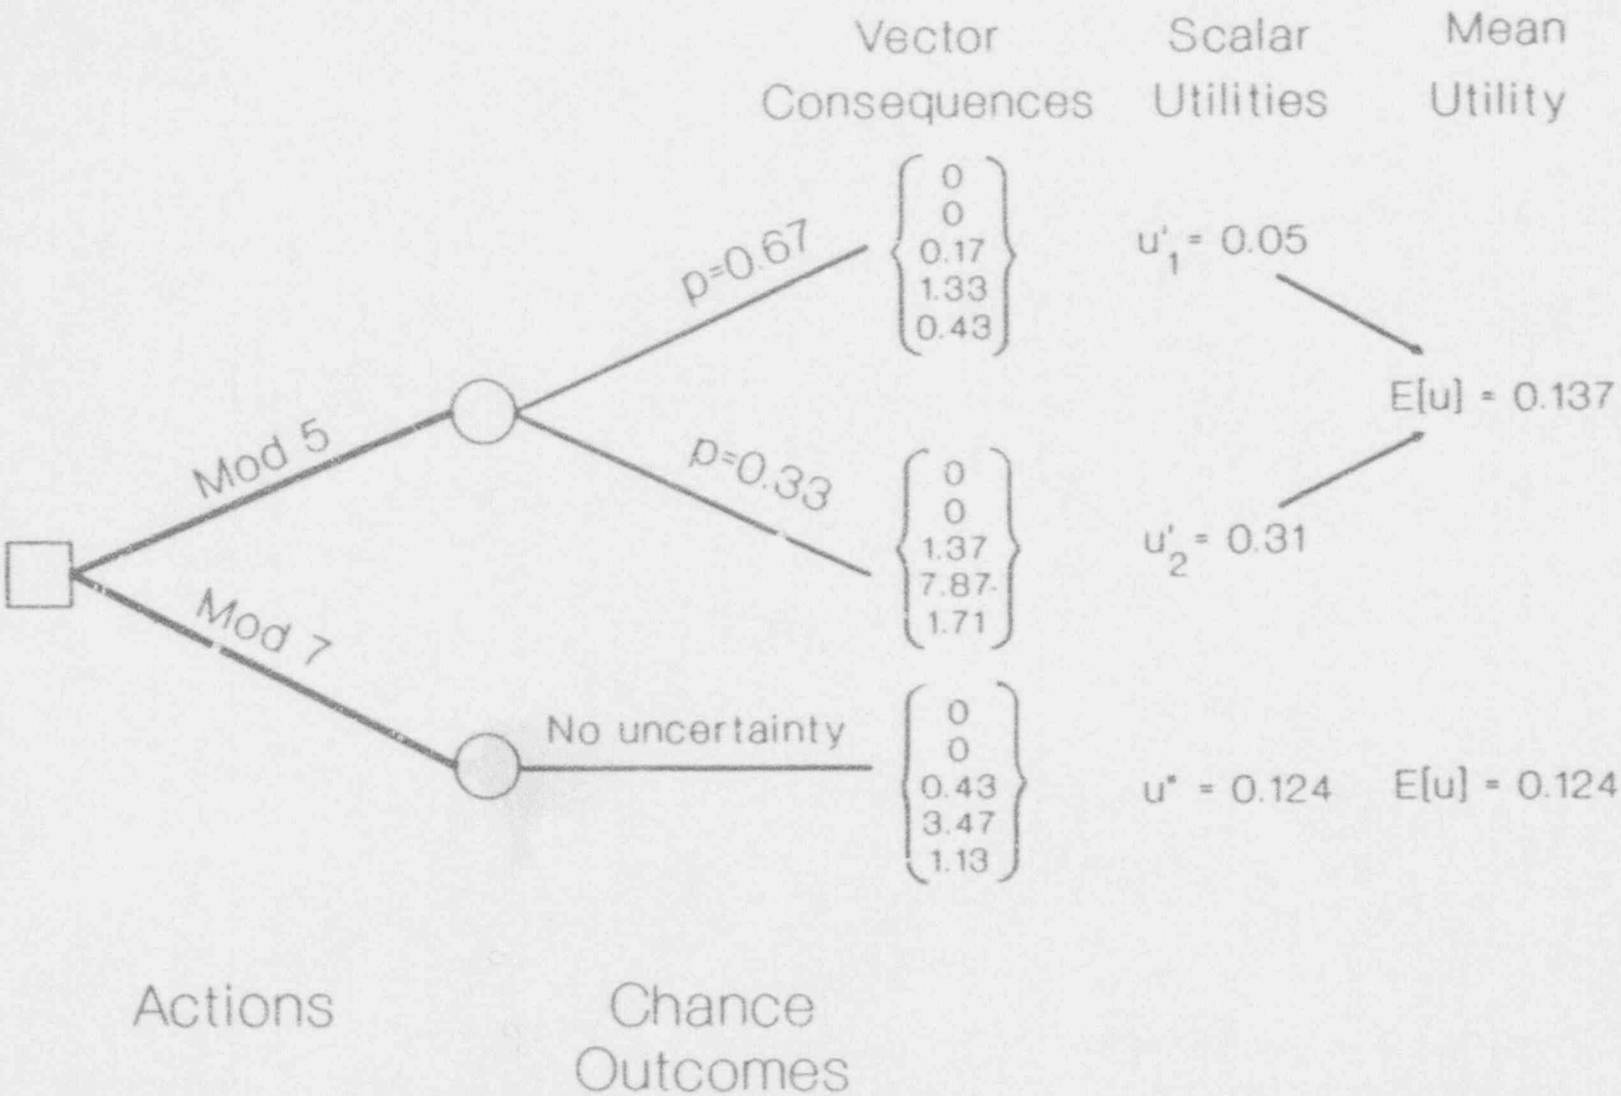

Figure 25 Decision Tree for Example in Section 3.4.2

4.0 SUMMARY AND RECOMMENDATIONS

Based on the investigations presented in this report and the examples developed using the actual cost/benefit data developed in the GI-57 project, it is clear that a formal decision-making methodology has great potential for assisting the US NRC in making retrofit (or front-fit) decisions in resolving the generic issues. It offers a number of advantages over a simple cost benefit ratio analysis, some of which include

a) It is able to clarify implicit assumptions made in the decision-making process. In particular, once the decision makers have established their preferences and value tradeoffs relative to the consequence vector attributes they wish to consider, then these preferences and tradeoffs may be taken as fixed and applied to <u>all</u> generic issues. That is, the development of utility functions and value tradeoffs are established <u>without reference</u> to the particular decision being made.

b) The formal decision-making process provides a structure which can be documented, prescribed and referenced, and if followed by different decision makers within the NRC (who utilize the same utility functions and value tradeoffs initially established) would provide unity and consistency between decisions made by different individuals, branches, or offices.

c) The decision tree formalism associating uncertain consequences to each retrofit action allows for the explicit inclusion of modeling uncertainty in the decision-making process. This modeling uncertainty is due to limitations in codes, tools, and assumptions, and these modeling uncertainties often overwhelm the variability in consequences computed when only random variabilities are propagated through the analysis.

   An additional advantage is that the assignment of relative probabilities between the various modeling uncertainty options can often be made by a combination of judgment and sensitivity studies performed by the analysts (rather than the decision makers) and this assignment of split fractions can be made referencable and transparent. In addition, it allows the inclusion of "worst case" scenarios in the decision-making process while allowing for a weighting of their likelihoods so that the final decision is not driven by the magnitude of the "worst case" scenarios hypothesized.

It is likely, then, that such a formalism has a place in the NRC decision-making process. Of course, in the examples presented in this

