NUREG/CR-6833

Formal Methods of Decision Analysis Applied to Prioritization of Research and Other Topics

Information Systems Laboratories, Inc.

U.S. Nuclear Regulatory Commission
Office of Nuclear Regulatory Research
Washington, DC 20555-0001

AVAILABILITY OF REFERENCE MATERIALS IN NRC PUBLICATIONS

NRC Reference Material

As of November 1999, you may electronically access NUREG-series publications and other NRC records at NRC's Public Electronic Reading Room at http://www.nrc.gov/reading-rm.html. Publicly released records include, to name a few, NUREG-series publications; *Federal Register* notices; applicant, licensee, and vendor documents and correspondence; NRC correspondence and internal memoranda; bulletins and information notices; inspection and investigative reports; licensee event reports; and Commission papers and their attachments.

NRC publications in the NUREG series, NRC regulations, and *Title 10, Energy*, in the Code of *Federal Regulations* may also be purchased from one of these two sources.

1. The Superintendent of Documents
   U.S. Government Printing Office
   Mail Stop SSOP
   Washington, DC 20402–0001
   Internet: bookstore.gpo.gov
   Telephone: 202-512-1800
   Fax: 202-512-2250
2. The National Technical Information Service
   Springfield, VA 22161–0002
   www.ntis.gov
   1–800–553–6847 or, locally, 703–605–6000

A single copy of each NRC draft report for comment is available free, to the extent of supply, upon written request as follows:

Address: Office of the Chief Information Officer,
Reproduction and Distribution
Services Section
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001
E-mail: DISTRIBUTION@nrc.gov
Facsimile: 301–415–2289

Some publications in the NUREG series that are posted at NRC's Web site address http://www.nrc.gov/reading-rm/doc-collections/nuregs are updated periodically and may differ from the last printed version. Although references to material found on a Web site bear the date the material was accessed, the material available on the date cited may subsequently be removed from the site.

Non-NRC Reference Material

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. Such documents as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings may be purchased from their sponsoring organization.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at—

The NRC Technical Library
Two White Flint North
11545 Rockville Pike
Rockville, MD 20852–2738

These standards are available in the library for reference 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—

American National Standards Institute
11 West 42nd Street
New York, NY 10036–8002
www.ansi.org
212–642–4900

Legally binding regulatory requirements are stated only in laws; NRC regulations; licenses, including technical specifications; or orders, not in NUREG-series publications. The views expressed in contractor-prepared publications in this series are not necessarily those of the NRC.

The NUREG series comprises (1) technical and administrative reports and books prepared by the staff (NUREG–XXXX) or agency contractors (NUREG/CR–XXXX), (2) proceedings of conferences (NUREG/CP–XXXX), (3) reports resulting from international agreements (NUREG/IA–XXXX), (4) brochures (NUREG/BR–XXXX), and (5) compilations of legal decisions and orders of the Commission and Atomic and Safety Licensing Boards and of Directors' decisions under Section 2.206 of NRC's regulations (NUREG–0750).

NUREG/CR-6833

Formal Methods of Decision Analysis Applied to Prioritization of Research and Other Topics

Manuscript Completed: August 2003
Date Published: September 2003

Prepared by
R.W. Youngblood, W.C. Arcieri, S.A. Faridi
N.P. Kadambi (NRC)

Information Systems Laboratories, Inc.
11140 Rockville Pike
Rockville, MD 20852

N.P. Kadambi, NRC Project Manager

Prepared for
Division of Systems Analysis and Regulatory Effectiveness
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001
NRC Job Code Y6406

ABSTRACT

This report discusses the application of methods of formal decision analysis to prioritization of research to be carried out in support of licensing of advanced reactor designs. Formal decision methods are useful in this area for two reasons. (1) Prioritization is a special case of decision-making. (2) Prioritization of safety research is closely related to safety decision-making, and formal analysis of safety decisions points the way more clearly to specific research tasks needed to support safety decisions. The report is presented in three main parts. Part I provides an overview of prioritization. Two main themes emerge from this overview: (1) the effect of uncertainty on decisions, and (2) the need to clarify decision objectives and carefully formulate decision alternatives. The first theme is taken up in Part II, and the latter in Part III. Specific topics considered include the development and use of an objectives hierarchy, kinds of performance measures, and the value of information to reduce uncertainty, including a discussion of hypothesis testing and "Receiver Operating Characteristic Curves." Many agency decisions could arguably benefit from application of individual tools discussed here, although application of the full suite of formal decision-analysis tools may not be warranted in all cases.

TABLE OF CONTENTS

List of Figures

List of Tables

EXECUTIVE SUMMARY

The primary subject of this report is the application of formal decision methods to prioritization of research to support review and licensing of advanced reactor designs. Because prioritization is a special case of general decision analysis, many of the tools and ideas that apply to prioritization apply much more broadly to agency decision-making, and selected instances of this are discussed in the report. The literature of this subject is vast; the present report aims to bring selected tools and ideas to the attention of potential users, and direct them to source material with the kind of detail needed to support real applications.

The report is divided into three parts: Tools for Research Prioritization (Part I), Treatment of Uncertainty (Part II), and Development and Application of the Objectives Hierarchy (Part III).

Part I, Tools for Research Prioritization, gives a brief overview of formal approaches to decision-making, and shows how prioritization is a special application of those approaches. One of the factors that complicates decision-making in many cases, and drives the need for formal approaches, is uncertainty regarding the probabilities and consequences of various outcomes associated with decision alternatives. This is certainly true for safety decisions that the agency must make; reduction of significant technical uncertainties is a key driver for research prioritization. This topic is discussed in the context of the "value of information," a concept that relates the potential worth of a research program to the effect of the subject uncertainty on agency decisions. Next, a general approach to prioritization is discussed. This discussion highlights the need for a considered formulation of the fundamental and means objectives that the decision-maker needs to address. The report does not execute this development on behalf of the agency, but mentions examples of other developments, and follows up with examples of performance measures that might well emerge from an agency development.

Two key themes emerging from Part I are (1) the treatment of uncertainty, and (2) the formulation and application of the objectives hierarchy. The concepts and ideas associated with each of these topics make them worth considering not only in actual decisions, but also in many supporting activities that are not necessarily explicitly formulated as "decision-making" even though they entail setting priorities and making choices. Accordingly, each is treated in a separate part of the report.

Part II, Treatment of Uncertainty, begins by presenting some essential results from the basic theory of hypothesis testing. The problem nominally treated by this theory is that of deciding which of two hypotheses is correct, given evidence that suggests, but does not prove conclusively, which is correct. Although called "hypothesis testing," this basic problem corresponds in principle to a broad spectrum of decision situations, because it shows how uncertainty limits the expected utility of a decision (i.e., creates some potential for adverse consequences resulting from the decision). This discussion is followed up with a brief overview of the subject of "Receiver Operating Characteristic" (ROC) curves. ROCs are very widely studied because of the insight that they afford into a broad spectrum of decision situations. Because of the perspective that they provide on uncertainty, they are of potential interest in a research organization responsible for reducing technical uncertainty. These formalisms support Part I's discussion of the value of information, and shed additional light on the value of reducing uncertainty regarding (for example) phenomenological issues in the context of specific decisions.

The standard treatment of hypothesis testing, and the standard treatment of uncertainty with ROCs, show that in order to optimize the expected consequences of this type of decision, one needs to apply

- prior probabilities of the hypotheses,
- information regarding the costs and benefits of correct and incorrect conclusions based on the evidence, and
- conditional probabilities linking the evidence to the hypotheses being tested.

Nevertheless, some decisions are made with less than complete attention to these factors.

Some do not accept using this "expected consequences" formulation, arguing instead that when there are substantial uncertainties regarding potentially large consequences, one ought to revert to an essentially conservative stance. This view can be approximated to some extent within an expected-consequences approach, but seems philosophically distinct. This position is discussed briefly. Finally, it is observed that even now, the *assessment* of uncertainty distributions is a much less well-developed area than the formalisms for applying uncertainty distributions once they have been obtained.

Part III, Development and Application of the Objectives Hierarchy, begins with a brief discussion of the objectives hierarchy, and illustrates it with reference to the formulation of the Cornerstones of Safety and associated Key Attributes in SECY 99-007, "Recommendations for Reactor Oversight Process Improvements." The objectives hierarchy is a presentation of the relationships between overarching goals, fundamental objectives, means objectives, and performance measures. It is obvious that in order to make decisions, one needs to have one's priorities straight, but certain subsequent points are initially less obvious. The kinds of performance measures that one selects can have an immense impact on how one's priorities are ultimately implemented; for example, "risk-informed" regulation is essentially an attempt to revert to natural measures, while prescriptive regulation is based on proxy measures. The next topic is formulation of decision alternatives. Choosing from a fixed list of alternatives, the best one can do is the best on that list; it is therefore crucial to formulate the best possible choices. It turns out that thinking carefully about the objectives hierarchy turns out to play into the formulation of better alternatives in the first place. As an example of the applicability of decision tools outside of narrowly-construed "decision-making," this topic was discussed at some length in NUREG/BR-0303, "Guidance for Performance-Based Regulation," in the context of formulating performance-based regulatory approaches. Another topic that arises is that of performance allocation (deciding what level of safety performance is needed in each area, in order that safety objectives be satisfied). This is a job that needs to be done not only in the formulation of safety cases in general, but also in many more specific regulatory tasks. Allocation is best performed on the basis of a considered development of the objectives hierarchy.

In Part I's overview of prioritization, the role of the objectives hierarchy in general prioritization was presented; in Part III, in order to support *research* prioritization, a summary of recent interesting work on long-term priorities is presented. This work addresses the difficult problem of dealing explicitly with long-term needs in the context of short-term program decision-making.

Finally, the Analytic Hierarchy Process is described briefly with respect to its use as a tool for ranking a list of items. In some applications, the AHP accomplishes this task while leaving much of the objectives hierarchy implicit. In other applications, it supports development of an objectives hierarchy and allocation of resources or performance over its elements.

FOREWORD

This document is intended to provide a starting point for guiding NRC staff toward the use of formal decision methods by extracting relevant methodologies from the literature, and explaining them without overwhelming the reader with technical detail. The specific focus on prioritization of research originated from the initial definition of the effort as helping with decision-making on advanced reactor infrastructure development (SECY-03-0059, "Advanced Reactor Research Program"). The existing infrastructure uses the PIRT (Phenomena Identification and Ranking Tables) process that has been used successfully for some time, and will be used to prioritize activities through the Planning, Budget, and Program Management (PBPM) process. As described in SECY-03-0059, the prioritization method is needed to help allocate available resources, and will be applied in conjunction with the PBPM process. The material presented in this document can be viewed as complementary to the PIRT process so as to accomplish the PBPM process more effectively, and thereby achieve the NRC's strategic performance goals.

NRC's strategic performance goals represent a real-life example of multiple attributes that need to be fulfilled simultaneously as part of organizational decision-making. Formal decision methods can help significantly with certain key aspects of such goals as conducting the regulatory process effectively, efficiently, and at the same time, openly with all stakeholders. The reason for this is the unique ways in which formal methods apply combinations of qualitative and quantitative information. The credibility of such a decision process depends on transparency in defining the process elements, repeatability of the process elements consistently across applications, and reproducibility of the results from implementing the process. These are shared characteristics with the outcome of other research conducted recently at NRC on performance-based regulation. The two documents that report on this research are NUREG/BR-0303, "Guidance for Performance-Based Regulation" and NUREG/CR-5392, "Elements of an Approach to Performance-Based Regulatory Oversight." What is common to these efforts is the careful thought given to structuring the objectives of the decision making process and iteratively obtaining stakeholder feedback along the way. Also, the systematic approach to deal with uncertainties that is possible with formal methods can improve the quality of decision-making. In terms of complementing the PIRT process, the formal methods approach will likely help answer an important question posed in SECY-03-0059: How conservative will a decision have to be if the information is not obtained or uncertainties are large?

This document has been reviewed by staff from other NRC Headquarters program offices. It has become clear from the review comments that offering this document as a resource to staff without additional guidance and elucidation using specific examples would not fulfill the objectives of this study, which included having the document be user friendly. It was recognized from the start that the abstract concepts in this document require specific context to promote understanding. The NRC will be considering a range of such applications as the next step.

Farouk Eltawila, Director
Division of Systems Analysis and Regulatory Effectiveness
Office of Nuclear Regulatory Research

ACKNOWLEDGMENTS

This report has benefited greatly from conversations with many colleagues and from insightful comments offered by NRC staff, including Dennis Damon, Susan Cooper, Makuteswara Srinivasan, and others. We regret that the limited objectives of the present effort did not allow us to implement all of the comments received, including some that we agree with.

ABBREVIATIONS

ACRS	Advisory Committee on Reactor Safeguards
AFWS	auxiliary feedwater system
AHP	analytic hierarchy process
ALARP	as low as reasonably practicable
BWR	boiling-water reactor
CDF	core damage frequency
CSAU	Code Scaling, Applicability, and Uncertainty Evaluation
DOE	U.S. Department of Energy
ECCS	emergency core cooling system
GNP	gross national product
LB	licensing basis
LERF	large early release frequency
LWR	light-water reactor
MAUT	multiattribute utility theory
MEU	maximum expected utility
NAS	DOE's Office of Environmental Management
NRC	U.S. Nuclear Regulatory Commission
PCT	peak cladding temperature
PIRT	Phenomena Identification and Ranking Table
PRA	probabilistic risk assessment
PrOACT	problem, objectives, alternatives, consequences, and tradeoffs
PWR	pressurized-water reactor
RES	NRC's Office of Nuclear Regulatory Research
RG	regulatory guide
ROC	receiver operating characteristic
SSCs	structures, systems, and components
SFF	suitability, feasibility & flexibility
VOI	value of (perfect) information

Part I: Overview of Prioritization

I.1 INTRODUCTION

I.1.1 Purpose

The purpose of this report is to support application of formal decision methods to prioritization of research related to advanced reactors. However, prioritization is a special case of decision-making, and the tools applicable to prioritization are applicable to other agency decisions as well. More fundamentally, safety research is ultimately aimed at safety decisions, and considering formal approaches to safety decisions also helps prioritization.

Therefore, two related kinds of decisions are addressed in this report:

1. safety decisions: decisions related to evaluation of designs, operations, ...;
2. programmatic decisions regarding the content of research programs aimed at supporting safety decisions.

One kind of decision is made regarding whether safety performance is satisfactory: whether a given system state is "success" or "failure," whether a given design has sufficient reliability, etc. This kind of decision is referred to in this report as a "safety" decision. Other kinds of decisions are also made, one kind being "programmatic." An example of a programmatic decision is one made regarding the priority of a given research activity. This kind of decision is referred to in this report as a "programmatic" decision. These two kinds of decisions are closely related, because NRC programmatic decisions are ultimately driven by safety decisions. Also, some tools apply to both. It might be possible to integrate them formally into a single overarching decision model. However, that integration task has not yet been accomplished. In order to avoid confusion, it is useful to think of them separately. Each kind of decision has its own drivers, which are related but distinct.

The purpose of this report is directly related to programmatic decisions, but both kinds of decisions need to be addressed in this report because the purpose of the programmatic decisions is to support the safety decisions.

The following activities culled from the Advanced Reactor Research Plan (Ref. I-1) especially lend themselves to the use of formal decision methods in prioritizing research:

- Independently confirm safety case
- Modifications to existing regulatory framework (in response to challenges posed by new designs)
- Develop new risk-informed, performance-based regulatory decision-making tools that will support modifications to the regulatory framework
- Develop new safety limits
- Upgrade data base to assess safety margins
- Development of expertise, tools, and methods that are needed to support the Agency's mission in understanding and resolving potential safety issues

The role for RES contemplated in the plan is:

- Conduct supporting research to help establish the technical basis and acceptance criteria for the safety case
- Explore issues involving large uncertainties
- Develop independent capabilities to review applicants' submittals
- NRC will recognize and, as appropriate, coordinate with research being conducted by others (e.g., industry, international efforts).

In a sense, formal decision analysis is applicable to all of the above activities, since all of them involve making choices. Indeed, some of the above topics arguably should be addressed as formal decision problems. However, this is not to say that all, or even most, choices should be subjected to the full decision process summarized later. Certain tools of decision analysis can beneficially be applied in particular contexts on a more or less stand-alone basis. That is the focus of this report: to emphasize certain tools and show how they are applicable to activities like those listed above.

There is a very large body of literature on formal decision methods. It is not the purpose of this report to substitute for any part of that large body of literature. Rather, it is the purpose of this report to present ideas culled from the literature that seem to offer help in meeting the needs identified above, and show how those ideas might help.

I.1.2 Overview of Part I

Part I focuses on the application of specific decision analysis tools to prioritization of research to support safety decisions for advanced reactors. Many of the tools can be helpful in research prioritization, even if a complete, formal decision analysis is not mounted for every issue that comes up. Part I of this report provides an overview of decision analysis in general and prioritization in particular, and discusses many of these tools. Two general areas in decision analysis have special importance for advanced-reactor work: uncertainty analysis, including its use in determining the value of additional research, and considered development of an objectives hierarchy. These general areas receive focused treatment in Parts II and III.

Section I.2 presents an overview of formal decision analysis as it is described in texts. (This formal approach is not always followed, perhaps not even typically followed.) This turns out to be much more than modeling; it entails much up-front thought about objectives and decisions, and these considerations then drive modeling (e.g., risk modeling) in particular ways.

Section I.3 then briefly introduces the concept of value of information. Analyzing the "value of (perfect) information" (VOI) provides an upper bound on what it is potentially worth to reduce uncertainty in a given decision situation. VOI could be an extremely useful tool in establishing research priorities. However, proper assessment of it requires an honest assessment of technical uncertainty, a subject that is not always thought of as part of "decision analysis" but one that plays a crucial role. Uncertainty analysis receives a more thorough treatment in Part II.

Arguably, many research projects indirectly help in ways that go beyond their impact on current decisions, helping (for example) to identify better decision alternatives in the future. VOI is therefore not the whole story on the value of research.

Next, in Section I.4, recent work in prioritization will be discussed. This not only serves to illustrate the mechanics of simple prioritization tasks, but also serves to introduce by example the need to focus on objectives. The objectives hierarchy is discussed much more extensively in Part III, not only with regard to prioritization, but also with regard to development of performance measures. There is broad consensus that this is a necessary first step in any systematic decision process, but it turns out that this is also a useful tool in certain developments that do not correspond to complete "decision analyses." An example is the identification of performance metrics for advanced designs. One objectives hierarchy is useful for thinking directly about the overall features of a given design, while another is useful for ranking research programs.

Finally, in Section I.5, performance measures for research will be discussed. It is beyond the mandate of this report to determine the set of measures for actual application, but some insights from the literature are noted for application by those whose job it is to determine and apply such measures.

Part I concludes in Section I.6 by collecting, in one place, possible applications of decision analysis tools that have been mentioned in the report. This summary is arranged in two tracks, mirroring the two kinds of decisions distinguished in Section I.1.

I.2 OVERVIEW OF FORMAL DECISION-MAKING

I.2.1 Context of Formal Decision Methods

In general, formal decision methods offer benefits in situations having one or more of the following characteristics (Ref. I-2):

- complexity of the situation;
- significant uncertainties;
- multiple objectives / tradeoffs;
- different perspectives (e.g., multiple stakeholders).

Prioritization of advanced-reactor work arguably involves all of these elements. Significant uncertainties are present, and prioritizing research to resolve them is an obvious need. Complexity is present, in that some research tasks offer benefits in different areas, some combinations may be more mutually reinforcing than others, etc. Multiple objectives are present in that safety, cost, and timeliness are all key factors in the decision. Different perspectives operate as well, even within broadly identified groups such as "public" and "industry."

Implicit in the above discussion is the idea that the justification for decision analysis is also influenced by the significance of the stakes. In reactor regulation, one scale of significance is set by the various consequences of a reactor accident, and (for example) corresponding monetizations of those consequences. Another consideration is the value of lost production.

However, some of the stakes arguably lie beyond simple notions of "cost" or "safety." Significant delay in a licensing action can have adverse societal impacts beyond the monetized cost to that specific applicant. On the other hand, a decision made hastily on an inadequate technical basis can have adverse impacts even if adverse radiological consequences do not immediately occur as a result of that specific decision. These consequence types might be monetizable, but their significance is qualitatively clear even without that step having been taken.

Formal decision methods therefore have something to offer to prioritizing advanced-reactor research.

"Decision-making" is discussed at several levels in this report. Ultimately, the agency makes decisions about safety, using the best available tools and methods. Formal decision methods can help here; they have been shown to be helpful in a wide variety of technical fields. Initially, this may not appear to bear on prioritization, but it does. Decision theory helps to establish, for a given design, what parameters are most appropriate to consider as the basis for decision-making. This is done based on a systematic consideration of decision objectives and system properties. Part of the job of research is to reduce uncertainty in these parameters, and decision theory provides a basis for quantifying the value of uncertainty reduction for specific parameters in the context of specific safety decisions. This capability has face value in research planning. However, just establishing upper bounds on the value of given research programs does not establish which of them ought to be pursued; deciding what programs actually to pursue is a separate, higher-level decision, having different intermediate-level objectives. For example, those objectives may involve timeliness. Formal methods for prioritization can help sort out research alternatives at this level.