report, the author has had to make assumptions as to possible utility functions and value tradeoffs associated with the specific sets of consequence attributes which have been the focus of this report. The next step in the evaluation of formal decision-making methods to NRC decision-making would be to convene a panel of decision makers (perhaps personnel from the Division of Safety Issues Resolution) and go through a formal querying process to identify which attributes are of significance to this set of decision makers and what utility functions and value tradeoffs they would select. This process is not trivial and requires careful planning and the exercise could be expected to take up to three calendar days for all panel members involved. Much of the time involved would be spent in establishing consistency of global values, utilities and value tradeoffs, and assuring that the decision makers understand the relative implications of the preferences they express. To assure the latter, it would be beneficial to develop a small set of standardized problems (somewhat along the lines of the examples presented in this report) to demonstrate the implications of the various preferences expressed. Given the results of this querying process, the examples in this report (which use data developed as part of the technical resolution of GI-57) would be reevaluated and the consequences reported and made available for study for all those involved.

In addition, a number of other multi-attribute consequence vectors should be studied for possible use. For example, one could utilize the benefit as subdivided into contributions due to different containment failure modes. This choice would allow the decision maker to express his preferences as to the mode of release he wishes to impact with his retrofits. For example, he could weight his preferences so as to avoid containment failure modes resulting in large releases high in the containment rather than containment releases through the basemat or through containment isolation failure which are less "undesirable" to the general public.

Finally, due to the limited nature of the preliminary investigation presented here, we were unable to pursue the consideration of the value of additional data or analysis as applied to a decision. However, the use of Bayesian-updating based on additional data or analysis (and its associated costs) can be used with only a small extension of the work presented here. Examples of the application of Bayesian-updating (in the single attribute situation) to decision-making are given in Reference 2. The usefulness of such an updating is that it can often show that, given the uncertainties involved in the decision the development of any reasonable amount of additional data or analysis may not affect the final decision, and hence there is no value in pursuing this added data. Hence this approach can often clarify the question of whether or not to pursue additional data.