Several of the above comments are aimed at identifying and culling focused individual research programs, at what might be called the tactical level. More fundamentally, formal methods can help in the process of developing good high-level program alternatives at the strategic level, usefully combining tactical program elements, and recognizing the difference between decision problems and decision opportunities.

I.2.2 Steps in Decision-Making

Most modern discussions of the overall process agree in many respects, and vary only in relative emphasis of certain topics and in the details of the presentation. This section presents an overview of a relatively formal decision-making process, based on the idea (discussed below) that there is a "utility function" that relates decision alternatives to the decision-maker's values and preferences, and the point of the formal process of decision-making is to optimize the expected outcome of the decision with respect to this figure of merit. Accordingly, this process devotes significant attention to explicitly clarifying objectives and values. Many different tools can be applied within this high-level process, although not all tools fit neatly into this paradigm.

The present discussion is based on a popular overview of decision-making recently provided by Hammond, Keeney, and Raiffa (Ref. I-3) (hereafter HKR). They decompose the overall decision process into eight elements, as follows:

1. Defining and working on the correct decision problem
2. Specifying and structuring objectives
3. Formulating alternatives
4. Understanding the consequences
5. Exploring tradeoffs
6. Clarifying uncertainties
7. Contemplating risk tolerance
8. Analyzing and examining linked decisions

HKR derive the acronym PrOACT from the first five elements, **Pr**oblem, **O**bjectives, **A**lternatives, **C**onsequences, and **T**radeoffs, which are the core of the process. At a high level, this discussion is representative of most academic presentations of decision analysis, in that it aims at choosing the alternative that has the best expected utility. (The concept of "utility" is discussed later in this section.) However, HKR's presentation is developed at a popular level (has less mathematics), and deals with the utility function at a much more intuitive level than is common practice in academic texts. In keeping with the emphasis of certain earlier works of these authors, HKR's presentation also pays a lot of attention to avoiding certain traps, such as settling for obvious default alternatives when superior ones might emerge from a considered thought process, and neglecting to carefully develop the objectives that need to be addressed in the decision. Although HKR minimize the mathematics, and in particular refrain from much discussion of mathematical utility functions, the approach is formal in the sense that it goes through a series of steps, each of which is meant to be taken seriously.

The first element, framing the "problem" (the Pr of the PrOACT method), introduces the logical beginning of every decision process: a creative problem definition. It is important to formulate the problem correctly. Defining a problem addresses such issues as: 1) the impetus for the particular problem at hand and 2) the definition of the problem and the constraints of the problem and whether they are available, necessary and relevant. Once these issues are addressed, the problem is decomposed into component pieces. Advice is solicited from professionals in the pertinent fields who have faced a similar problem. The problem definition is reevaluated as the decision is formulated as a final step.

HKR emphasize that the way in which a problem is formulated has substantial impact in the framing of the decision and in the determination and evaluation of the alternatives. Not giving ample consideration to problem formulation can force participants involved in the decision making process to plunge into the other elements of decision making based on a sub-optimal formulation of the problem. Problem formulation for the prioritization of research for advanced reactors might beneficially recognize the desirability of proceeding more effectively and efficiently as compared with light-water reactor (LWR) experience. This involves explicitly taking into account the safety significance and timing of decisions. Rather than guessing at research resources needed for each concept, a more structured approach would consider the significance of topics to specific milestones and account for associated uncertainties

The second element, specifying and structuring objectives, is important because objectives form the bases of evaluating alternatives. Elsewhere in this report, this subject receives considerable attention because it is an interesting topic even outside the context of a formal decision process. In decision analysis, it informs the development of alternatives and the selection of attributes that drive the choice between alternatives being considered. In order to do this, of course, it drives various intermediate steps, including modeling and uncertainty analysis.

The third element is formulating alternatives. There is a significant body of work that recognizes the formulation of alternatives as a key to effective decision-making. Some common pitfalls experienced when devising alternatives include extensive dependence on past experience, failure to capture ideas and insights, early (premature) closure on problem definition, criticizing new ideas as they are offered, ruling out alternatives too early in the decision-making process and failure to re-consider discarded alternatives if conditions change. HKR identify some of the common mistakes made in devising alternatives, such as falling back upon a default alternative, or simply choosing among alternatives presented by other people.

The fourth element, understanding the consequences, is directly linked to the formulation of alternatives. As stated in HKR, "be sure you really understand the consequences of your alternatives before you make a choice. If you don't, you surely will afterwards, and you may not be very happy with them." Aimed at popular audiences who will not typically "model" the consequences of their decision alternatives, HKR's popular account does not address "risk analysis" at this point, but in reactor licensing, this is the stage at which risk analysis would enter.

In general discussions of decision analysis, the phrase "risk analysis" is used in a more general sense than that frequently intended by agency users, who frequently mean to imply fault tree or event tree analysis. Quite generally, the phrase "risk analysis" refers to the quantification of expected consequences in light of uncertainty. This can be complex, or it can be very simple, but if called "risk analysis," it involves uncertainty.

The fifth element of the PrOACT process is evaluating tradeoffs. This element basically involves ranking the alternatives and making tradeoffs among the various alternatives. Since many practical decisions entail consideration of diverse alternatives and diverse objectives, an approach to evaluating tradeoffs is needed. Evaluating tradeoffs can be one of the more challenging aspects of decision-making, because each objective may have its own basis of comparison. In one case, the alternatives may be compared using numerical values. In other cases, relational judgements (high, medium, low) or descriptive terms may be used. There may be a need to consider input from different stakeholders. Making tradeoffs among disparate objectives can be a significant undertaking.

In eliminating alternatives, one approach is to identify dominant alternatives using an argument such as "if alternative A is better than alternative B in one objective and no worse in the remaining objectives, then B can be eliminated."

Elimination of objectives can also be done to simplify the decision problem. The problem of making tradeoffs among objectives is not new. According to HKR (page 88-89), Benjamin Franklin outlined a method for making decision tradeoffs between two alternatives in a letter to Joseph Priestley (the discoverer of oxygen). Franklin's approach was to divide a sheet of paper into two columns, listing one "Pro" and the other column "Con." Then, over the course of three to four days, he would write short hints of the positive and negative aspects of the decision. After he organized his thoughts (hints) for or against the decision, he would estimate their weights, eliminating factors on each side that he judged to be equal in weight. After this process was completed, he would make his decision accordingly. Franklin felt that he could make better judgements using this process, which he dubbed "moral or prudential algebra." HKR suggests an approach similar to Franklin's in making tradeoffs: the "even swap method." In this method, the pros and cons of various alternatives are bartered against the other, forcing the decision-maker to think about the value of one objective against other objectives. In already simple problems, this method allows the decision-maker to eliminate objectives, further simplifying the problem.

Overall, the strategy is to eliminate a sufficient number of alternatives and objectives that a decision can be reached by inspection. HKR's popular account is aimed at being able to do this by hand; more difficult decision problems may call for developing complex models and carefully-considered utility functions, so that the final choice can be made based on a quantitative ranking of alternatives' expected utilities.

The remaining three elements, uncertainty, risk tolerance, and linked decisions, are used to clarify decisions in evolving environments. Prioritization of research needs in a research program would be an evolving environment.

I.2.3 Discussion

The above philosophy of decision-making seeks to produce a decision that "best" matches the true objectives of the decision-maker, whatever they are. Substantial energy is spent determining what these objectives are, because actual application of objectives can change one's understanding of them, or even change them. At the end of the process, a utility function based on the objectives is used to rank the alternatives in order to support the selection of a preferred one.

In some situations, there is no single, coherent decision-maker whose objectives need to be understood and implemented. Instead, there are "stakeholders" whose diverse interests need somehow to be reflected. The non-trivial effort of eliciting a coherent set of objectives from a single, purportedly sane decision-maker grows significantly when diverse stakeholders are involved; an example discussed later in the report (under "prioritization") illustrates this. In yet other situations, decisions are made on behalf of stakeholders who are not official "decision-makers," though they may have input through public comment processes. For cases in which decisions are being made on behalf of others, some authors argue on ethical and philosophical grounds that different approaches to ranking alternatives are appropriate. These other approaches are presented as if different from utility-function-based approaches. These approaches are discussed elsewhere in the report.

The issues raised in the preceding paragraph pertain directly to safety decision-making, and only indirectly to research decision-making. Nevertheless, their effect on research decision-making is significant, even though indirect. Two main themes of this report – uncertainty analysis and considered formulation of decision objectives – are discussed within straight utility-based decision-making. The above issues only reinforce the emphasis placed on these topics.

Utility Theory

For several reasons, ranking decision alternatives is frequently harder than it may at first seem. Competing objectives may need to be addressed, making it nontrivial to rank-order alternatives on a single axis, because tradeoffs between objectives need to be assessed. Uncertainties may need to be assessed, and these can have subtle effects on ranking when a decision-maker's relative preference regarding certain alternatives depends non-linearly on some attribute such as cost. The concept of "utility" is introduced to deal with these and other issues.

"Utility" is a figure of merit for a decision alternative that reflects how successfully the decision-maker's values and preferences will be addressed by implementing that alternative. Utility theory has its roots in economics, where the term "utility" refers to the real or illusory ability of a good or service to satisfy a human want.

The principle of maximum expected utility (MEU) states that a rational decision-maker should choose an action that maximizes expected utility. By the "expected utility" of a given alternative, we mean the utility of that alternative averaged over the uncertainties in the consequences of that alternative.

When uncertainties are present, it is useful to think of each decision alternative as a "lottery:" analogously to various state-sponsored lotteries, it has different possible outcomes, each with associated consequences and associated probability. The concept of utility can then be developed with reference to the idea that a given decisionmaker may prefer some lotteries (alternatives) to others, even if they have the same *expected* consequences. For example, a risk-averse (but still "rational") decisionmaker may prefer an alternative guaranteeing a so-so outcome, if other alternatives having equivalent *expected* consequences involve even a small chance of a very bad outcome. This idea is related to current agency discussions of "structuralist" and "rationalist" approaches, as will be seen later in this report when defense in depth is discussed.

Based on the principle of MEU, utility theory can be used to foresee or direct the choice that the decision-maker will make, or should make, among the available alternatives. This goal can be accomplished by assigning a utility to each of the possible (and mutually exclusive) consequences of each decision alternative. The utility reflects a decision maker's preferences such that: (i) the numerical order of utilities for consequences preserves the decision maker's preference order among the consequences; (ii) the numerical order of expected utilities of alternatives preserves the decision maker's preference order among these alternatives.

The ability to map preferences into a single number for ranking purposes is not a given. It does, however, follow from certain axioms (Ref. I-4, pp 474-475):

- Orderability: Given any two states, a person must either prefer one to the other or else rank the two as equally preferable. In short, a person should know what he/she wants.
- Transitivity: Given any three states, if a person prefers A to B and prefers B to C, then the person must prefer A to C.
- Continuity: If some state B is between A and C in preference, then there is some probability p for which the person will be indifferent between getting B for sure and the lottery that yields A with probability p and C with probability 1-p.
- Substitutability: If a person is indifferent between two lotteries, A and B, then the person is indifferent between two more complex lotteries that are the same except that B is substituted for A in one of them.
- Monotonicity: Suppose there are two lotteries that have the same two outcomes, A and B. If a person prefers A to B, then the person must prefer the lottery that has a higher probability for A (and vice versa).
- Decomposability: Compound lotteries can be reduced to simpler ones using the laws of probability.

The axioms only discuss preferences, and not an actual utility function.

Two principles are then used to determine the existence of the utility function:

- Utility principle: if a person obeys the above axioms of utility, then there exists a real-valued function U that operates on states such that U(A) > U(B) if and only if A is preferred to B and U(A) = U(B) if and only if there is no preference between A and B.
- Maximum Expected Utility Principle: The utility of a lottery (or complex scenario) is given as:

$$U([p_1, S_1; \ldots; p_n, S_n]) = \sum_i p_i U(S_i)$$

where

$U(S_i)$ is the utility of possible outcome state S_i,

p_i is the probability of outcome i.

Once the probabilities and utilities of the possible outcome states are specified, the utility of a compound lottery involving these states can be computed.

Determining the general form of the utility function (e.g., whether it is a weighted sum of individual-attribute utilities, or other form) requires more assumptions and arguments. Quantifying its *parameters* is yet another significant step. Finally, a good deal of literature explores departures from certain of these axioms. These topics are beyond the scope of the present report. The main point is that in all but very simple situations, establishing the decisionmaker's preferences so that alternatives can be appropriately ranked is a significant undertaking. As a practical matter, most decisions have a number of attributes that must be considered. Decisions where the outcome is characterized by two or more attributes are handled using multiattribute utility theory (MAUT). Prioritization is a case of this, and different attributes are discussed in Section I.4.

I.3 HOW UNCERTAINTY AFFECTS SAFETY DECISIONS

The expected utility of decisions is limited by uncertainty. In the case of a safety decision, uncertainty may drive a decision to a conservative result that may impose burden that would not be considered necessary if better information were available. Alternatively, uncertainty could allow insufficiently conservative decisions to be made. Therefore, the utility of safety decisions is reduced by uncertainty, in that either unnecessary risk is incurred, or unnecessary preventive measures are taken.

Research may be undertaken to reduce uncertainty, in the belief that this will lead to decisions having greater utility. Formal decision analysis offers a way to think systematically about the actual value of doing such research (i.e., the "value of information"). The present section briefly discusses this idea. It has obvious potential value in setting research priorities, but in addition, understanding the mechanics of the process can shed light on how to think more productively about uncertainty itself. In some contexts, the thought process in this example is familiar to many regulators, without the trappings of decision theory. One point here is to show how more formal decision theory provides a way of looking at uncertainty in much more complex examples. A more thorough discussion of uncertainty is given in Part II.

Consider a hypothetical decision regarding what power level to allow in a particular LWR design, assuming essentially fixed emergency core cooling system (ECCS) capability. Higher power has more benefits: it enables more electricity production. But it also requires more cooling if fuel damage is to be avoided in certain postulated accidents. The hypothetical decision weighs these two factors. Accident consequences, simplified here down to consideration of peak cladding temperature (PCT), are modeled for various power levels using a complex code. There is uncertainty in the code prediction, and depending on how criteria are to be applied, this can be a complicating factor.

The simplicity of this example is not intended to suggest that safety is simply being traded off against economics without regard to other priorities. Reasonable assurance of adequate protection is the standard to be applied in licensing. But licensing a commercial plant in the first place is predicated on some net benefit associated with power plant operation, and the net benefit is clearly correlated with power output. The point of this example is that decision analysis helps to discuss the significance of residual epistemic uncertainties that affect the licensed power level. (Kinds of uncertainty, e.g., epistemic, are discussed in Part II.)

Figure I.1 shows a simple influence diagram for this decision. (See Ref. I-2 for a discussion of influence diagrams, and references to original work.) The influence diagram shows the flow of the calculations needed to quantify the expected consequences of each decision alternative. (Influence diagrams are widely used for this, but this example is too simple to illustrate why such diagrams might be useful.) On this diagram, the yellow block is a "decision" node (what power level to allow); the rounded blue rectangles are "value" nodes, which are either data or calculations. The node labeled "UTILITY" represents the figure of merit by which this decision will be made in this example. The green oval is a "chance" node, corresponding to a key input uncertainty for which there is an underlying prior probability distribution. In this case, the key uncertainty is PCT in a particular accident scenario for a particular power level. (In a real case, many more variables would be involved. This example is oversimplified for purposes of illustration.)

If we decide to restrict power to a low level, the outcome is that we forgo a certain amount of production, and perhaps lose economies of scale. If we decide to allow higher power, we accept a greater chance of a higher PCT.

Node	Significance
ECONOMICS	Value node that quantifies economic benefit in terms of power level
SAFETY	Value node that quantifies the level of “safety,” which in this case is taken to related to the probability that PCT exceeds 2200 degrees, and by how much)
PCT	Value node: Peak cladding temperature (conditional on accident) for a particular licensed power level
LICENSED_POWER_LEVEL	Decision node: what power level to allow
UTILITY	Value node: quantifies utility of each decision alternative

If there were no uncertainty in PT, the power level in this example would be derived simply by solving the model for the power level that gives a PCT just under 2200. Since there is uncertainty, the decision on power level must weigh the probability that each given power level would lead to exceedance of the PCT criterion (conditional on the accident, in this example). The utility function, which needs care in its development, reflects the tradeoff between safety and economics. It will have higher values if the economics are more favorable at a given level of safety, and lower values if PCT is exceeded at a given level of economics. A “conservative” utility function would insist on driving this probability down to a very low value, which (refer to Fig. 1.2) would mean driving the power level down until code uncertainty essentially did not overlap the PCT criterion.

If code uncertainty is large, there may be substantial benefit to reducing it, because a higher power level could be allowed. The “value of (perfect) information” in this case is the increase in utility to be “expected” if we could narrow down code uncertainty to zero. This expected increase is a function of the code uncertainty distribution and the utility function; there is some probability that (in reality) allowable power could increase by a lot, some probability that allowable power could increase only a little, some probability that the allowable increase would be in between. There may even be a chance that the allowed power level would *decrease*. Decision analysis software carries out the implied calculation to quantify the value of perfect information, in units of the utility function.

By “perfect” information, we mean the elimination of all uncertainty. The actual value of a program to reduce uncertainty will likely be less than the value of “perfect” information, because it is not typically practical to eliminate *all* uncertainty. The above type of calculation is therefore an upper bound on the value of any real activity.

Figure I.1
Influence Diagram for Analyzing the Value of Reducing Code Uncertainty

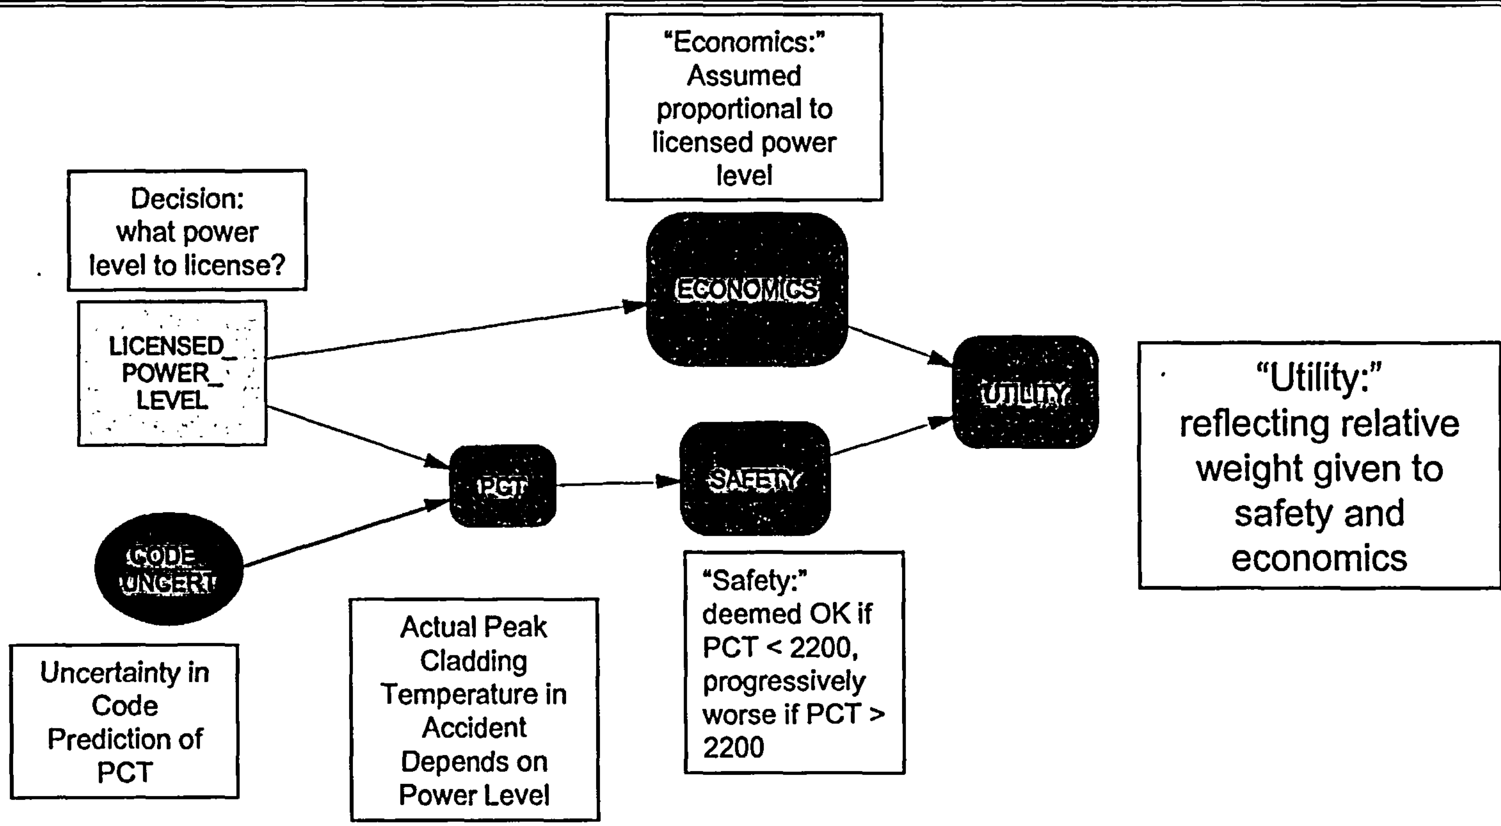

Reduction in Code Uncertainty Means Greater Utility of Decision:
Good Chance that Higher Power Level Can Be Licensed, *Even at the Original Confidence Level*

Figure I.2
Example of Value of Information:Value of Reducing Code Uncertainty

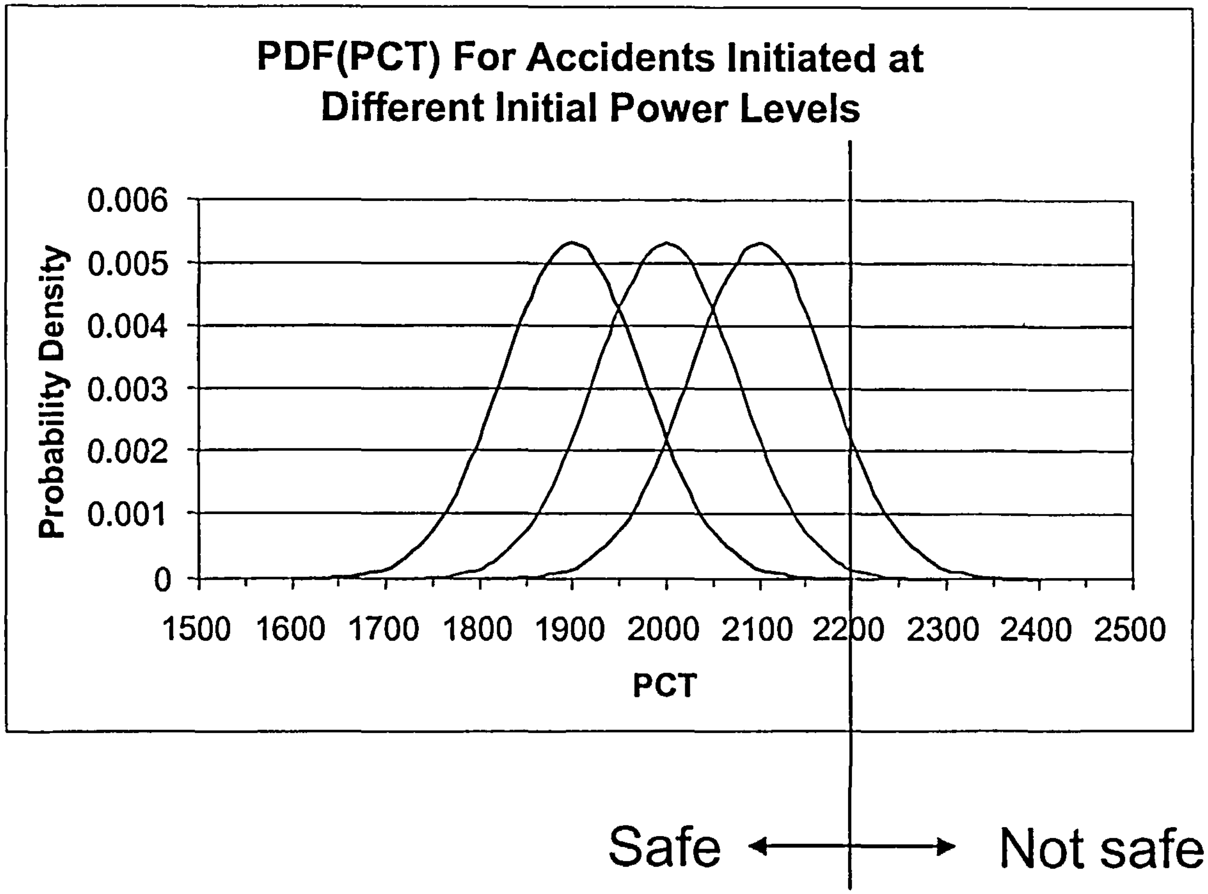

Observations

In principle, even as an upper bound calculation, this thought process has potential application to prioritization of all kinds of research programs. In the case of an advanced reactor, many types of uncertainty may be involved, and reasoning of this kind may help to set priorities in addressing them.

The value of reducing uncertainty can be *zero*. If other factors limited power to a level at which PCT were not an issue, then reducing code uncertainty would not be valuable in the context of this specific decision. This does not necessarily mean that a program in that area would be wholly without merit, but it would not be useful for purposes of the decision at hand.

If this approach is to be used to set priorities, the uncertainty distribution has to be reasonably honest. Deliberately overstating uncertainty is potentially costly. Understating uncertainty is potentially unsafe. Within the decision paradigm discussed here, conservatism belongs in the utility function (if anywhere), not in the uncertainty distributions.

Summary Observations:

Uncertainty limits the utility of decisions, and reducing it has value. Better knowledge (reduced uncertainty) leads to a better allocation of safety resources and agency resources. In the context of a decision model, one can actually quantify an upper bound on the value of a candidate research project that would, if approved, be undertaken in order to reduce uncertainty. Therefore, evaluating the "value of information" to a given safety decision has the potential to inform a research prioritization decision.

In order for this potential to be realized, the uncertainty must itself be understood. Uncertainty is discussed in some depth in Part II.

I.4 OVERVIEW OF PRIORITIZATION

I.4.1 Prioritization as a Special Case of Decision-Making

By "prioritization," we mean rank-ordering a set of items according to a metric that reflects "priority."

Assignment of "priority" to an item is a decision, and prioritization is therefore a special case of decision analysis. Table 4.1 below shows how the general steps in decision analysis specialize in the application of decision methods to prioritization.

Section I.4.2 discusses simple cases of prioritization. By "simple," we mean cases in which the absolute priority one would assign to a given item is not dependent on the absolute priority assigned to other items. In general, research planning is more complicated than this. It is not sufficient to assign a "priority" to each of a list of possible research activities. No single research activity will meet all the demands placed on RES, so a collection of activities must be identified, and the elements of this collection need to fit together in a satisfactory way. For example, the program should adequately cover the range of issues facing decision-makers in the near term, adequately balance near and long term research interests, and so on. These issues will be taken up later in the report.

I.4.2 Steps in Prioritization

This subsection discusses simple prioritization, based on four sources: a recent paper by Weil and Apostolakis (hereafter WA) (Ref. I-5), presenting a methodology for prioritization of the review of operating experience; a paper by Apostolakis and Pickett (hereafter AP) (Ref. I-6), discussing in some depth the issues encountered in dealing with multiple stakeholders; a National Academy report on research management, including prioritization, within DOE's Office of Environmental Management (hereafter NAS) (Ref. I-7); and a recent DOE Standard on risk-based prioritization (Ref. I-8), which offers comparative perspective on high-level objectives that can usefully be considered when developing the objectives hierarchy.

The papers cited above are not the only recent papers on prioritization. Many papers are written on the subject. Some consider only safety and cost metrics, and essentially establish simple cost-benefit ratios for use in selecting among possible safety backfits. Others assign priority based on a much broader range of objectives. Even among the latter works, not all papers adhere to conventions regarding the development of utility functions, or carefully consider the form of the utility function. The present selection of the papers cited above is driven by the breadth of the range of objectives considered, and the attention paid to conventions regarding development of utility functions.

Earlier in this report, it was remarked that the discussion covers two kinds of decision-making: safety decision-making and programmatic decision-making. The present subsection is aimed primarily at programmatic decision-making. In addition to considered formulation of the appropriate objectives hierarchy and performance measures, safety decision-making requires more substantial effort in modeling certain other areas that are addressed later in this report.

The following steps to accomplish simple prioritization are given by WA:

structuring the objectives;
determining appropriate performance measures;
weighing objectives and performance measures;
assessing utility functions of performance measures;
performing consistency checks;
validating the results.

These steps result in a formula that can be used to assign a numerical rating to each alternative evaluated. Each step will be discussed below, bringing in considerations remarked by the other sources cited above.

Table I.1 Specializing Decision Analysis to Prioritization

Step	Application in General Decision-Making	Special Focus in Research Prioritization
Problem Formulation	In general, this refers to characterization of a situation that requires some kind of response from a decision-maker, or to a situation in which pro-active measures could be taken (see "value-focused thinking" in Part III).	The "problem" in prioritization is to determine how best to support near-term and longer-term agency decision-making. This has two elements: determining what technical uncertainties most complicate safety decisions, and determining how best to address those uncertainties.
Objectives	Formulation of the goals, fundamental objectives, and means objectives in light of which the decision alternatives are to be evaluated.	There are two kinds of objectives to consider: objectives driving safety decisions, and objectives driving programmatic decisions.
Alternatives Formulation	Considered formulation of alternatives is usefully undertaken after consideration of the decision objectives. As discussed elsewhere in this report, once no more alternatives are developed, the utility of the outcome is limited by the set of alternatives under consideration; formulation of a manifestly preferable alternative can lead to a very good outcome with no need for further analysis.	At one level, this refers to formulation of good individual research program alternatives; at another level, to formulation of good *combinations* of research programs.
Consequence Modeling	This refers to predicting how well a given decision alternative will perform, if implemented.	The consequences of many research programs will be reductions in technical uncertainties that would otherwise limit the expected utility of safety decisions. Some research programs develop new methods. The projected consequences of these decisions need to be considered below under "linked decisions."
Tradeoffs	Given multiple objectives to be addressed, it may occur that one of two alternatives is superior in one respect and inferior in the other. "Tradeoffs" refers to the process of establishing which objective has greater weight.	In prioritization, the cost of proposed programs and their expected benefits need to be assessed against each other. Proposed research in different areas may address different safety objectives, calling for tradeoffs at other levels.

Table I.1 (Continued)

Step	Application in General Decision-Making	Special Focus in Research Prioritization
Uncertainty	Uncertainty in the quantification of expected consequences limits the utility of decisions in general. Assessing the uncertainty is a key part of safety decision-making.	The uncertainty that limits the utility of safety decisions is the driver of many research programs. The uncertainty in the consequences of research program decisions limits the utility of prioritization decisions. Quantitative assessment of program uncertainties is in its comparative infancy, but is being more widely addressed by many federal agencies.
Risk Tolerance	Depending on their utility functions, decision-makers may have different levels of risk tolerance, even if most of their underlying objectives are fairly similar. This means that some decision-makers will be willing to opt for alternatives having a significant potential down-side if the potential up-side is worth it. Others may prefer an alternative with a more limited down-side, even if the up-side is also more limited.	In research prioritization, this refers to several topics. One is the willingness to invest in relatively novel ideas. Projects based on minor extensions of existing knowledge promise small but guaranteed returns; more ambitious programs may simply fail to produce useful results. Such an outcome could have long-range consequences for agency decision-making, whether or not it is explicitly recognized. If a program is explicitly judged negatively in retrospect, this could result in adverse consequences for the decision-maker who chose to invest in it. On the other hand, programs with little or no technical risk are unlikely to make substantial progress.
Linked Decisions	In general, decisions to be made in the future will face different options and perhaps different objectives, depending on the outcomes of decisions made today.	Research decisions in particular are linked to subsequent decisions, and thinking of them as linked decisions is essential. This is discussed further in Part III of this report.

I.4.2.1 Structuring the Objectives

The discussion of the objectives hierarchy provided later in this report implicitly contemplates potentially fairly complex objectives hierarchies, intended to foster completeness in a set of performance measures aimed at covering safety issues. For simple prioritization, a simpler construct may well suffice, as in WA.

The problem that WA are trying to solve is deciding which of many hundreds of items of industry-wide operational experience need detailed evaluation at a particular station, in order to support safe and reliable operation of the station. WA settle on the following objectives and associated performance measures as representing their clients' priorities:

Objective	Performance Measures
Perform timely evaluation	Time cushion (how soon the plant is likely to experience the condition)
Relevance to station	Precursors in common, design similarities (degree to which this can apparently happen at the plant)
Minimize costs	Cost
Protect environment	Anticipated environmental impact (if the event occurs at the plant)
Maintain public confidence	Anticipated public criticism (if the event occurs at the plant)
Safety	Industrial safety significance, "nuclear" safety significance (if the event occurs at the plant)
Maintain positive relationships with regulator	Regulator interest

Several of these arguably pertain to research planning. The "timeliness" objective specifically pertains to advanced reactor work, albeit on a time scale different from that contemplated by WA.

For comparison, following is a set of objectives identified by DOE (Ref. I-8) for application in "risk-based prioritization" of DOE work (these are meant to be refined in the context of specific applications). The context of these objectives is deciding what tasks most need doing at DOE facilities, examples being site decontamination or facility backfits. This sufficiently resembles WA's application to make comparison of the two lists illuminating.

High-Level Decision Objectives (Ref. I-8):

Maximize Accomplishment of Mission
Minimize Adverse Effects upon Public Health and Worker Safety
Minimize Adverse Effects upon the Environment
Maximize Compliance with Regulations
Minimize Adverse/Maximize Desirable Socioeconomic Impacts
Maximize Safeguards and Security Integrity
Maximize Cost Effectiveness
Maximize Public Trust and Confidence

This list addresses both facility performance and facility safety with respect to several risk metrics, and a suite of external indices (public confidence, etc.) that compare with the list used by WA.

Institutional objectives are addressed more explicitly in the chapter by Dehmer ("Assessing the Value of Research at DOE") in Ref. I-9. The DOE Office of Science's mission includes both research and construction and operation of facilities as top-level fundamentals; its goals bear comparison with those of NRC Research. Its four goals are

(1) leadership in scientific research relevant to energy (including environmental impact),
(2) fostering dissemination of results,
(3) world-class scientific user facilities, and
(4) "steward[ship]" of human resources, essential scientific disciplines, institutions, and premier scientific facilities.

Although pitched at a very high level, these objectives could, if taken literally, have a significant practical impact in the formulation of research alternatives and in ranking research proposals. This would come about because the resulting performance index would tend to favor research proposals whose side effects included training of staff in specific disciplines, maintaining funding to key test facilities, and so on. In a similar vein, the Advanced Reactor Research Plan alludes to coordinating research with other entities that sponsor research. It is clear that a modification of the objectives considered above could reflect this consideration.

I.4.2.2 Determining Appropriate Performance Measures

The measures defined by WA for their objectives are also shown in the above table.

While the measures illustrated in WA are designed to be simple to apply, some measures may require in-depth evaluation. A main theme of the present report is that one key performance measure to be considered for research programs is the "value of information" (VOI) associated with the program, as evaluated in the context of safety decisions. (VOI is discussed later in Part I and more extensively in Part II.) Prioritization is a decision task that runs somewhat in parallel with another decision task, "safety decision-making." The safety decision-making task identifies performance measures relating to safety, and assesses their uncertainties as part of formulating research activities to sharpen the technical basis for safety decisions.

The NAS report (Ref. I-7) cites measures that might usefully be considered in prioritizing research. They discuss a figure of merit for proposed research projects, determined as

$$\frac{P_T * P_D * \Delta_{COST}}{C_D},$$

where

P_T is the probability of technical success of the research project,

P_D is the probability of successful "deployment" (i.e., implementation) of the development by the customer (e.g., the probability of actually making beneficial use of a developed software tool in subsequent work),

Δ_{COST} is the cost saving relative to the existing way of doing things, evidently conditional on program success and deployment success,

C_D is the cost of the project.

This clearly stems from a particular set of program objectives, but is noteworthy in that it recognizes the need to think about the probability of technical success and the usefulness of thinking about the prospects of successful deployment.

Separately, a useful measure cited by NAS is the health risk reduction associated with the work.

I.4.2.3 Weighing Objectives and Performance Measures

The step of actually assigning a ranking to a given item involves quantifying the weight of each performance measure (this step) and the utility functions (next step) in order to support evaluation of the following "performance index" for each item:

$$PI_j = \sum_{i}^{K_{pm}} w_i u_{ij} ,$$

where

PI_j is the performance index for the jth item (the measure of its "priority"),

w_i is the weight of performance measure i,

u_{ij} is the utility of performance measure i for item j (see Section I.2),

K_{pm} is the number of performance measures.

The present step is to obtain a set of weights. The method used by WA is the Analytic Hierarchy Process (Ref. I-10), which is discussed in Part III.

The additive form of this utility function receives a good deal of attention in the literature in general, and certainly from WA. It is also discussed by NAS. The summary perspective offered by NAS is essentially the following. This type of utility function has been criticized by many (see NAS, p. 82), and the AHP is considered by some to be an alternative. (Note that the use of the AHP contemplated in this statement is as an alternative to the whole utility function, not just for determining the weights in the above formula. The AHP is discussed in Part III.) The NAS report declines to take a position on the form of the utility function, but cites a study suggesting that differences between results obtained with the AHP and with the simple utility function have a lot to do with how the problem is structured and how the weighting factors are elicited. There seems to be a broad base of support for the idea that if one is thinking systematically and correctly about one's objectives, and satisfying the "additive independence" condition discussed by WA, this kind of utility function is adequate. Use of the AHP for such things as determining weights as done by WA seems quite widespread, although using the AHP in this way also has its critics (Ref. I-2).

Note that the last two steps in the WA process – consistency checks and "validation" (or benchmarking) – have the potential to identify some kinds of problems with the utility function or ranking procedure, however it is formulated.

I.4.2.4 Assessing Utility Functions of Performance Measures

In WA, constructed scales (discussed in Part III) were used for the performance measures. "Assessing utility functions" means essentially mapping each constructed scale into a range between 0 and 1 (so that the PI formula above will behave correctly), and doing so in a possibly non-linear way that reflects the decision-maker's attitudes and values regarding different levels of performance on that particular measure. This topic was discussed in Section I.2.

I.4.2.5 Performing Consistency Checks

WA, and AP before them, spent a great deal of effort discussing objectives and values with decision-makers and stakeholders, and learned a lot about the difficulty of achieving consensus and/or consistency in the weights when multiple and disparate objectives are being addressed. "Consistency" means that the different sets of inputs (values of performance measures) that can all lead to the same output (performance index) are, in fact, equivalent for purposes of the decision, as implied by their having the same values of the PI. One can derive a set of weights for a group of stakeholders from a single pass through the AHP, but the experience of WA and AP suggest that this first cut is not likely to survive stakeholder or decision-maker review without modification.

ACRS members (e.g., Apostolakis) have recently (Ref. I-11) advocated such consistency checks for the way in which the ROP's Action Matrix processes performance information to arrive at a net measure of plant performance (e.g., should two whites equate to a yellow?).

I.4.2.6 Validating the Results

This refers essentially to "benchmarking:" applying the prioritization process to items for which the priority is already "known," in order to see whether the process gives the "right" answer. According to WA, failure in such benchmarking tests may well relate to a disconnect between the objectives that are built into the performance index and those used originally to obtain the "right" answer. The literature of this subject seems implicitly to suggest that there is much experience with latent objectives emerging only gradually through benchmarking exercises. This point seems especially important when diverse stakeholders are involved. WA and AP devote significant space to discussion of it. This kind of experience seems to drive recommendations such as Keeney's (discussed in Part III) that seem extremely exhaustive to persons who have not themselves had to deal with such disconnects between the recommendations of a decision process and the supposed objectives of the decisionmaker.

I.4.3 Summary Observations on Simple Prioritization

The process described by WA is generically applicable to simple prioritization: that is, it covers the topic of assigning an index to each item in a list for purposes of ranking. Even for this simple application, a considered formulation of objectives plays a critical role. This has been shown by WA and predecessors to be a difficult task when multiple stakeholders are involved. Not only is it necessary to reconcile diverse viewpoints, but also significant iteration may be necessary when the stakeholders begin to see how their stated priorities actually play out in application. The subject of formulation of objectives will receive substantial attention in Part III.

The task performed by WA did not require them to develop the items being prioritized. In WA's application, the "items" were instances of operational experience, and it was only necessary to map them into a priority scheme to determine how they would be addressed. In research planning, the items to be prioritized are research activities, which have to be formulated before they can be prioritized. This process, too, can benefit from a considered development of the objectives.