5.0 REFERENCES

1. Ericson, D.M., et al., Shutdown Decay Heat Removal Analysis: Plant Case Studies and Special Issues, NUREG/CR-5230, April, 1989.

2. Benjamin, Jack R., and C. Allin Cornell, Probability, Statistics, and Decision for Civil Engineers, McGraw-Hill, Inc., 1970.

3. U.S. Nuclear Regulatory Commission, Severe Accident Risks: An Assessment for Five U. S. Nuclear Power Plants, NUREG-1150, Vols. 1 and 2, June, 1989.

4. Lambright, J. A., et al., Risk Evaluation for a Westinghouse Pressurized Water Reactor, Effects of Fire Protection Systems Actuation on Safety-Related Equipment (Evaluation of Generic Issue 57), NUREG/CR-5789, SAND91-1534, December, 1992.

5. Lambright, J. A., et al., Risk Evaluation for a General Electric Boiling Water Reactor, Effects of Fire Protection System Actuation on Safety-Related Equipment (Evaluation of Generic Issu. 57), NUREG/CR-5791, SAND91-1536, December, 1992.

6. Lambright J. A., et al , Risk Evaluation i.. a Babcock & Wilcox Pressurized Water Reactor, Effects of Fire Protection System Actuation on Safety-Related Equipment (Evaluation of Generic Issue 57), NUREG/CR-5790, SAND91-1535, September, 1992.

7. Zionts, S. and Wallenius, J., "An interactive Programming Method for Solving the Multiple Criteria Problem", Management Science, Vol. 22, No. 6 February 1976.

8. von Neumann, J. and Morgenstern, O., Theory of Games and Economic Behavior, Princeton University Press, 1947.

9. Luce, R. Duncan, and Howard Raiffa, Games and Decisions, Introduction and Critical Survey, Dover Publications, Inc., New York, 1957.

10. Keeney, Ralph L., and Howard Raiffa, Decisions with Multiple Objectives: Preferences and Value Tradeoffs, John Wiley & Sons, New York, 1976.

11. Chernoff, Herman, and Lincoln E. Moses, Elementary Decision Theory, Dover Publications, Inc., New York, 1959.

12. Ang, Alfredo H-S., and Wilson H. Tang, Probability Concepts in Engineering Planning and Design: Basic Principles, Vol. I, John Wylie & Sons, New York, 1975.

13. Merkhofer, Miley W., Decision Science and Social Risk Management, D. Reidel Publishing Company, Boston, 1987.

14. Schlaifer, Robert, Analysis of Decisions Under Uncertainty, McGraw-Hill, New York, 1969.

15. White, D. J., Decision Methodology: A Formalization of the O R Process, John Wiley & Sons, New York, 1975.

Bibliography

Daellenbach, Hans G., and Earl J. Bell, User's Guide to Linear Programming, Prentice-Hall, Inc., Englewood Cliffs, N. J., 1970.

Hahn, Gerald J., and Samuel S. Shapiro, Statistical Models in Engineering, John Wiley & Sons, Inc., New York, 1967.

Levin, Richard I., and Charles A. Kirkpatrick, Quantitative Approaches to Management, McGraw-Hill, New York, 1965.

Mittra, Sitansu S., Decision Support Systems, Tools and Techniques, John Wiley & Sons, New York, 1986.

Rudwick, Bernard H., Systems Analysis for Effective Planning: Principles and Cases, John Wiley & Sons, Inc., 1969.

Schlaifer, Robert, Computer Programs for Elementary Decision Analysis, Graduate School of Business Administration, Harvard University, Boston, 1971.

DISTRIBUTION:

U. S. Nuclear Regulatory Commission
Attn: Roy Woods
Office of Nuclear Regulatory Research
Washington, DC 20555
(4 Copies)

Brookhaven National Laboratories
Attn: John Boccio
Bld. 130
Upton, NY 11793

Professional Loss Control
Attn: Kenneth Dungan
Michael Mowrer
PO Box 446
Oak Ridge, TN 37830
(2 Copies)

Electric Power Research Institute
Attn: Jean-Pierre Sursock
Nuclear Power Division
3412 Hillview Ave.
Palo Alto, CA 94303

American Nuclear Insurers
Attn: D. Sherman, Library
Exchange Building, Suite 245
270 Farmington Ave.
Farmington, CT 06032

Underwriters Laboratories
Attn: Leon Przybyla
333 Pfingston Rd.
Northbrook, IL 60062

Patton Fire Suppression Systems
Attn: Richard Patton
531(Roseville Rd., Suite P
Nor Highlands, CA 95660

Factory Mutual Research Corporation
Attn: Jeff Newman
1151 Boston-Providence Hwy.
Norwood, MA 02062