While prioritization of individual tasks may be necessary for research program management, it is not a sufficient basis for optimizing the overall research portfolio. (The best portfolio will not, in general, result from selection of the highest priority research tasks, if priority of each task is assessed without regard to the status of other tasks.) It is necessary to think at levels above the individual program level. This will be taken up in Section I.5 (under "Research Performance Measures").

I.5 PERFORMANCE METRICS FOR RESEARCH PRIORITIZATION

Although it is beyond the scope of this report to propose a specific objectives hierarchy to be used in research planning, this section discusses considerations in the development of such an objectives hierarchy. Appropriately weighted, the figures of merit discussed below, possibly supplemented by others, might beneficially be used as performance measures in research planning. It is emphasized that these are offered as an illustrative set.

Two kinds of performance measures can be identified. One kind pertains to the merits of individual programs. The other kind of measure relates to the performance of RES at the office level. To see why the two kinds of measures are different, begin by considering the "portfolio problem" in investment. In the portfolio problem, one decides how to allocate one's investment funds over diverse investment options in such a way as to achieve a certain level of expected return while limiting the downside risk of the overall investment portfolio. A good solution will have the property that the portfolio's downside risk is lower than that of many of the individual investments, but the expected overall return is still good. Analogously, a collection of individually high-risk research activities may be significantly less risky than some of the individual elements.

In addition, there are measures such as balance across technical area that clearly do not pertain to individual programs, and must be evaluated at the portfolio level. Examples of each kind of performance measure are discussed below.

Although it is beyond the scope of this report to discuss research management per se, some of the measures discussed here are useful both in prospect (deciding on programs) and in retrospect (gathering feedback to improve the process of research planning).

I.5.1 Measures for Individual Programs

Numerous figures of merit can be identified that might bear on the desirability of funding a particular program proposal. Some are discussed below. Some of these figures of merit can be combined mathematically (as illustrated in Section I.4.2.2) to yield a more targeted figure of merit.

Figures of merit that apply to individual programs, even summed over all funded programs in a research portfolio, do not necessarily measure all agency objectives operating at the portfolio level. For example, coverage of issues may not be an appropriate demand on any single program. Some portfolio-level objectives are discussed in Section I.5.2.

Potential Value of Information / Area under ROC

The "value of information" was discussed in Section I.3. When a regulatory decision is influenced by uncertainty, and a decision model can be constructed for that decision, it is possible to quantify the value of completely eliminating that uncertainty. This will be an upper bound on the value of research on that item in the context of that specific decision. (The research may, of course, have value in other ways.)

In some cases, a more value-free measure of program impact on uncertainty reduction could be based on a concept discussed in Section II.2.2, namely the "Receiver Operating Characteristic" (ROC) curve for a given decision situation.

Cost of Program

This simply recognizes that real value is only rarely achieved without some price, and the merits of a specific program have to be judged in light of its costs.

Probability of Program Success

This item recognizes that program intent may not be realized in execution. This needs to be recognized in program evaluation, because a hopelessly quixotic program could in theory get high marks in potential value of information, which would need to be cancelled by a low success probability.

While this needs to be a consideration at some level, this measure in particular lends itself to being gamed. Trivial programs get a high score on this measure. Long-term research having an innovative character might get a low score on this measure. It is appropriate to offset this measure by recognizing at the portfolio level that some innovative research is necessary.

Probability of Program Implementation / Timeliness

A program targeted at a specific decision is useful to that decision only if its results come before that decision is made. However, "timeliness" arguments sometimes foreclose programs that would have been useful for other reasons. Moreover, it needs to be recognized that some decision timetables are not met, and foreclosing potentially useful research on the basis of overly-optimistic program schedules is undesirable.

Potential of Program to Foster the Development of Improved Decision Alternatives

This is a difficult metric to judge in prospect. But programs whose value of information is speculative could still have significant benefits of this type.

I.5.2 Measures for Portfolios

Some pertinent considerations do not pertain at the level of individual programs. It is necessary not only to judge each program on its own merits, as discussed above, but also to formulate a portfolio having certain properties that can only be evaluated at the portfolio level. This has elements of "formulate decision alternatives," which was discussed extensively in connection with the objectives hierarchy.

Coverage of Safety Issues Facing the Agency

Clearly, a research portfolio needs to cover high-priority safety issues. The high-priority issues list is derived from issues identified in safety decisions, or perhaps tool development needs that are derived from those issues. This relates to the "hedging" measure below.

Hedging Against Technical Uncertainty – Multiplicity of Technical Approaches

Some organizations recommend hedging against technical uncertainty by conducting programs in parallel. If one approach fails, the other may succeed. This is common in time-critical development projects, where alternative design activities may be carried out in parallel until a downselect occurs. A series approach has the potential to save money, e.g., take up approach B only if approach A fails, but this risks project delay in the event that A fails.

Maintaining Proper Balance Between Long-Term Research (Innovative, Anticipatory) and Short-Term Research (Near-Term Issue Resolution)

Some organizations simply budget a certain fraction of research for longer-term research activities. In proposal review schemes that focus on immediate needs, probability of project success, and so on, it is easy for short-term or even trivial research to prevail in review over more challenging research projects. There may be an argument for formulating a long-term portfolio with its own metrics.

Maintaining Intellectual Capital, Research Facilities

This measure is discussed in Part III. It reflects a proxy objective for the decision needs facing tomorrow's decision makers, whose needs may have to be addressed in today's research programs, even though we do not know in detail what those needs are. The usefulness of this objective in prioritization is that there is something to be said for a research portfolio that does this, relative to a research portfolio that does not. This also relates to the "continuity" measure below.

Continuity

Consider two funding scenarios. In one funding scenario, an effort is planned and funded consistently from year to year. In an alternative funding scenario, an effort is funded in an alternating pattern: one year at twice the yearly funding of scenario 1; the next year at zero; the following year at twice the yearly funding of scenario 1; and so on. The two scenarios have the same average funding levels, but different outcomes. For many reasons, the latter scenario is a highly ineffective way to conduct research. Qualified personnel leave during the "zero" years, leading to an inefficient ramping-up process at the beginning of every "on" year. Having fixed deliverables at the end of each "on" year also compresses the time available for real work in "on" years. (Not having deliverables at the end of each "on" year would be even worse, in this scenario.)

The alternative funding scenario may sound unlikely, but has been known to occur.

Considering this measure in the formulation of the research portfolio not only evaluates the portfolio across programs, but along the time axis.

I.6 SUMMARY: APPLICATION OF KEY IDEAS TO ADVANCED REACTOR RESEARCH PRIORITIZATION

I.6.1 Overview

This report has focused on two main areas: development and application of objectives hierarchies, and assessment of uncertainty and what to do about it. The Advanced Reactor Research Plan explicitly recognizes these areas as important, and contemplates activities to deal with them. The present report differs from the Plan in discussing these areas in an explicitly decision-analytic way, and showing where specific tools can help in specific ways with the job of prioritizing research to support decision-making for advanced designs.

It is possible to apply the ideas and tools discussed in this report piecemeal. However, it is useful to summarize the discussion as if they were to be applied as part of a coherent approach. Accordingly, the ideas contained in the main body of this report are summarized below in the context of two essentially parallel "tracks" of activity related to prioritization. These activities use tools that have been discussed in this report to respond to needs that have been discussed in the Advanced Reactor Research Plan.

Before discussing these two tracks, the topic of "defense in depth" will be discussed briefly (Section I.6.2) as a specific example of an area known to be important for some advanced reactor designs. (This topic is also discussed in Part II.) Consideration of this topic from a decision-analysis point of view sheds light on agency decision-making and points to ways of setting research priorities that beneficially focus on issue resolution in this area.

One track (Section I.6.3) is the formulation of a design-specific decision model addressing design performance. The decision is basically the extent to which (or the level of assurance at which) the design meets safety objectives. This development activity results in the formulation of key elements of a regulatory approach, and the prioritization of key technical uncertainties limiting the utility of a decision regarding design adequacy. This track leads to formulation of a regulatory approach, and identifies and ranks technical issues that need to be resolved in the course of implementing that approach. These latter outputs are inputs to "track two."

The other track (Section I.6.4) is formulation and application of a decision model for prioritization. Its activities focus on creating, and then ranking, research programs that best respond to agency objectives in the near term and in the longer term.

I.6.2 Consideration of Defense in Depth in Advanced Reactor Designs

Defense in depth continues to be discussed extensively for existing reactors, for reasons that arguably operate for advanced reactors as well. Therefore, defense in depth is likely to be an issue for advanced reactors. Moreover, it is an issue that relates closely to the major themes of this report: clarification and implementation of fundamental objectives, and treatment of uncertainty. Decision theory can, in principle, improve the way in which defense in depth is handled in regulatory practice.

This subsection presents an abbreviated discussion of this issue, and suggests ways of confronting it in advanced-reactor decision-making.

Consider the hypothetical case of whether a particular reactor design ought to be required to have a containment, and suppose temporarily that "yes" or "no" are the only two alternatives. Suppose that the design promises a very low frequency of significant release of radioactivity from the *fuel*, and that taken at face value, this frequency would easily meet goals on large release frequency from the *plant*. Taking this release frequency at face value, one might consider not requiring a containment. However, on any of three possible grounds – uncertainty in the frequency, balance between prevention and mitigation, and the need for multiple barriers – traditional defense in depth arguments might be brought to bear to defend a requirement for containment.

Formal decision methods and associated tools ought to be of some help, specifically in the following ways.

1. Clarification of the objectives of all stakeholders is well within the purview of formal decision methods (even if the implicit criticism of the precautionary approach is that some of those objectives are neglected in application). This would include objectives of all kinds (safety, cost, common defense and security, ...). Apart from being interesting in its own right, the exercise would support the following steps.

2. Creation of improved alternatives (e.g., confinement, but there is no reason to stop there) is within the purview of formal decision methods.

3. Mainstream decision analysis suggests that construction of a utility function can and should be done in this case. It would be useful to explore the implications of biasing a utility function more in favor of benefits, as done by the British (see Part II).

4. An honest assessment of the uncertainties, including uncertainties other than the merely parametric (i.e., modeling, completeness, ...), would be important. This would require stepping outside of the conceptual framework of the engineering models to contemplate whether the models themselves are adequate.

5. Given the uncertainties and an improved set of objectives, either VOI formalism itself or an analog could be used to assess the change in utility associated with the various decision alternatives, including at least one based on defense in depth.

These recommendations generally comport with the thrust of the "rationalist" option mentioned above, but go beyond the previously-stated rationalist option in two ways: (1) taking a more structured and formal decision-theoretic approach, and (2) in doing so, trying to implement standard methods in a way that folds in the concerns of the precautionist school.

I.6.3 Use of Formal Decision Methods to Assess Design Adequacy

As mentioned above, formal decision methods can be applied on more than one track. On one track, formal methods are used to set up decision models for safety decisions such as "design adequacy." On the way to improved decision-making, these models also point the way to research needed to support these decisions. Specifically, they provide a structured approach to the identification of key uncertainties, and a method for quantifying the value of information that is potentially available from research programs. This information could be applied directly in research prioritization.

Following are specific activities in support of this track:

- Develop design-specific objectives hierarchy for the safety case (iteratively with vendor);
- Establish top-level safety / risk metrics to be used to judge the design;
- Allocate performance at a high level, and far enough down to establish meaningful guidance for systems reviewers and for formulation of regulatory approach (AHP, other methods);
- Assess key uncertainties ("Expert Information," Bayesian nets, AHP);
- Develop decision model (e.g., influence diagram) for "design adequacy";
  (Top level reflects selected metrics;
  Lower levels reflect key uncertainties)
- Use the decision model to determine value of information and value of control for key uncertain variables.

I.6.4 Use of Formal Decision Methods to Improve Research Portfolio

Using inputs from the activities summarized in Section I.3, it is possible to formulate and apply a decision model for development and prioritization of candidate research projects.

Following are specific activities in support of this track:

- Develop a technology-neutral objectives hierarchy reflecting fundamental RES objectives in the context of agency strategic goals and RES's role within the agency, considering research performance measures such as those in Section I.6;

- Determine performance measures for research programs;
  Performance measures should reflect:

  - Value of information associated with uncertainties to be resolved
  - Longer-term priorities
  - Utility of program in addressing cross-cutting issues
  - Institutional Issues

- Determine weights for the performance measures, possibly using the AHP

- Use the objectives hierarchy in the formulation of candidate research programs

Note: Development of objectives hierarchy and assignment of weights require very substantial stakeholder interaction.

I.7 REFERENCES

I.1. Advanced Reactor Research Plan, Office of Nuclear Regulatory Research (USNRC, March 2002).

I.2. Robert T. Clemen, *Making Hard Decisions: An Introduction to Decision Analysis*, Second edition (Duxbury Press, 1996).

I.3. John S. Hammond, Ralph L. Keeney, and Howard Raiffa, *Smart Choices: A Practical Guide to Making Better Decisions* (Harvard University Press, 1998).

I.4. Stuart Russell and Peter Norvig, *Artificial Intelligence, A Modern Approach* (Prentice-Hall, Englewood Cliffs, NJ, 1995).

I.5. "A methodology for the prioritization of operating experience in nuclear power plants," R Weil and G. Apostolakis, Reliability Engineering and System Safety **74**, pp. 23-42 (Elsevier, 2001).

I.6. "Deliberation: integrating analytical results into environmental decisions involving multiple stakeholders," G. E. Apostolakis and S. E. Pickett, Risk Analysis 18, pp. 621-634 (Society for Risk Analysis, 1998).

I.7. *Decision Making in the U. S. Department of Energy's Environmental Management Office of Science and Technology*, National Research Council (National Academy of Sciences, 1999).

I.8. Guidelines for Risk-Based Prioritization of DOE Activities, DOE-DP-STD-3023-98 (US Department of Energy, 1998).

I.9. *Assessing the Value of Research in the Chemical Sciences* (National Academy Press, 1998).

I.10. T. L. Saaty and L. G. Vargas, *The Logic of Priorities* (Kluwer-Nijhoff Publishing, 1982).

I.11. G. Apostolakis, Presentation to Nuclear Safety Research Conference, Proceedings To Be Published (2002).

Part II: Treatment of Uncertainty

II.1 OVERVIEW

This part of the report is focused on the treatment of uncertainty in decision-making. As noted in Part I, uncertainty is one of the factors that warrants taking a formal approach to decision-making, and the literature on the subject is already considerable and still growing. The purpose of this part of the report is not to substitute for that literature, but to note certain ideas that seem relevant for agency decision-making in general, and especially research prioritization.

For some purposes, what matters most about a key parameter is its mean value. However, it should be clear from the discussion in the main body that uncertainty can drive decisions, and an honest assessment of uncertainty is obviously essential to a meaningful VOI calculation. This makes uncertainty assessment a key element of decision-making for research programs. It is not enough to say that a parameter is uncertain; it is necessary to say how uncertain it is, in which direction(s), and what the impact of those uncertainties is on decision-making. This is not the same thing as the usual assertion that quantification is not complete without addressing uncertainty; formal decision-making drives the uncertainty analysis in a much more focused way than when modeling is undertaken outside of a specific decision context.

Part I of the report provided a conceptual overview of "value of information:" the value of reducing uncertainty in a specific decision context. That discussion shows how uncertainty limits the expected utility of decisions, and how the worth of reducing uncertainty can be analyzed.[1] For some purposes, one could stop there: those concepts suffice to enable an analyst to set up a decision model in modern software (assuming that the analyst is competent to assess uncertainty), determine which alternative maximizes utility, and obtain a "value of information" output. However, a significant body of fundamental work from the latter half of the previous century provides a somewhat deeper understanding of how prior information and loss functions (utility, cost, ...) interact to affect decisions. It provides an improved perspective on why decision analysis models give the answers that they do. That work is discussed briefly in this part of the report.

The outline is as follows.

The first topic is "simple" binary hypothesis testing: the problem of deciding, on the basis of current information, which of two competing hypotheses is correct in a particular situation. (Example: whether pump reliability has declined from its previous level. One hypothesis is that it has, and the other is that it has not.) This problem contains many of the concepts that one needs in order to cope with uncertainty in many decision contexts. Within Bayesian approaches, the best decision in this problem is a function of prior probabilities of the competing hypotheses, conditional probabilities relating the current evidence to the hypotheses, and the benefits and the costs associated with correct and incorrect decisions.

Next, "Receiver Operating Characteristic" (ROC) curves are briefly summarized. ROCs are a very interesting and general way to look at the binary hypothesis testing problem. Viewing a given decision problem through the lens of the ROC formalism shows visually how difficult the decision is made by the prevailing uncertainties. The area under the ROC is, in a sense, a measure of how well current information supports the current decision.

---

[1] Over the years, many types of uncertainty have been discussed. Recent discussion (RG 1.174) (Ref. II-1) has referred to two types: aleatory and epistemic. "Epistemic" refers to state-of-knowledge uncertainty; "aleatory" refers to randomness in the variables being modeled. Reliability modeling addresses the aleatory character of the timing of system failures that arise from causes that are modeled as occurring stochastically. However, where we are unsure of the values of certain parameters, our uncertainty is epistemic. Some model parameters characterize *population* variability in the thing being modeled; this corresponds to aleatory uncertainty (Ref. II-2). For research prioritization, interest focuses on reducing epistemic uncertainty.

In Section II.3, other approaches to inference are discussed briefly. It is necessary to recognize that some stakeholders do not necessarily accept the paradigm informing most of the present report – that decisions should minimize an overall risk function, or optimize an overall utility function. Some stakeholders seek essentially to pre-empt utility-based arguments under some circumstances. This will be discussed briefly.

Most of this report (and much of the literature) addresses what to do about uncertainty, *given a distribution.* In Section II.4, some discussion is provided of how one might *assess* uncertainty distributions. This subject is less well developed than working with uncertainty distributions once they are given. Code Scaling, Applicability, and Uncertainty Evaluation (CSAU) is mentioned briefly. A vast literature on expert elicitation exists; an old idea of Kaplan's is also cited as a way of highlighting some of the issues.

II.2 THE IMPACT OF UNCERTAINTY ON UTILITY-BASED DECISION-MAKING

II.2.1 Statistical Hypothesis Testing

Statistical hypothesis testing is a fundamental problem whose essential ingredients are present in many important decision-making contexts. Its basic theory is typical of the utility-based (academic mainstream) approach to decision-making under uncertainty. The problem is to decide which of two competing hypotheses is correct in a given situation, based on observed (but not necessarily conclusive) evidence. A useful overview of statistical hypothesis testing and related subjects is provided by Van Trees in Ref. II-3. The present subsection simply states the essential result; derivation of it is beyond the scope of the present report.

For present purposes, it suffices to focus on "binary" hypothesis testing (that is, choosing which of two hypotheses is correct). Label the two hypotheses *H0* and *H1*. In looking at the general results below, it will be useful to carry along the following example. Suppose that the decision problem is one of performance assessment, and the issue is whether current reliability data for a given pump are, or are not, consistent with that pump's historical performance. *H0* is the hypothesis that the pump's performance has not changed; *H1* is the hypothesis that it has.

The essential result of the theory is a *threshold* value (in the running example, the threshold value corresponds to a number of observed failures.). When the number of observed failures is below this threshold, the "best" decision is to declare *H0* (i.e., yes, the data are consistent with the industry average. When the observed number of failures is above this threshold, the best decision is to declare *H1* (no, the data are not consistent with the industry average). The "best" decision is the one that minimizes *risk*. In some decision contexts, the risk metric will be logically equivalent to simply minimizing the overall probability of error in declaring one of the hypotheses to be correct; but in general, some ways of being wrong are more costly than others, and it is necessary to address this.

In the running example, we can identify four possible outcomes, each having its own costs and benefits:

1. Correct declaration of *H0*. Cost: collecting and processing the information.
2. Correct declaration of *H1*. Cost: collecting and processing the information, and coping with the performance issue. Benefit: correctly responding to declining performance.
3. Incorrect declaration of *H0*. Cost: collecting and processing the information, and incurring risk while degraded performance persists uncorrected.
4. Incorrect declaration of *H1*. Cost: collecting and processing the information, and responding to a non-existing performance issue, with all that this entails in lost confidence, costs to licensees, costs to the agency, etc.

The theory addresses the consequences of misdiagnosis by quantifying risk in terms of the *expected consequences* of the decision, as follows. (The derivation is skipped; only the result is given below.) Let

P_0 be the prior probability that *H0* is correct,

P_1 be the prior probability that *H1* is correct,

C_{00} be the cost associated with correctly declaring *H0*,

C_{11} be the cost associated with correctly declaring *H1*,

C_{01} be the cost associated with incorrectly declaring *H0* when *H1* is correct,

C_{10} be the cost associated with declaring *H1* when *H0* is correct,

$p_{r|H_1}(R|H_1)$ is the probability density of the observed variable, given that H_1 is correct, and

$P_{r|H_0}(R|H_0)$ is the probability density of the observed variable, given that H_0 is correct.

(In Van Trees' treatment, all of the outcomes are telescoped into a single parameter called "cost."). Then the *threshold* value of the observed variable that minimizes risk is determined by whether

$$\frac{p_{r|H_1}(R|H_1)}{p_{r|H_0}(R|H_0} > \frac{P_0(C_{10} - C_{00})}{P_1(C_{01} - C_{11})}.$$

In this treatment, costs and benefits of each outcome have been rolled into the "C" parameters, with benefits treated as negative costs. Van Trees adopts this formulation because the inequalities $C_{10} > C_{00}$ and $C_{01} > C_{11}$ (i.e., it's always worse to be wrong) play a role in the derivation of the above relationship determining the threshold.

Note that in the terminology of decision analysis, "risk" does not necessarily equate to CDF. Here, it means the expected value of "cost." In the present context, it is this overall "risk" metric that needs to be minimized.

If the above inequality is satisfied, then *H1* is declared; otherwise, *H0* is declared. This requirement compares the likelihood ratio of observed evidence (the left-hand side) with a quantity that depends on the consequences (expected costs) of the possible assignments. The next subsection shows how this maps back into a threshold value of the observed variable. The practical effect of minimizing the "risk" in the running example is that the setting of the failure threshold depends on how the costs of false alarms weigh against the risk of unidentified performance issues.

An essential feature of this formalism is its use of prior probabilities. Kelly (Ref. II-4) has provided an interesting comparison of a simplified version of this formalism with a frequentist approach that does not make use of prior probabilities, and correspondingly yields a different threshold. Consistent with the above formalism, his conclusion is that use of prior information yields better decisions.

II.2.2 "Receiver Operating Characteristic" (ROC) Curves

The discussion in this subsection is based on discussions provided by Swets *et al.* (Ref. II-5).

Long before decision theory reached its present stage of development, "receiver operating characteristic" (ROC) curves were used for making a certain kind of decision in the face of a certain kind of uncertainty. ROC curves essentially go one step beyond the discussion provided in 2.1. The ROC formalism will be summarized here, because it illustrates certain issues surrounding uncertainty. This formalism does NOT (necessarily) differ in content from more modern methods applied to the same problem type; it differs in presentation and emphasis.

In the previous subsection, the discussion was carried out in terms of "hypothesis testing," perhaps the most general and abstract formulation of the essential problem. For purposes of developing intuition, other kinds of examples are useful. In general, we need to decide which state of nature obtains in a given situation, based on information (e.g., a measurement) that does not give us an unambiguous answer. One type of problem treatable by ROCs is typified by a certain kind of issue in medical diagnosis. Suppose that we are faced with deciding whether a particular patient has a particular medical condition, and suppose further that a diagnostic test is available for this condition. For some conditions and some tests, the test result may be completely unambiguous, but this situation does not always obtain. It may be the case that test results from a population of patients having this condition can fall anywhere within a broad range of values; it may also be the case that test results from the population of patients NOT having this condition can also fall anywhere within a broad range of values. If the two ranges overlap, and the test result falls within the region of this overlap, then the test alone has not conclusively established whether the patient has the condition, and further thought is needed. There is a penalty for failing to identify a condition needing treatment; the patient will suffer consequences that could have been avoided. There is also a penalty for treating a condition that is not present, measured in cost, inconvenience, the episodic risk associated with treatment, and unnecessary chronic side effects.

The type of problem originally addressed by ROCs was akin to deciding whether or not a given radar echo represents an enemy aircraft. In this case as in medical diagnosis, there can be variability in the kinds of echoes produced by the two kinds of targets, and there are penalties associated with an incorrect interpretation. For present purposes, the problem type is a decision regarding whether an adverse condition is present, based on measurement of a particular variable, subject to uncertainties of the kinds mentioned above.

What the ROC formalism does is show how to determine a threshold value, such that if the measurement comes in above this threshold value, the decision is that the adverse condition is present, and if the measurement comes in below this value, the decision is that the adverse condition is not present. The threshold that maximizes expected utility depends on the consequences of correct and incorrect interpretations, and on the prior probability of having the adverse condition.

<u>Details</u>

Key points of the ROC formalism are illustrated on Figures II.1-II.4. Figure II.1 shows how the ROC reflects the overlap between two probability density functions and the decision threshold. Figure II.2 goes into more detail on false positive and false negative decisions. The example considered on Figure II.2 (loosely based on Swets *et al.*) is that of diagnosing glaucoma based on intraocular pressure. In this case, the horizontal axis is pressure, the leftmost distribution schematically represents the distribution of pressure seen in "normal" individuals (for convenience, the curve actually drawn is a pure Gaussian, which is not necessarily representative of reality), and the rightmost distribution schematically represents the distribution of pressures seen in individuals who have glaucoma (again, the curve actually drawn is a pure Gaussian).

The question is where to set the decision threshold. If the threshold is set as in Figure II.2, then some individuals who have glaucoma will be diagnosed incorrectly as being normal (represented by the shaded area to the left of the decision threshold). This is a "false negative." Some normal individuals (represented by the diagonally-hatched area to the right of the decision threshold) will be incorrectly diagnosed as having glaucoma; this is a "false positive." The conditional probabilities of false negative and false positive diagnoses correspond to the *areas* indicated on the figure.

The best threshold setting will represent an optimal balance between these outcomes. For a given decision situation, it is determined as follows.

1. Plot the ROC. As shown on Figure II.2, each point on the ROC corresponds to a particular threshold setting, and for that setting, its coordinates are (P(false positive), P(true positive)). As shown on Figure II.3, the shape of the ROC reflects the overlap between the two uncertainty distributions. If there is very little overlap, as in Figure II.3(a), it is easy to find a setting where the true positive probability is essentially unity and the false positive probability is essentially zero, and the ROC is almost a right-angle elbow. If there is essentially complete overlap, as in Figure II.3 (c), there is essentially no good threshold setting; the probability of true positive only slightly exceeds the value of false positive, and the ROC is almost a diagonal line. For the case of partial overlap, as in Figure II.3 (b), the ROC has an in-between shape.

Figure II.1
"Receiver Operating Characteristic" (ROC) Curve

The ROC is a plot of

P(False positive | Threshold setting AND Adverse Condition Is NOT Present)
vs.
P(True Positive | Threshold setting, Adverse Condition IS Present) for all possible threshold settings

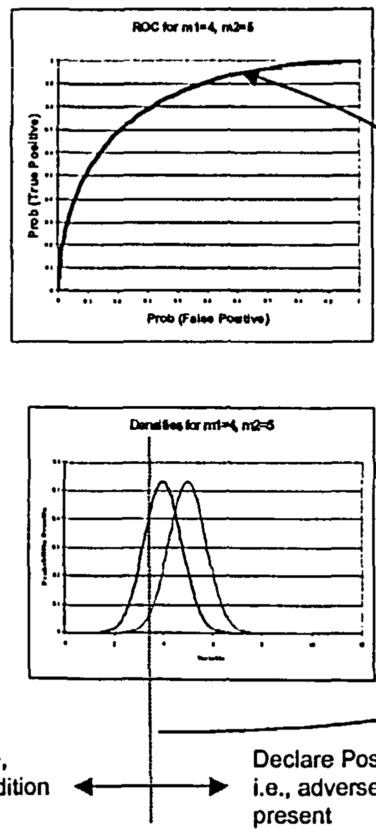

Each threshold setting corresponds to a point on the ROC

Declare Negative, i.e., adverse condition is *not* present

Declare Positive, i.e., adverse condition is present

Figure II.2
Details of ROC Development

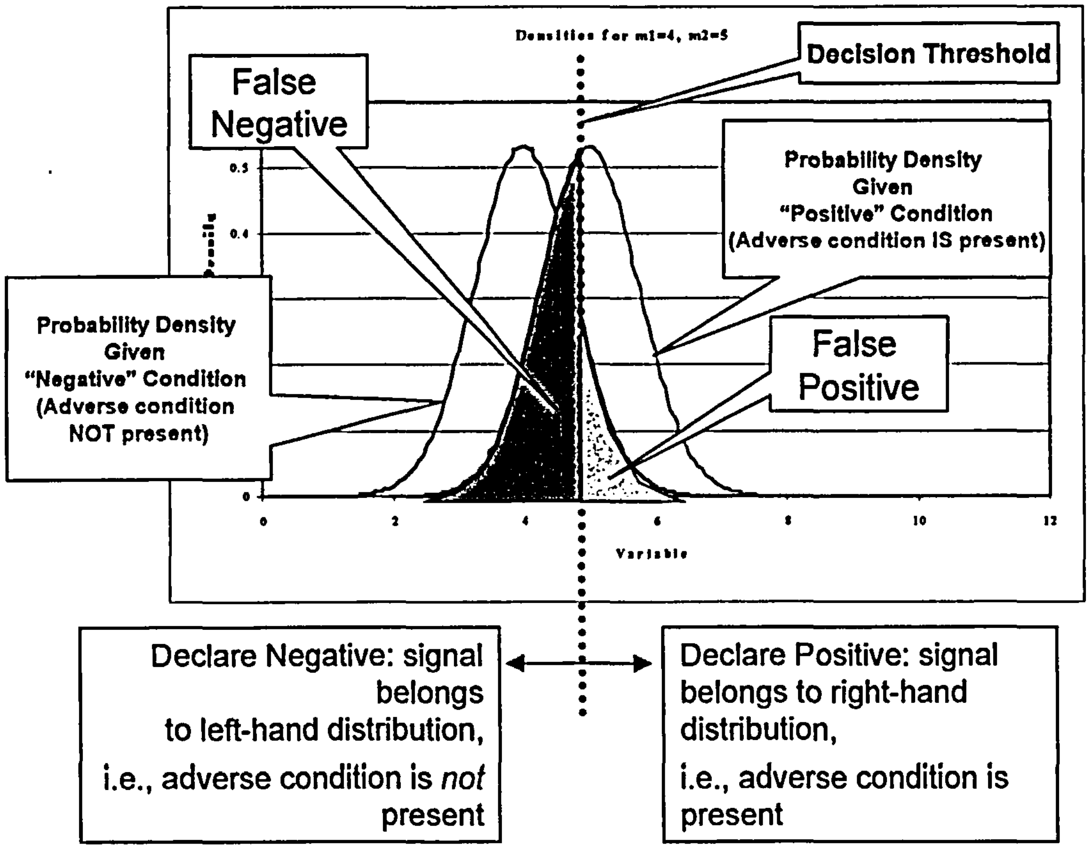

Example: Diagnosis of glaucoma based on measurement of intraocular pressure. High pressure correlates with glaucoma.
The horizontal axis variable is then "pressure," and the vertical axis is probability density. For the leftmost curve, the probability density is representative of the "normal" population; the rightmost curve (centered at higher pressure) is representative of the population having glaucoma. The question is where to set the decision threshold.

Figure II.3
Shape of ROC Describes The Overlap Between Two Uncertainty Distributions