Grinnell Fire Protection Co.
Attn: Joe Priest
10 Dorrance St.
Providence, RI 02903

Edison Electric Institute
Attn: Jim Evans
111 19th St., NW
Washington, DC 20036-3691

EQE Engineering
Attn: Steven P. Harris
San Francisco Regional Office
595 Market Street, 18th Floor
San Francisco, CA 94105

Ogden Environmental Resources
Attn: Doug Brousseau
7301A Indian School Rd. N. E.
Albuquerque, NM 87110

Virginia Power
Attn: L. T. Warnick, P. E.
5000 Dominion Boulevard
Glen Allen, VA 23060

Science & Engineering Associates
6100 Uptown Blvd. N. E.
Albuquerque, NM 87110
(4 Copies)

Philadelphia Electric Company
Attn: Gregory A. Krueger, P. E.
Chesterbrook Blvd., 62B-3
Wayne, PA 19087-5691

Arkansas Power and Light
Attn: Ron Rospoli
1312 North Nashville St.
Russelville, AR 72801

Internal SNL:
6400 N. R. Ortiz
6405 D. A. Dahlgren
6411 D. D. Carlson
6412 A. L. Camp
6412 S. L. Daniel
6413 F. T. Harper
6449 M. P. Bohn
6449 J. A. Lambright (10)
6449 E. W. Klamerus
7141 Technical Library (5)
7151 Technical Publications
8523-2 Central Technical Files

NRC FORM 335 (2-89) NRCM 1102, 3201, 3202

U.S. NUCLEAR REGULATORY COMMISSION

BIBLIOGRAPHIC DATA SHEET

(See instructions on the reverse)

1. REPORT NUMBER (Assigned by NRC. Add Vol., Supp., Rev., and Addendum Numbers, if any.)

NUREG/CR-5906
SAND92-1547

2. TITLE AND SUBTITLE

Decision Making Under Uncertainty: An Investigation Into the Application of Formal Decision-Making Methods to Safety Issue Decisions

3. DATE REPORT PUBLISHED

MONTH	YEAR
December	1992

4. FIN OR GRANT NUMBER

L1334

5. AUTHOR(S)

Michael P. Bohn

6. TYPE OF REPORT

Technical

7. PERIOD COVERED *(Inclusive Dates)*

8. PERFORMING ORGANIZATION – NAME AND ADDRESS *(If NRC, provide Division, Office or Region, U.S. Nuclear Regulatory Commission, and mailing address; if contractor, provide name and mailing address.)*

Sandia National Laboratories
Albuquerque, NM 87185

9. SPONSORING ORGANIZATION – NAME AND ADDRESS *(If NRC, type "Same as above"; if contractor, provide NRC Division, Office or Region, U.S. Nuclear Regulatory Commission, and mailing address.)*

Division of Safety Issues Resolution
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555

10. SUPPLEMENTARY NOTES

11. ABSTRACT *(200 words or less)*

As part of the NRC-sponsored program to study the implications of Generic Issue 57, "Effects of Fire Protection System Actuation on Safety-Related Equipment", a subtask was performed to evaluate the applicability of formal decision analysis methods to generic issues cost/benefit type decisions and to apply these methods to GI-57 results. In this report, the numerical results obtained from the analysis of three plants in the GI-57 technical resolution program were studied. For each plant, these results included a calculation of the man-REM averted due to various accident scenarios and various modifications proposed to mitigate the accident scenarios identified. These results were recomputed to break out the benefit in terms of contributions due to random event scenarios, fire event scenarios, and seismic event scenarios.

Given this data, formal decision methodologies involving decision trees, value functions, and utility functions were applied to this basic data. Examples are given (based on assumed NRC decision maker preferences) in which the decision between several retrofits is changed from that resulting from a simple cost/benefit ratio criterion. The examples given in this report demonstrate that the application of formal decision making methodology to the NRC decision-making process can potentially make such decisions more consistent, transparent and referenceable as well as allowing for explicit inclusion of the decision-maker's preferences. Recommendations for further work in this area are included.

12. KEY WORDS/DESCRIPTORS *(List words or phrases that will assist researchers in locating the report.)*

Decision Analysis
Cost/Benefit Decisions
Risk Analysis
PRA Applications

13. AVAILABILITY STATEMENT

Unlimited

14. SECURITY CLASSIFICATION

(This Page) Unclassified

(This Report) Unclassified

15. NUMBER OF PAGES

16. PRICE

NRC FORM 335 (2-89)

Printed on recycled paper

Federal Recycling Program

UNITED STATES
NUCLEAR REGULATORY COMMISSION
WASHINGTON, D.C. 20555-0001

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE, $300

SPECIAL FOURTH-CLASS RATE
POSTAGE AND FEES PAID
USNRC
PERMIT NO. G-67