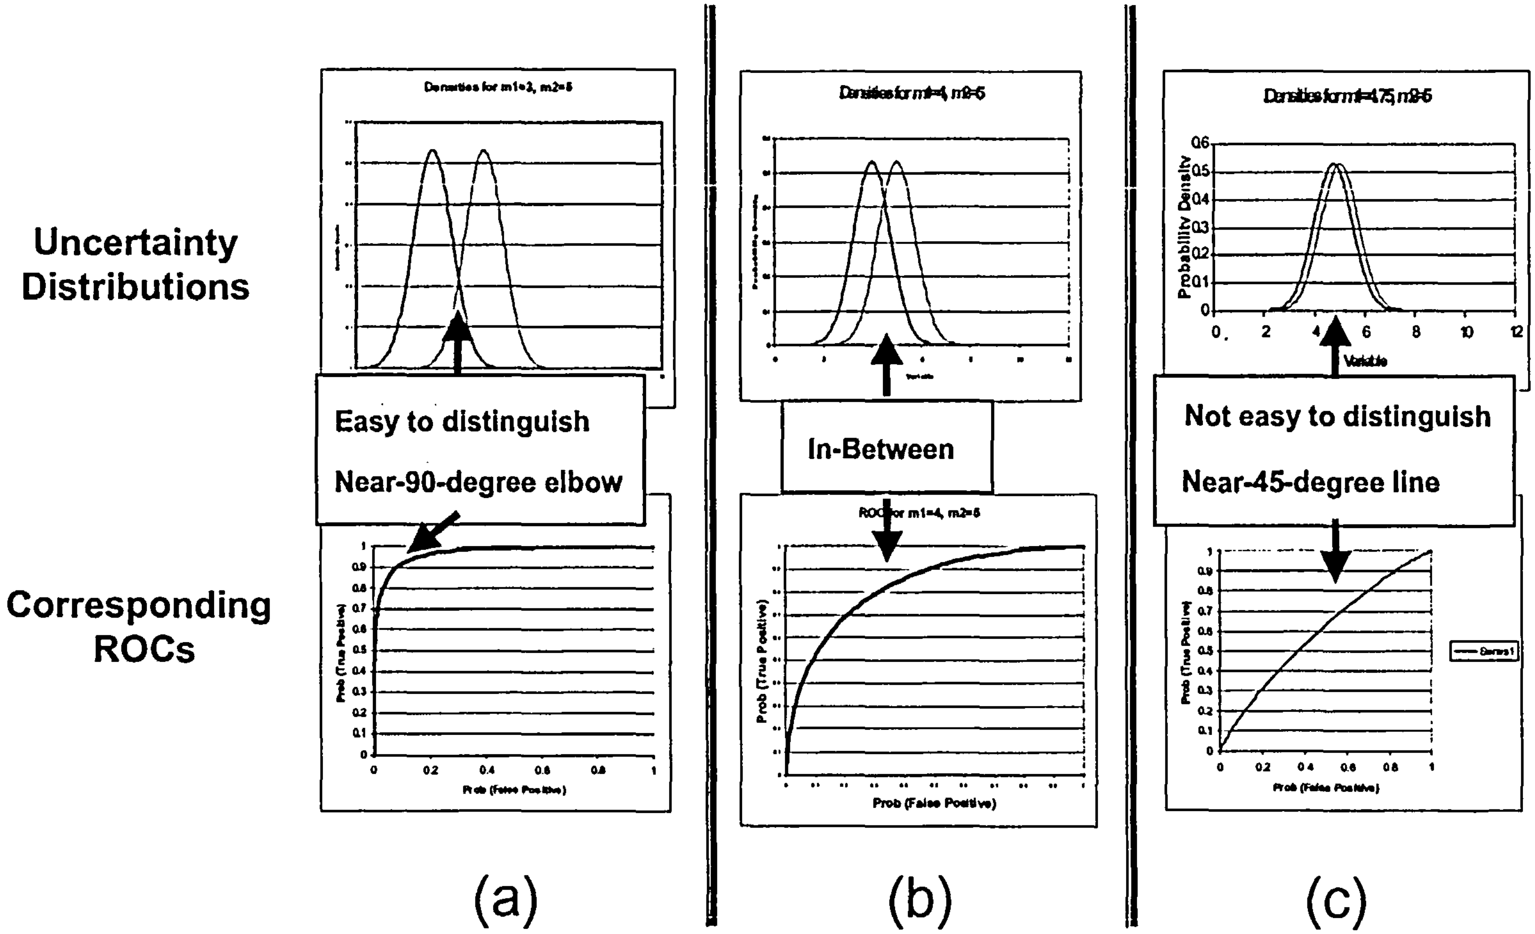

Figure II.4
ROC Description of Margin: Sheet 1

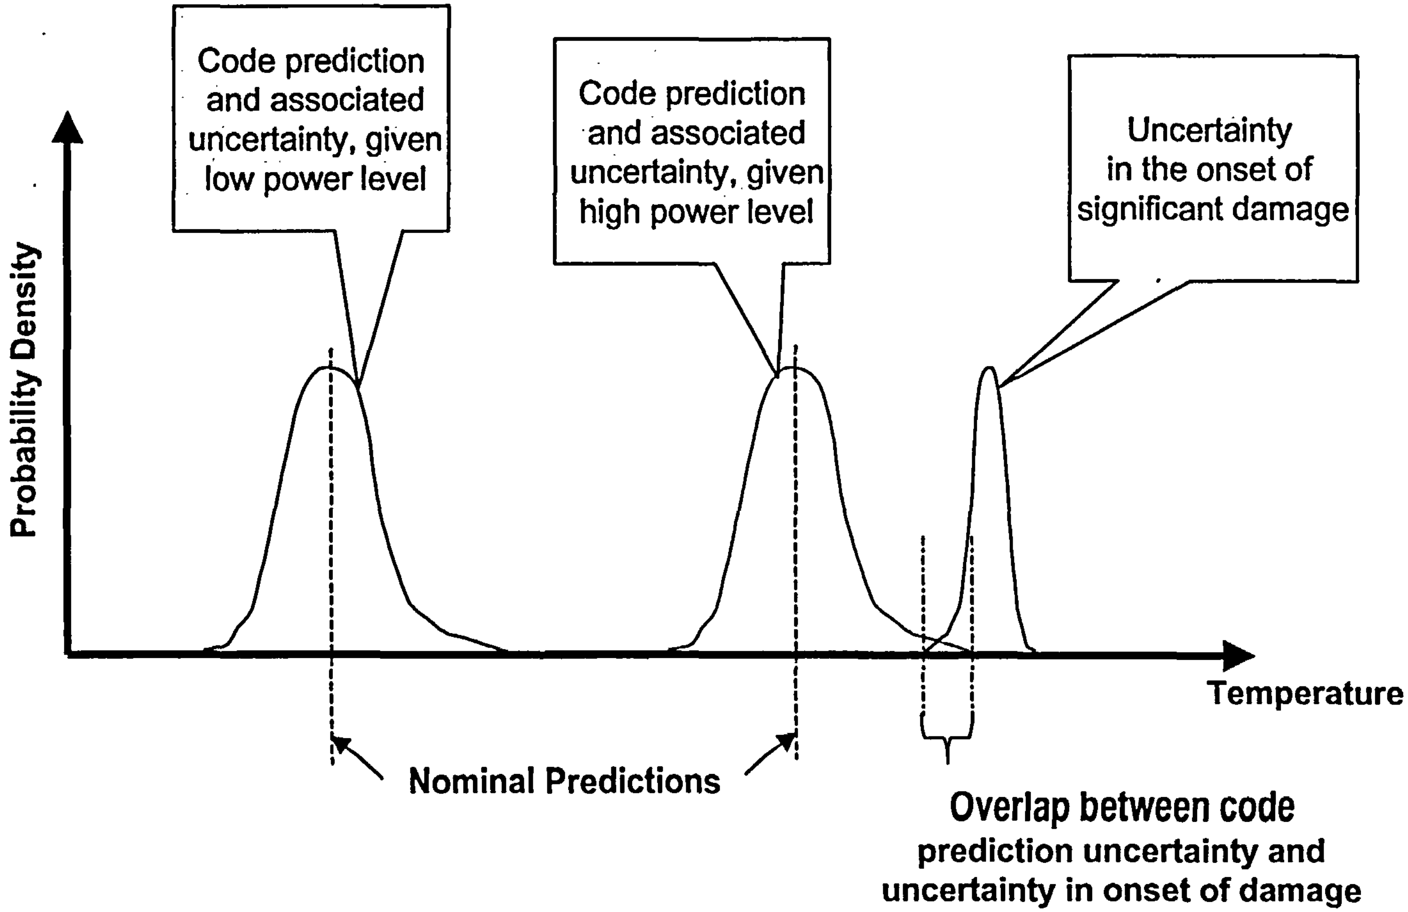

Figure II.4
ROC Description of Margin: Sheet 2

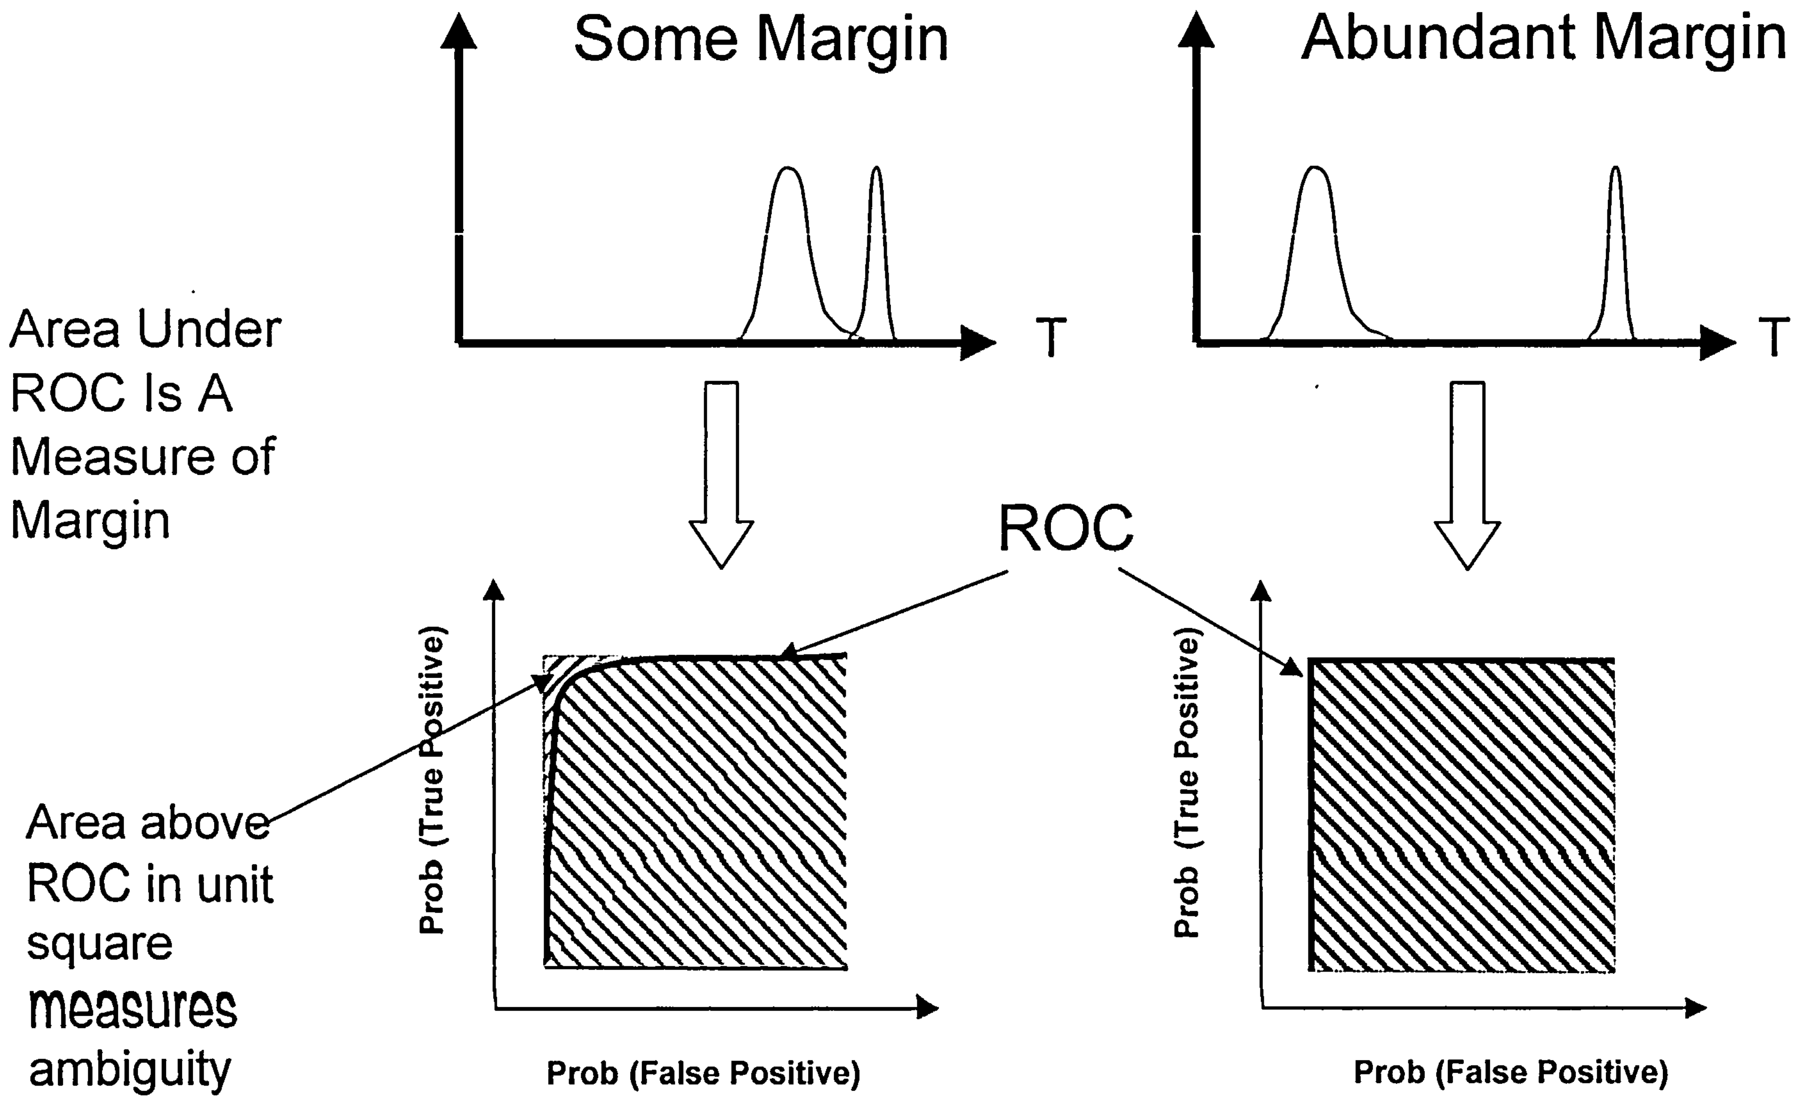

2. Determine the point on the ROC whose slope is given by

$$S_{optimal} = \frac{P(T-)}{P(T+)} \times \frac{B(T-\&D-) + C(T-\&D+)}{B(T+\&D+) + C(T+\&D-)}$$

where

P(T–) is the prior probability of a true negative,
P(T+) is the prior probability of a true positive,
B(T– &D–) is the benefit of assigning "negative" given that the condition is truly negative,
B(T+ &D+) is the benefit of assigning "positive" given that the condition is truly positive,
C(T– &D+) is the cost of assigning "positive" given that the condition is truly negative,
C(T+ &D–) is the cost of assigning "negative" given that the condition is truly positive.

According to Swets *et al.*, this point corresponds to the optimal threshold setting. This corresponds to the formula in Section II.2.1, although Swets and Van Trees use different notation.

The numerator contains the prior probability of true negative, the benefits of correctly declaring negative, and the costs of incorrectly declaring positive when the true condition is negative. Thus, this quantity has a certain intuitive significance: everything in the numerator is a reason to bias the decision "negative", and everything in the denominator is a reason to bias the decision "positive." A large number biases the decision negative by choosing a point well on the left of the ROC, accepting a low true positive probability in order to achieve a low false positive probability; a small number biases the decision positive by doing the opposite.

This formalism is very old, but does not contradict current utility-based decision theory; it is merely a special case of it, worked out in a relatively visual way.

Discussion

Some safety decisions can be cast in the form of a decision whether an adverse condition is present, and uncertainty affects some of these decisions. For example, perhaps symptom-based severe accident management guidelines could make beneficial use of this formalism (e.g., determine the threshold at which a particular function should be actuated, or emergency evacuation ordered).

An example of an *agency* decision of this kind could be review of accident analysis to decide whether the consequences of particular events are acceptable. Take the case of peak cladding temperature (PCT) reached in certain accident scenarios. Suppose that current methodological assumptions are not already in place, and that we need to determine whether certain outcomes are acceptable, given uncertainties in the calculation of PCT. Further suppose (for the sake of argument) that there is some uncertainty in the point at which real damage would be deemed to have occurred. Up to a point, this case corresponds to the ROC formalism discussed above. The horizontal axis is temperature. The leftmost uncertainty distribution corresponds to uncertainty in PCT for a particular scenario type; the rightmost uncertainty distribution would be the distribution of the lowest temperature corresponding to the onset of real damage. One could imagine coping with evaluations in which the two uncertainties overlapped, and trying to derive a threshold temperature, above which a calculated PCT would be deemed unacceptable.

In licensing review of the accident analysis of current-generation plants, agency practice is to conservatively avoid the kind of ambiguity that ROCs are designed to address. Where uncertainty is present, "margin" is required. For example, if operating power is reduced, the curve representing uncertainty in the PCT calculation will shift to the left; when power is reduced to where there is essentially no overlap between the calculated uncertainty and the damage curve, one can be confident that at that power level, the consequences of the postulated event are acceptable. By this process, a condition even more favorable than that shown in Figure II.3(a) is reached, rather than the condition shown in Figure II.3(c). In other words, rather than settling for an ambiguous evaluation, the agency requires that the initial conditions be controlled in such a way that the ambiguity in interpretation of the calculation does not arise.

It is easily seen that the magnitude of the uncertainties determines the power level at which the overlap goes to zero. Thus, uncertainty reduction has value, in that a higher power level can be tolerated if uncertainty is less. This is related to the "value of information" idea summarized earlier in this report, but this treatment comes at it rather differently, and shows visually how the economics of a partial reduction in uncertainty would play out.

Area Under the ROC

When the observed data do not help the decision process at all – that is, when the two pdf's overlap completely - the ROC is a diagonal line, having area 0.5. When the two pdf's have no overlap at all, the observed data unambiguously establish which hypothesis is correct, and the ROC is a right angle having unit area. It turns out that the area under the ROC is a useful measure of what Swets *et al.* call the "accuracy" of the decision rule. This is independent of the decision threshold: it is a measure of how the two pdf's overlap.

One use of such a measure is in comparing two decision rules. For example, it is interesting to compare a rule based on one set of data with another rule based on these data plus additional data. If the additional data help to resolve the two hypotheses, then the corresponding ROC will more nearly resemble a right angle, and have a higher value of A. This is related to, but distinct from, the value of information, which as presented earlier is evaluated for a given threshold; comparing A for different ROCs reflects more than just the VOI associated with a given threshold.

II.2.3 "Empirical" ROCs

Given the conditional probability density functions that represent the two hypotheses, one can construct the ROC as illustrated in the previous figures *a priori*. However, given sufficient data, one can also construct an ROC empirically. It may be of interest to do this in order to assess a given decision situation, e.g., to shed light on a discussion of the value of reducing a particular uncertainty, or it may be of interest to assess the area under the ROC in order to gain perspective on how good a job a particular decision rule is doing.

Constructing an empirical ROC requires that a number of cases be available, such that for each of a series of candidate thresholds (values of the observed variable), it is possible to assess a probability of false positive and a probability of true positive. According to Swets *et al.*, even fairly crude binning on the observed-variable axis can lead to a useful result. The ROC tool therefore suggests itself as a way of assessing diagnostic processes in application. For the case of non-destructive evaluation of metallurgical specimens to determine whether cracks are present, the attempt to construct an empirical ROC taking data from different examiners showed graphically that there is very substantial variability among the examiners, suggesting empirically that very different decision thresholds are being used.

II.2.4 Summary

The above discussion is only a brief summary of the formalism associated with this subject. Much attention has been paid to this area in the last half century, because the underlying concepts relate to many subjects, some of which are not all normally thought of as "decisions." Parameter estimation, for example, is related to these ideas.

Taking a step back from the mathematical machinery, it is important to note which variables emerge as important. They are the prior probabilities of the two hypotheses, the conditional probabilities of true and false diagnoses as a function of the threshold setting, and the costs of correct and incorrect diagnoses.

Note that in many contexts, we are not accustomed to quantifying the prior probabilities. For example, returning to the example used early in this section, while many analysts can quote a mean pump failure probability, few would be able to quote the prior probability of declining reliability performance. However, within the framework described above, this probability is needed to support the setting of thresholds for corrective action. Similarly, in many contexts, the cost functions are not discussed explicitly (although the hardware cost of proposed backfits is an exception to this).

II.3 ALTERNATIVE APPROACHES TO DECISION-MAKING UNDER UNCERTAINTY

Many texts in decision analysis present the view that expected utility is *the* basis for decision-making. The use of the term "expected" means that uncertainties are quantified, and based on these distributions, the expectation value of the utility function is quantified for each alternative.

This expected-utility-based approach is not universally embraced. In the interest of informing the reader's perspective, *but not to advocate these views,* the present section gives examples of alternative points of view. They differ in how uncertainty plays in the decision rule. All are aimed at situations in which the spectrum of consequences includes some very adverse outcomes, and where regulatory bodies must address the needs of diverse stakeholders.

It seems possible that some of the criticisms of the expected-utility-based approach that drive some of the alternatives mentioned below relate to flaws in implementation of the theory, rather than flaws in its underlying philosophical framework. It seems possible that some of the supposed benefits of these other ideas can be realized within a properly designed and implemented framework based on expected utility. The British model (discussed below) is an example of this.

II.3.1 "Decision-Making When Science Is Ambiguous"

In a recent Science article discussing policy towards mad cow disease, Anand (Ref. II-6) suggests that in dealing with high-consequence events whose probabilities are highly uncertain, the expected-utility-based approach (choosing the alternative that maximizes expected utility) has "fatal" difficulties [Anand]. He argues instead for "decisionmaking under ambiguity." He discusses several different decision rules that are said to operate without consideration of actual probability:

1. pick the alternative having the least onerous of the worst possible outcomes, based on examining all states of nature independent of their assessed probabilities;
2. pick the alternative having the best possible outcome, based on examining all states of nature independent of their assessed probabilities;
3. assign equal probabilities to all states of nature, and go with the alternative having the best "expected" outcome on that basis.

An important dimension of Anand's discussion is the involvement of a broad range of stakeholders in a politicized environment. His discussion leads him to the conclusion that when uncertainties in probability are large, preferences become even more important, and assessment of consequences ought to dominate the thought process. This is reflected in ideas cited below, originating in some cases from philosophical or ethical points of view that are argued by some people not to be captured in utility functions.

It is doubtful that the concept of "worst possible outcome" really avoids probability. Looking more closely at "worst case" scenarios typically reveals embedded assumptions about probability.

II.3.2 The Precautionary Principle

According to Davidson (Ref. II-7),

> In the political liberalism of John Rawls, decision-making under conditions of uncertainty plays a key role. To design a just society, Rawls advocates application of a so-called maximin principle: "The maximin rule tells us to rank alternatives by their worst possible outcomes: we are to adopt the alternative the worst outcome of which is superior to the worst outcomes of the others" (Rawls, 1972, 1974; Shrader-Frechette, 1991). Utilitarians disagree here and advocate a *Bayesian strategy,* involving optimisation of expected aggregate social utility on the basis of subjective probabilities (Harsanyi, 1975, 1977, 1978). Finally, there

is the 'precautionary principle'. Although many definitions of the precautionary principle exist, it boils down roughly to the principle *in dubio pro natura.*

There is a considerable literature on the "Precautionary Principle." The essence of the principle appears to be that it is appropriate to take conservative action in the presence of uncertainty, rather than requiring proof of the existence of a problem before action is taken.

Based on some web searching, it appears that much of the discussion of this principle is carried out by environmental groups in the context of balancing environmental interests against economic ones. However, moderated forms of the principle seem established in some quarters of Europe and the UK. The UK's Health and Safety Executive devotes significant attention to a form of the principle (Ref. II-8); their discussion embeds the principle in a much less absolutist conceptual frame than do some of the environmental groups. Unlike the USNRC, the British apply a test of "gross disproportion" to determine whether a backfit needs to be made.[2] (That is, a proposed benefit must be implemented unless costs are in "gross disproportion" to benefits, a somewhat precautionary approach.) This moderate approach seems to be reconcilable with expected-utility approaches.

II.3.3 Integrated Decision-Making

In utility-based decision-making, "risk analysis" is the body of work done to compare alternatives, taking proper account of uncertainties. The alternatives cited above step outside of this paradigm when uncertainties are large and stakes are high.

In a sense, the decision-making paradigm articulated in RG 1.174 more nearly resembles these alternatives than it does mainstream utility theory. As explained in the following excerpts from Regulatory Guide 1.174, risk analysis is but one factor (item 4 in the list below) informing an integrated decision (Ref. II-10).

> In implementing risk-informed decisionmaking, LB [licensing basis] changes are expected to meet a set of key principles. Some of these principles are written in terms typically used in traditional engineering decisions (e.g., defense in depth). While written in these terms, it should be understood that risk analysis techniques can be, and are encouraged to be, used to help ensure and show that these principles are met. These principles are:

---

[2] For example (Ref. II-9):

> The regulation of safety in the United Kingdom is based upon the principle that risks must be reduced to a level which is "As Low As Reasonably Practicable" (the so called ALARP principle). The meaning of "reasonably practicable" is well established in English case law:
>
> > *"Reasonably practicable" is a narrower term than "physically possible" and seems to me to imply that a computation must be made by the owner in which the quantum of risk is placed on one scale and the sacrifice involved in the measures necessary for averting the risk (whether in money, time or trouble) is placed in the other, and that, if it be shown that there is a gross disproportion between them -- the risk being insignificant in relation to the sacrifice -- the defendants discharge the onus on them. "(Judge Asquith, Edwards v. National Coal Board, All England Law Reports Vol.* 1, p. 747 (1949)).
>
> Thus the ALARP principle allows cost to be taken into account in determining how far to go in the pursuit of safety, so that if a risk reduction measure involves "grossly disproportionate" cost, it is not "reasonably practicable". This principle was adopted in the Health and Safety at Work Act (1974), and is the basis of the approach adopted by the Health and Safety Executive in their regulation of the major hazard industries, such as the offshore oil and gas industry. It has also been adopted by BNM in its application of the Railways (Safety Case) Regulations (1994).

1. The proposed change meets the current regulations unless it is explicitly related to a requested exemption or rule change, i.e., a "specific exemption" under 10 CFR 50.12 or a "petition for rulemaking" under 10 CFR 2.802.
2. The proposed change is consistent with the defense-in-depth philosophy.
3. The proposed change maintains sufficient safety margins.
4. When proposed changes result in an increase in core damage frequency or risk, the increases should be small and consistent with the intent of the Commission's Safety Goal Policy Statement (Ref. 5 [in the original]).
5. The impact of the proposed change should be monitored using performance measurement strategies.

...

Consistency with the defense-in-depth philosophy is maintained if:

- A reasonable balance is preserved among prevention of core damage, prevention of containment failure, and consequence mitigation.
- Over-reliance on programmatic activities to compensate for weaknesses in plant design is avoided.
- System redundancy, independence, and diversity are preserved commensurate with the expected frequency, consequences of challenges to the system, and uncertainties (e.g., no risk outliers).
- Defenses against potential common cause failures are preserved, and the potential for the introduction of new common cause failure mechanisms is assessed.
- Independence of barriers is not degraded.
- Defenses against human errors are preserved.
- The intent of the General Design Criteria in Appendix A to 10 CFR Part 50 is maintained.

In a letter to the Commission (Ref. II-10), the Advisory Committee on Reactor Safeguards (ACRS) implicitly remarked the essential difference in practice between integrated decisionmaking and basing decisions on expected utility:

> We believe that two different perceptions of defense in depth are prominent. In one view (the "structuralist" view as described in the attached paper), defense in depth is considered to be the application of multiple and redundant measures to identify, prevent, or mitigate accidents to such a degree that the design meets the safety objectives. This is the general view taken by the plant designers. The other view (the "rationalist"), sees the proper role of defense in depth in a risk-informed regulatory scheme as compensation for inadequacies, incompleteness, and omissions of risk analyses. We choose here to refer to the inadequacies, incompleteness, and omissions collectively as uncertainties. Defense-in-depth measures are those that are applied to the design or operation of a plant in order to reduce the uncertainties in the determination of the overall regulatory objectives to acceptable levels. Ideally then, there would be an inverse correlation between the uncertainty in the results of risk assessments and the extent to which defense in depth is applied. For those uncertainties that can be directly evaluated, this inverse correlation between defense in depth and the uncertainty should be manifest in a sophisticated probabilistic risk assessment (PRA) uncertainty analysis.

In essence, invoking defense in depth is a precautionary response to uncertainty in risk analysis.

II.3.4 Analytic Hierarchy Process (AHP)

The AHP (Ref. II-11) is discussed in Part III in connection with objectives hierarchies. It deserves mention here because it is an approach to decision-making that leaves certain things, including particular uncertainties, implicit.

II.3.5 Summary: Different Approaches to Accommodating Uncertainty in Decision-Making

In some quarters, there is significant distrust of the process of treating uncertainty by propagating them through risk models and using utility functions. Some writers also argue from ethical considerations that environmental regulatory decisions need to treat large uncertainties in a different way. These factors drive the formulation of decision rules that differ from utility-function-based decision-making. Arguably, NRC's own "integrated decision-making" process is one instance of this. As noted by ACRS, there are problems with some of these alternative approaches. But improving on such decision rules means modeling, including quantification of uncertainties, of a sort that apparently cannot be accomplished to universal satisfaction. Some of the resistance to approaches based on expected utility relates to the real difficulty of characterizing uncertainty. Section II.4 below briefly discusses this difficulty.

It is not the purpose of this report to criticize integrated decision-making, or insist that it is a departure from the mainstream. It is only the purpose of the present section to note that the elements of integrated decision-making somewhat resemble the elements of the "alternative" approaches, in that when the decision is made, the quantitative risk metrics are trumped by higher external principles (defense in depth). At least some of the literature seems to regard these alternatives as "decision rules" that are different from "maximize utility." It would be interesting to develop a utility function that captures agency thinking regarding these external principles.

II.4 ASSESSING UNCERTAINTY

II.4.1 Code Scaling, Applicability, and Uncertainty Evaluation Methodology (CSAU)

The Code Scaling, Applicability, and Uncertainty Evaluation Methodology (CSAU) is a systematic procedure to quantify the uncertainty in calculated safety parameters from computer models. This methodology has many of the elements and illustrates some of the steps of a decision-making process even though it is intended for use with evaluating results from a computer model. The discussion in this section is based on Ref. II-12.

Figure II.5 presents a diagram illustrating the CSAU methodology (taken from Figure 1 of Ref. II-12). The methodology is organized into three elements as shown on the figure: 1) Requirements and Capabilities, 2) Assessment and Ranging of Parameters, and 3) Sensitivity and Uncertainty Analysis. There are 14 steps that are incorporated into the three elements indicated by number on the figure. After specifying a scenario and nuclear power plant for analysis (Steps 1 and 2), a Phenomena Identification and Ranking Table (PIRT) is generated (Step 3). A PIRT is developed for many thermal hydraulic analyses to identify important phenomena for subsequent evaluation. In the CSAU methodology, the PIRT is used to guide the uncertainty quantification. NUREG/CR-5249 identifies several techniques for accomplishing the ranking in a PIRT, including expert opinion, subjective decision-making methods, and scoping calculations.

Prioritization, as it is discussed in Part I, is a key part of what goes on in PIRT ("Ranking"). Therefore, many of the decision-making processes described in this report can be used in developing a PIRT.

Selection of a documented computer code to be used for the analysis constitutes Steps 5 and 6. Given the selection of the computer code, the applicability of the code to the transient that is to be analyzed is evaluated (Step 6). This evaluation includes a review of the code capability. Note that the evaluation of the code capability is performed in conjunction with the phenomena identified in the PIRT. A goal of this evaluation is to determine whether the code is suitable to be used for analyzing the transient of interest. These first six steps are included in the Requirements and Capabilities Element.

Steps 7 to 10 are part of the Assessment and Ranging of Parameters Element. These steps are needed to quantify the effects of the individual contributors to uncertainty. These contributors include code limitations, scaling effects embedded in experimental data and code, and the initial state of the reactor. An assessment matrix is developed that identifies tests that best address the important phenomena identified in the PIRT in Step 7. In Step 8, a plant model is developed using a nodalization that is sufficiently fine to capture important phenomena and plant design characteristics without being impractical from a computer execution time perspective. Code and experimental accuracy are evaluated in Step 9 and the effect of scale is determined in Step 10.

The remaining four steps are part of the Sensitivity and Uncertainty Analysis Element. Step 11 is the determination of the effect of the reactor input parameters and operating state. Uncertainty in fuel geometry and tolerances are an example of input parameters that would be considered. In Step 12, sensitivity calculations are performed to evaluate the impact of key thermal hydraulic results on the parameters identified during Step 11. Step 13 is the determination of combined bias and uncertainty. In this step, a formal uncertainty analysis is performed using Monte Carlo analysis. Determination of the total uncertainty is performed in Step 14.

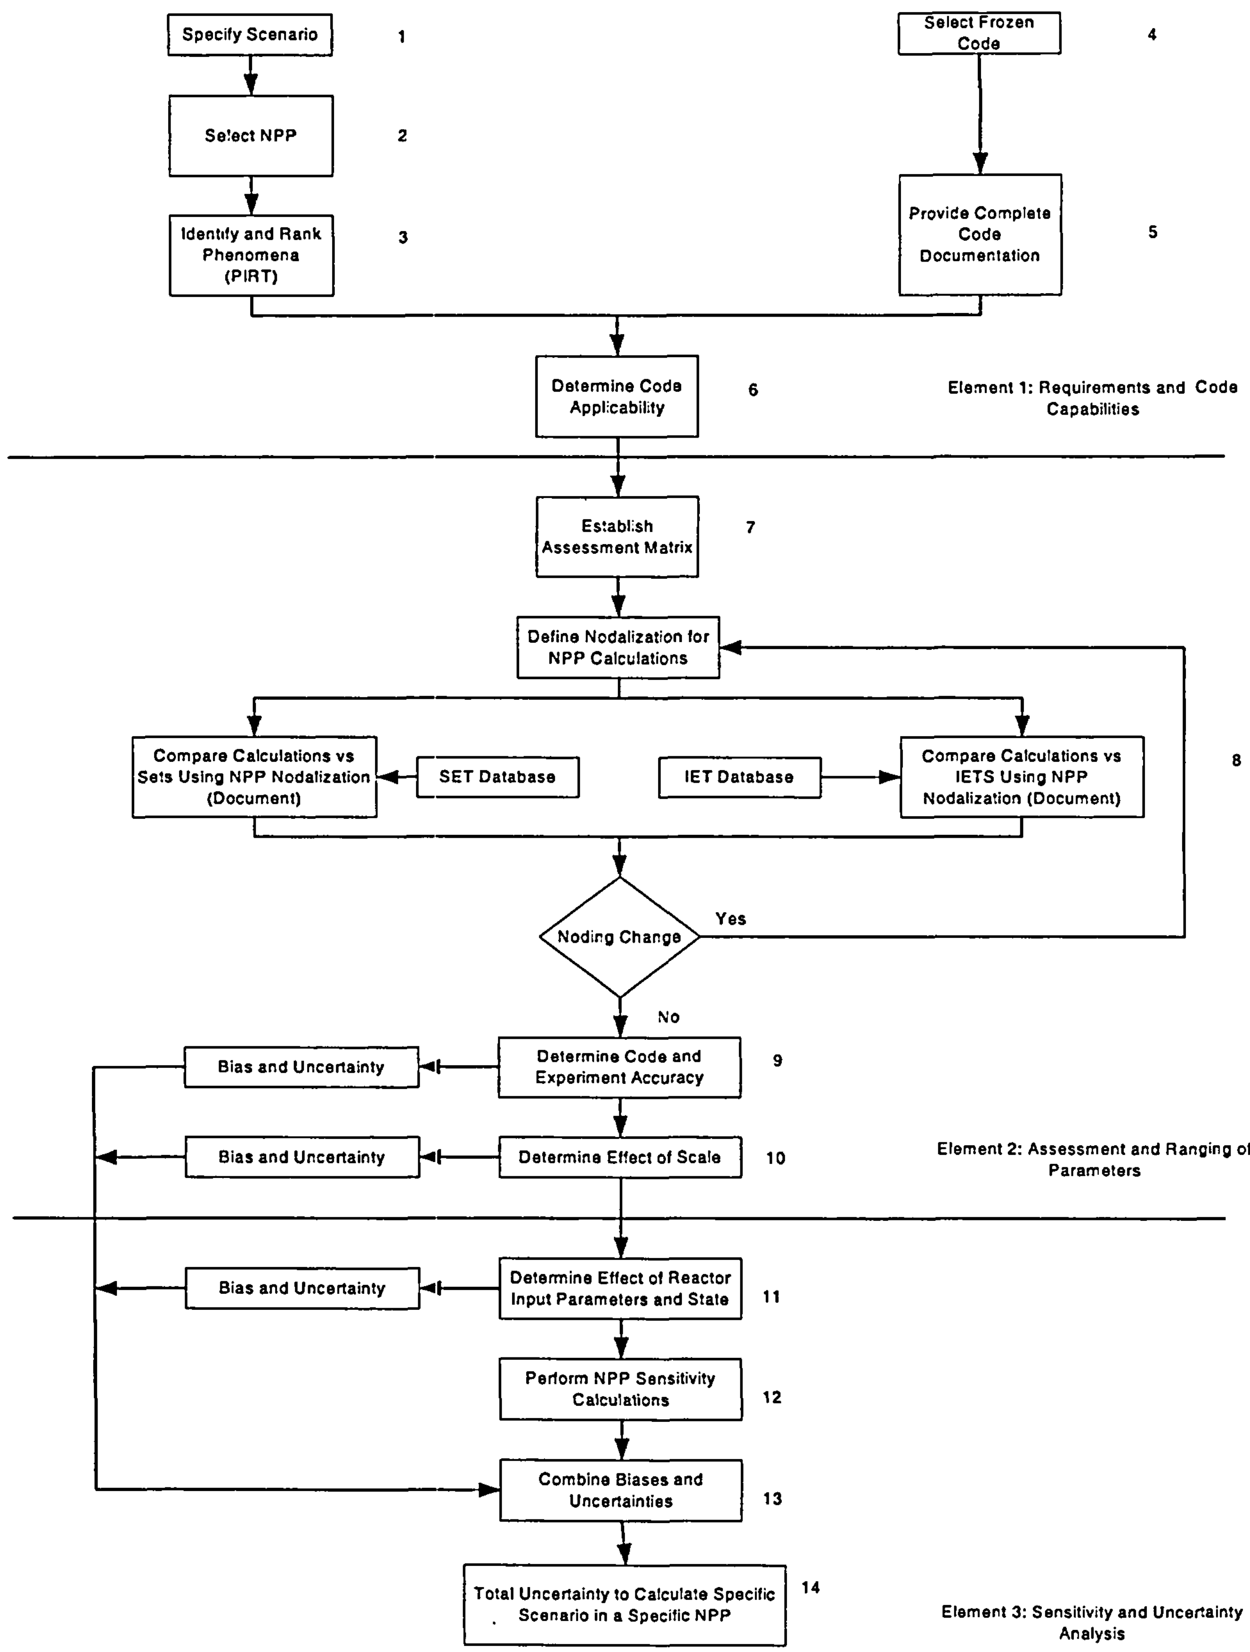

Figure II.5 Illustration of the CSAU Methodology

The CSAU methodology in NUREG/CR-5249 is presented from the perspective of thermal hydraulic code applications. However, the CSAU methodology has many elements in common with decision-theoretic approaches. A comparison of the PrOACT process (refer to Part I of this report) to CSAU, for example, shows many common elements (ignoring the differences in terminology). The approach for developing a PIRT would utilize the approach outlined in the first five steps in the PrOACT process. The use of statistical methods to treat uncertainty is a key element in both decision-making and in CSAU. The CSAU process suggests the approach of sampling from a response surface developed from an experimental design to assess uncertainty. It would be an interesting exercise to develop a PIRT using a belief network or an influence diagram to gain experience and to assess the insights that may result.

II.4.2 Other Methods for Developing Uncertainty Distributions

In the last half century, machinery for dealing with the implications of uncertainty distributions has become highly developed. However, methods for developing honest uncertainty distributions in the first place (at least, methods that are both generally understood and accepted) appear still to be lacking. The literature of this subject is too vast to permit adequately summarizing it in this report.

There is a very substantial literature discussing Bayesian methods for development of uncertainty distributions. Much of that discussion treats cases in which evidence relating to a given parameter is unambiguously applicable, but limited. A typical example is the problem of assessing the reliability of components in a particular group whose essential equivalence is not in any doubt, given a limited set of operating reliability data. Given a prior distribution, it is clear how to update it in light of current data. Where to get the prior is still a research topic.

A more problematic case, and one that can be expected to arise for advanced reactors, is the case where not only are the relevant data limited, but also their applicability is imperfect or controversial. In cases where data are extremely limited, one technique used is the elicitation of expert opinions regarding the value of the unknown parameter, followed by the integration of these opinions into a single uncertainty distribution. This is somewhat like treating the expert opinions as equivalent to experimental measurements of the parameters. There is a very substantial literature on this topic as well. Such approaches are controversial in some quarters both because of certain biases that are known to affect the way in which individuals process their own experiences, and because of controversies regarding how the aggregation should be performed.

An idea due to Kaplan seems worth trying on issues of this kind (Ref. II-13). Kaplan's idea is that instead of eliciting and aggregating experts' opinions regarding the unknown parameter, one ought instead to elicit the evidence informing each expert's opinion, aggregate this body of evidence, and then, using Bayesian methods, derive the desired uncertainty distribution from this body of evidence. Instead of mentally quantifying the parameter individually, the experts' job is to evaluate the significance of the available evidence; instead of a distribution reflecting the aggregate of the individuals' quantifications, the process yields a distribution reflecting the evidence as weighed by the consensus of the panel. Although this approach uses experts, it focuses more directly on evidence, and Kaplan calls it the "expert information" approach to distinguish it from the "expert opinion" approach.

As stated earlier, this approach suggested itself originally in the context of issues for which data are limited. But it seems also to have the potential to deal with issues in which evidence is not only sparse, but of uncertain applicability. Qualification tests suggest themselves as instances of this. In some cases, qualification tests can replicate the situation in which performance needs to be assessed; in other cases, the replication may be imperfect, for reasons of scale or of system topology. To what extent is experience with Design A applicable to Design B, if B is similar but not identical to A? Kaplan presents an example similar to this, in which the problem is to characterize a performance parameter for a new design that has been adapted from an earlier model for which there is substantial operating experience. It would appear that formally, this approach could deal not only with test data, but also with the problem of integrating test data with information derived from simulations.

Belief networks suggest themselves as a possible tool for carrying out Kaplan's approach, perhaps in combination with the AHP. Belief networks and the AHP are discussed in the Appendix. We have not yet developed an example illustrating this potential.

II.5 SUMMARY

This part of the report has described widely-practiced formal methods for treating uncertainty in order to minimize the risk associated with selecting a decision alternative. Here, "risk" associated with a given decision alternative is to be understood in a general sense, referring to the expected consequences of various types, expressed in common terms (e.g., monetized). The discussion showed how optimizing on a risk metric inherently involves the use of prior probability information, as well as information regarding the benefits and costs of decision alternatives. Decision thresholds were shown to depend on an interaction between the underlying uncertainties and the costs and benefits. The discussion was carried out for cases in which only two alternatives are possible, but the main ideas generalize to more complicated situations.

"Minimizing expected risk (in the general sense)" is the paradigm most widely discussed in connection with formal decision-making. It takes for granted the reduction of the consequences, and the quantification of their uncertainty, into a description that is acceptable to all involved stakeholders. For several reasons, this paradigm is not universally accepted. When uncertainties are great and potential consequences are large, some do not follow the utility-based paradigm, preferring simply to hard-wire conservatism into the decision-making process in one way or another. For example, some advocate selecting the alternative having the least possible adverse consequences, as if this could be assessed independently of probability. Another approach is to supplement utility-type metrics in decision-making with proxy attributes (refer to Part III of this report) that amount to litmus tests on the potential vulnerability of decision alternatives to unidentified or underestimated sources of uncertainty. One could argue that thee approaches merely redefine "risk," and evaluate this new "risk" in essentially the usual way. However, these two approaches are formulated in terms that explicitly go beyond normal uncertainty analysis; indeed, they essentially decline to accept uncertainty analysis at face value. Distinguishing them from utility-based decision-making therefore seems warranted. These ideas are not presented here in order to suggest that they are practical alternative approaches, but rather to make the reader aware that treatment of uncertainty remains, for many, an unsettled question.

Although the machinery for reflecting uncertainty distributions in decision models is highly developed, the machinery for obtaining the uncertainty distributions in the first place is much less developed. There is an enormous literature on expert elicitation, which is largely beyond the scope of the present report. At this point, expert elicitation has not been reduced to a simple, widely-accepted approach. A related idea due to Kaplan ("expert information") has intuitive appeal, but still calls for assimilation of disparate types of evidence into a Bayesian evaluaion, and this process of assimilation has also not been reduced to a simple, widely-accepted approach. NRC's own "Code Scaling, Applicability, and Uncertainty" (CSAU) methodology represents one approach to a very focused type of evaluation (characterizing uncertainties in specific modeling contexts).

Part I showed how research priorities can be informed by consideration of the value of reducing uncertainty. Clearly, in order for this to work, the uncertainties need to be assessed properly in the first place. Assessment of uncertainty is an area in which significant gains remain to be made.

II.6 REFERENCES

II.1. U.S. Nuclear Regulatory Commission, An approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis, Regulatory Guide 1.174 (USNRC, 1998).

II.2. G. Apostolakis, "The Distinction Between Aleatory and Epistemic Uncertainties Is Important: An Example From The Inclusion of Aging Effects Into PSA," Proceedings of PSA '99, International Topical Meeting on Probabilistic Safety Assessment (American Nuclear Society, 1999).

II.3. H. L. Van Trees, *Detection, Estimation, and Modulation Theory*, John Wiley & Sons, Inc. (1968).

II.4. D. L. Kelly, "Bayesian Hypothesis Testing and the Maintenance Rule," pp. 1245-1250, Proceedings of the 4th International Conference on Probabilistic Safety Assessment and Management, 13-18 September 1998, New York City, USA [sic] (Springer, 1998).

II.5. J. A. Swets et al., "Psychological Science Can Improve Diagnostic Decisions," pp. 1-26, Psychological Science in the Public Interest, Vol. 1, No. 1, May 2000.

II.6. Anand, P., Decision Making When Science Is Ambiguous, Science Magazine, Volume 295, March 8, 2002.

II.7. M. D. Davidson, "Treating future generations as equals with respect to risk and uncertainty," Risk and Uncertainty in Environmental and Resource Economics, June 5-7, 2002, Wagenlingen University, The Netherlands.

II.8. Health and Safety Executive, Reducing Risks, Protecting People – HSE's Decision-Making Process, Crown Copyright 2001, http://www.hse.gov.uk/dst/r2p2.pdf.

II.9. T. A. W. Geyer *et al.*, "Improving Passenger Safety at Platforms," 1996 Anual Meeting of the Society for Risk Analysis – Europe, Summary of Meeting Paper (Society for Risk Analysis – Europe, 1996).

II.10. Letter, Dana A. Powers (ACRS) to The Honorable Shirley A. Jackson (USNRC Chairman), "The Role of Defense In Depth in a Risk-Informed Regulatory System," May 19, 1999.

II.11. T. L. Saaty and L. G. Vargas, *The Logic of Priorities* (Kluwer-Nijhoff Publishing, 1982).

II.12. Idaho National Engineering Laboratory, Quantifying Reactor Safety Margins, NUREG/CR-5249, October 1989.

II.13. S. Kaplan, " 'Expert Information' Versus 'Expert Opinions:' Another Approach to the Problem of Eliciting / Combining / Using Expert Opinion in PRA," Proceedings of PSA '89, International Topical Meeting on Probability, Reliability, and Safety Assessment (American Nuclear Society, 1989).

Part III: Formulation and Application of the Objectives Hierarchy

III.1 FORMULATION OF THE OBJECTIVES HIERARCHY

III.1.1 Overview

This part of the report discusses the formulation of alternatives, selection of performance measures, and the setting of performance goals. All of these processes are aided significantly by the development of an objectives hierarchy.

The objectives hierarchy is treated extensively in mainstream discussions of decision theory, because selection of a preferred alternative should be carried out in light of specific objectives. For example, choosing among research programs depends on what we are trying to accomplish. However, well before comparison of alternative proposals is undertaken, developing an objectives hierarchy can improve the very formulation of alternative proposals. Clarifying objectives not only helps to choose among pre-defined research programs, it supports formulating better programs in the first place.

At a more detailed level, the application of the objectives hierarchy to selection of performance measures has definite potential for application in advanced reactor work, in selection of metrics by which new designs are to be evaluated. Metrics used for LWRs include such things as peak clad temperature in selected scenarios, departure from nucleate boiling ratio (DNBR), clad oxidation, core damage frequency, large early release frequency, etc. For non-LWR designs, it is necessary to reconsider these metrics. Recently, application of the objectives hierarchy to selection of performance measures has been discussed in connection with the development of so-called "performance-based" approaches, and elements of that discussion will be adapted for present purposes.

The subject of allocation will be discussed briefly. By "allocation," we mean essentially parsing out high-level functional performance goals over lower levels, in order to support implementation guidance at those lower levels to promote satisfaction of the higher-level goals. General decision-theoretic treatments do not necessarily address this particular topic, but it is a particular instance of the "optimization" motive that drives mainstream decision analysis, and it is implicit in the application of decision methods to reactor regulation.

Finally, an important element of objectives to be faced in research prioritization relates to the dynamics of the research program. Priorities change as time passes: new knowledge causes shifts in perceived uncertainties, implying changes in research priorities; programmatic changes can re-order technical priorities in the research program. Recent work responding to these considerations is discussed.

III.1.2 Development of an Objectives Hierarchy and Selection of Performance Measures

This subsection briefly discusses development of an objectives hierarchy. In recent years, use of this kind of construct has become increasingly common within the agency. For example, development of the revised Reactor Oversight Process (SECY 99-007) (Ref. III-1, discussed briefly below) made use of an objectives hierarchy in a way that helped to foster completeness in the set of performance areas considered for attention as part of oversight. Similarly, the formulation of more general approaches to performance-based regulation (Refs. III-2 through III-4) has relied fundamentally on the idea of objectives hierarchies in deciding what elements of current safety cases should receive attention. These ideas are directly applicable to review of advanced reactor designs, and to the formulation of regulatory approaches to licensing and oversight.

The ideas would be applicable even if the basic regulatory criteria had already been formulated (as they have been for current plants); but they are especially applicable when the regulatory criteria have not yet been formulated. One can consider more systematically the selection of performance measures: whether natural measures are available, or if not, what proxy measures to work with.

III.1.3 The Objectives Hierarchy

An objectives hierarchy is a diagram representing the relationships and dependencies between goals, top-level fundamental objectives, lower-level fundamental objectives, and means objectives. Fundamental objectives are ends in themselves; means objectives are things that are desirable because they promote satisfaction of the fundamental objectives (Refs. III-5, III-6). An example of a goal is "protection of the health and safety of the public;" an example of a fundamental objective is "protection of the public from excessive radiological exposures;" an example of a means objective is "reliability of safety systems." These are examples from the structure of the Reactor Oversight Process, which makes use of an objectives hierarchy.

Figures III.1-III.3, based in part on the Reactor Oversight Process (ROP) (Ref. III-1), illustrate concepts of an objectives hierarchy. Figure III.1 shows goals, fundamental objectives, and means objectives. Figure III.2 presents more detail of the ROP. The cornerstone areas identified on Figure III.2 are intended to be a complete set of key performance areas affecting safety. The key attributes identified within each cornerstone are likewise intended to be a complete set. Completeness is one of the reasons to pursue such a systematic development.

On Figure III.2, consideration of the different cornerstone areas also illustrates how the implicit underlying allocation of performance addresses defense-in-depth at a high level. Balance between prevention and mitigation is shown by the presence of cornerstones addressing initiating events, mitigating systems, and emergency preparedness; the additional consideration of barrier integrity further reinforces defense-in-depth.

Analogously to logic tree development, each level of the objectives hierarchy is derived from the level above by decomposing each node into constituent elements. Each means objective relates to an objective above it on the hierarchy, in that it answers the question "how is the higher-level objective to be accomplished?" (Example: How will Safety function X be accomplished? Answer: By reliable function of systems A, B, and C. The reliable function of A, B, and C are means objectives supporting Safety Function X.) In fact, a system reliability model developed hierarchically and expressed in "success space" is essentially a partial objectives hierarchy. It is "partial" because it addresses only safety performance, and because even within "safety," a logic model does not usually address cross-cutting programmatic issues of the sort that appear at the bottom of Figure III.3.

Figure III.3 shows levels of a hierarchy applicable to many issues involving safety assessments. A complete and explicit development of all hardware and programs involved in accomplishing high-level safety functions would be a significant undertaking. However, the concepts can apply even if the development exists only in the abstract or in a qualitative way. The value of this construct will be clearer in light of the discussion below.

Figure III.1
Sheet 1 of Objectives Hierarchy

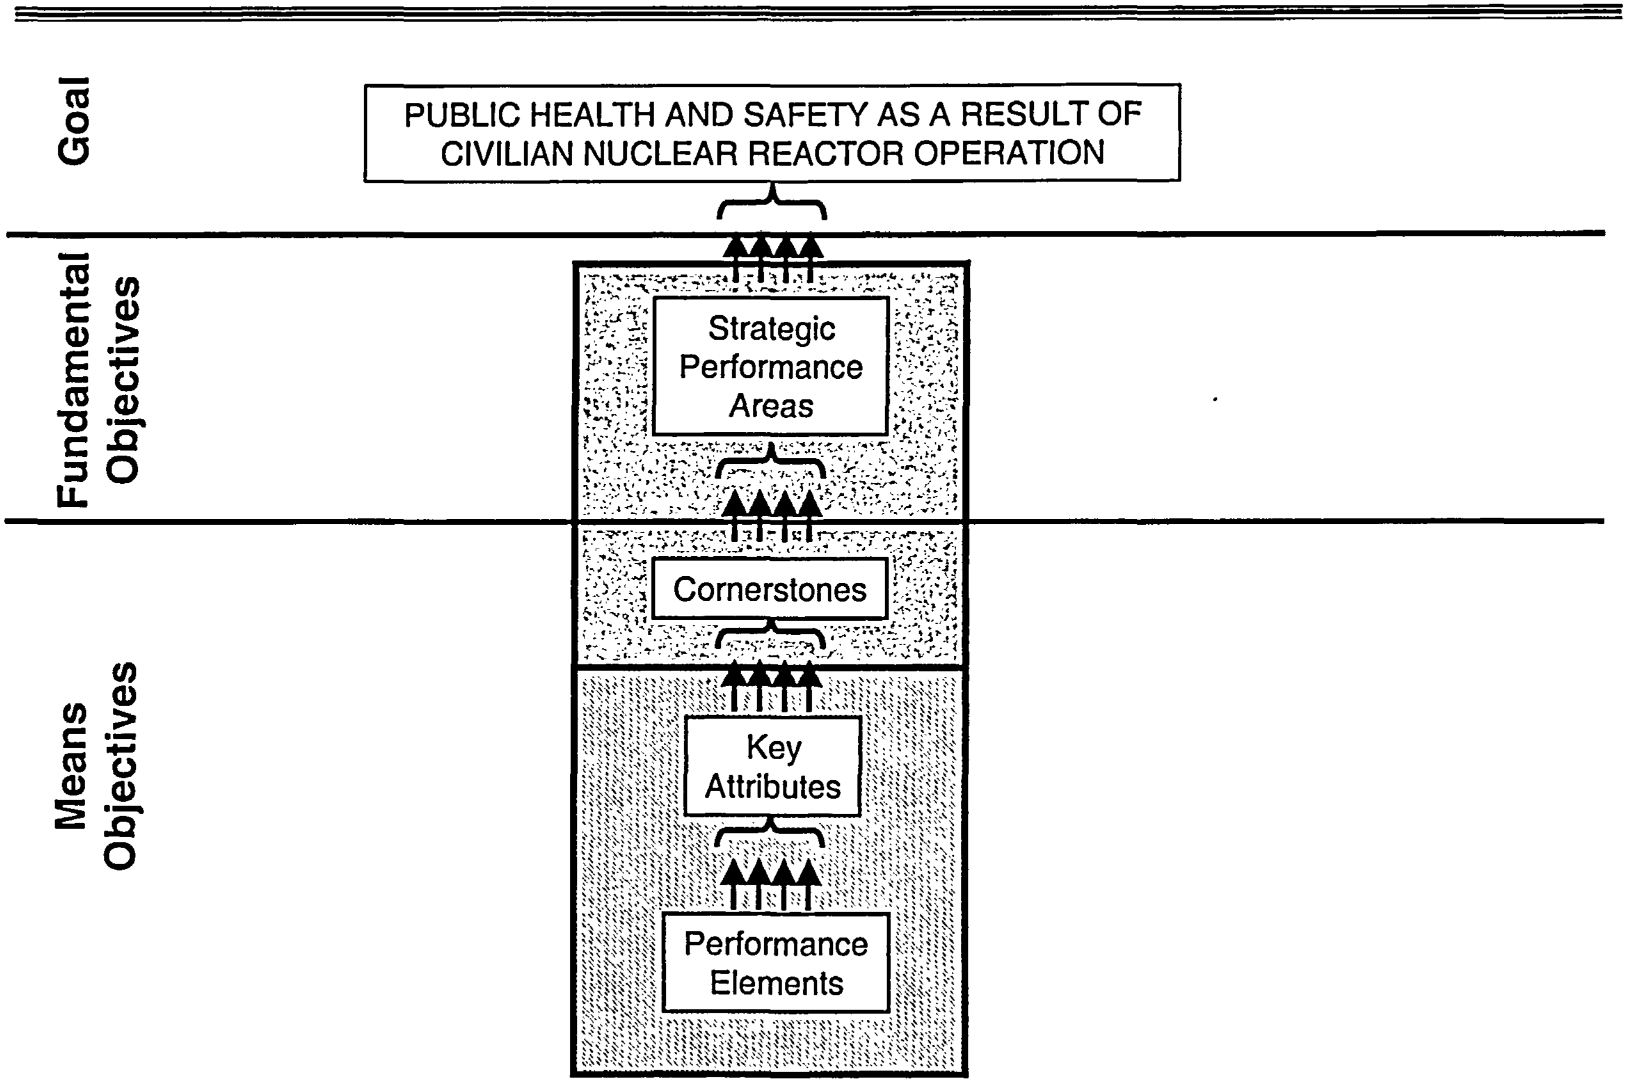

Figure III.2
Sheet 2 of Objectives Hierarchy

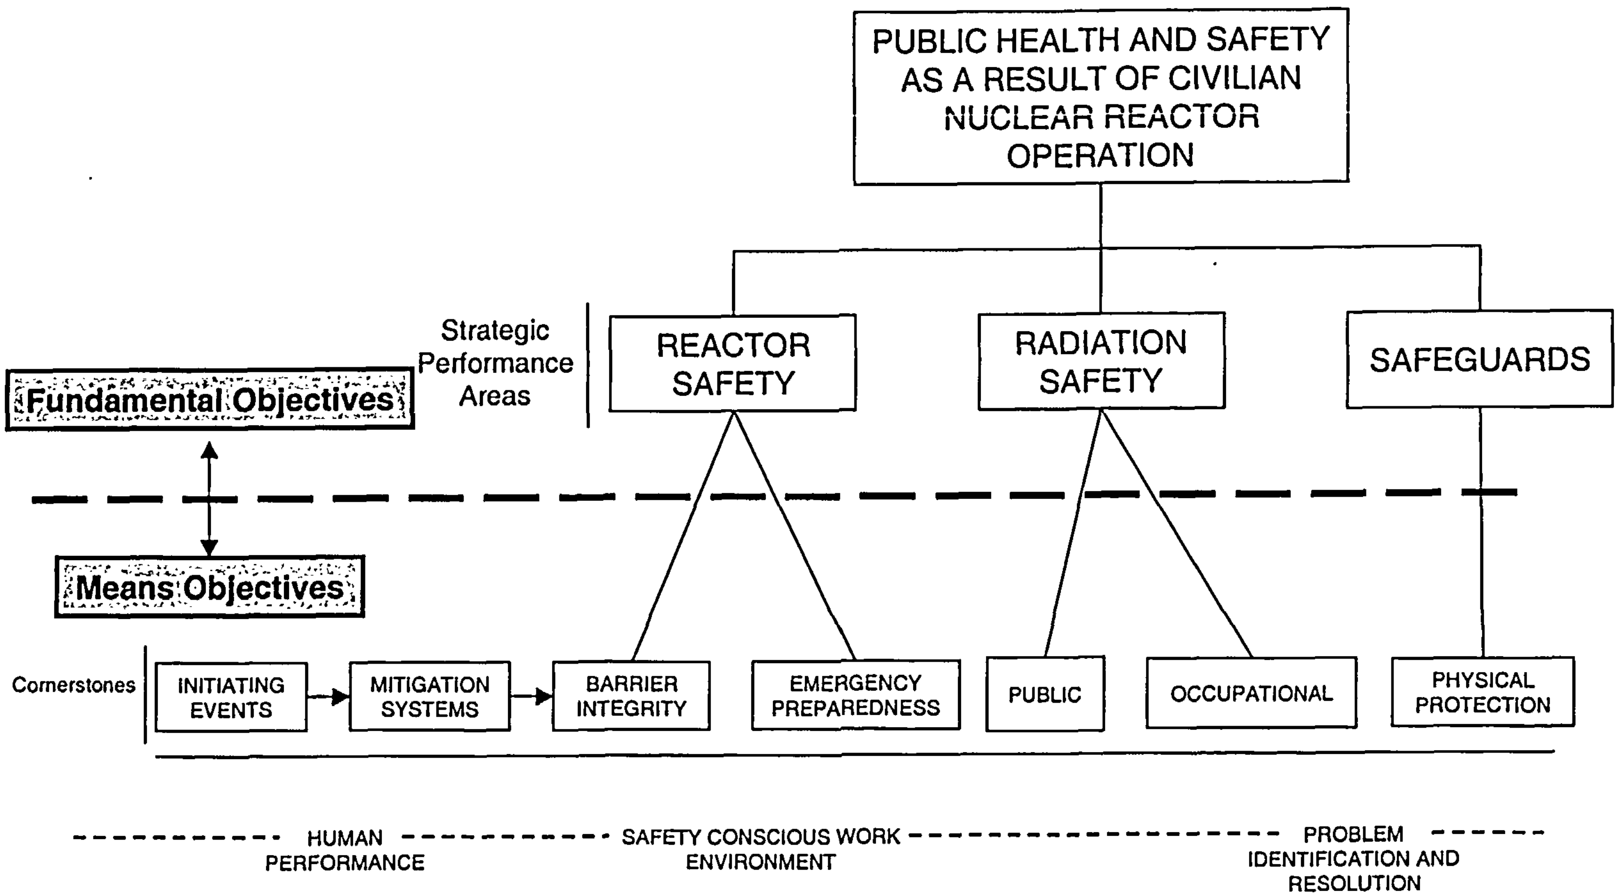

Figure III.3
Sheet 3 of Objectives Hierarchy (Lower-Level Portion)

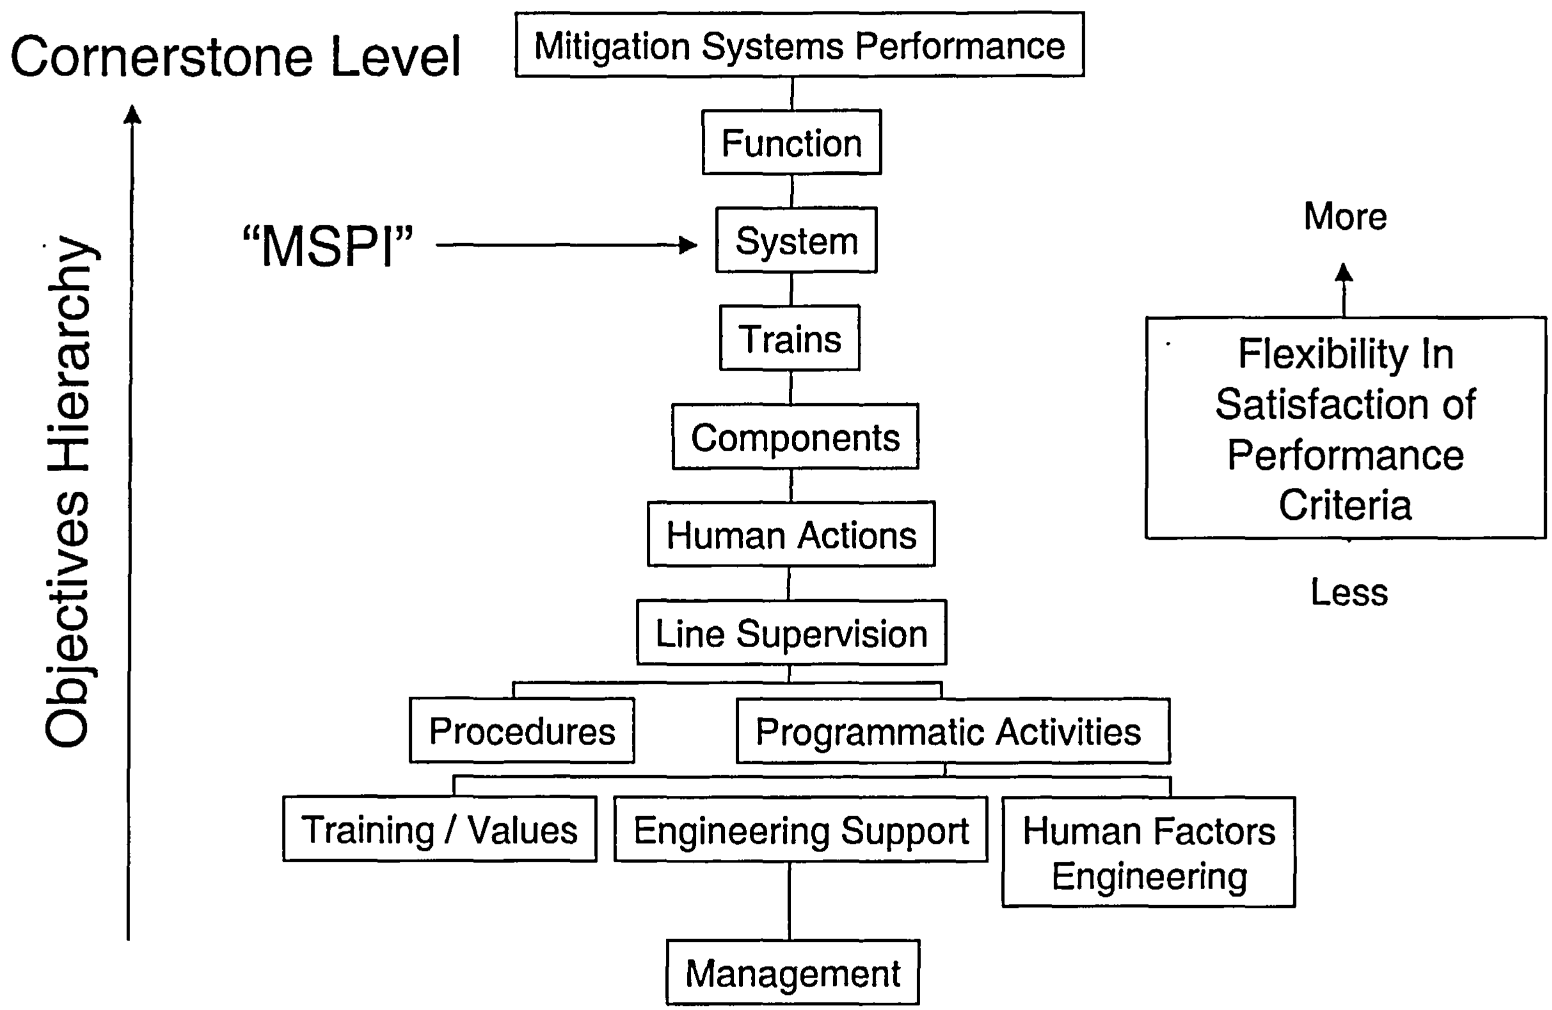

III.2 NATURAL MEASURES, CONSTRUCTED MEASURES, AND PROXY MEASURES

The subject of this subsection is the definition of quantities of different types that can be used to characterize the performance of different decision alternatives with respect to the fundamental objectives. Keeney (Ref. III-6) refers to these quantities as "attributes," a usage that comports with the idea that alternatives have attributes. Within the agency, some people prefer the term "measure" to "attribute," and this usage will be followed here. Per Keeney, different kinds of measures can be identified: natural, constructed, and proxy.

A natural measure is typically one that is in general use and is commonly understood, and corresponds directly to a fundamental objective. For example, the *cost* of implementing an alternative is an example of a natural measure that corresponds directly to achievement of objectives that are formulated in terms of cost (e.g., "minimize cost"). In many situations, it is less straightforward to measure accomplishment of an objective, and there are different kinds of reasons for this. The following paragraphs contain examples of this.

Where natural measures cannot be identified, or where it is impractical to model and quantify them, it may be appropriate to work with "constructed" measures. A constructed measure is defined in a specific decision context to (approximately) measure some objective. As with natural measures, a constructed measure relates different levels of some index to different levels of accomplishment of an objective; the difference is that a "constructed" measure is defined artificially by the decisionmakers in order to have the desired properties. For example, Keeney considers gross national product (GNP) to be an example of a constructed measure that reflects "economic health." There are many qualitative considerations like "economic health" that are not straightforwardly captured in any obvious natural measures. The formulation of GNP as a measure of economic health, and the decision to use it in certain contexts, is the domain of designated experts.

Where natural measures and constructed measures cannot be identified or quantified, it may be necessary to use "proxy" measures. A proxy measure is an indirect measure. For example, for purposes of the objective "minimize damage to stone statues and historic buildings from acid rain," Keeney considers "atmospheric concentration of sulfur dioxide in the vicinity of buildings" to be a proxy measure. It does not directly measure damage, but within simple models, arguably correlates with damage. Analogously, worker population dose might be a proxy for radiological worker safety.

Natural measures are preferred; constructed measures can be used where natural measures are lacking; proxy measures are the last resort.

What kinds of measures can be used in reactor regulation? One objective of regulation is to "protect the health and safety of the public." This translates into the prevention of various adverse consequences, including releases of radioactivity. Of course, many regulatory requirements do not explicitly refer to prevention of radiological release, and in fact may relate to it only indirectly. Traditional reactor regulation can arguably be summarized as a regulatory approach based on a particular set of *proxy* measures relating to means objectives (peak cladding temperature in design basis scenarios, containment response to non-mechanistic challenges, single-failure-proof "reliability") and prescriptive requirements aimed at promoting performance *as reflected in those specific measures.*

Risk-based regulation (if implemented) would differ fundamentally in the choice of measures used (CDF, LERF, expected fatalities, expected person-rem). It could be debated whether these metrics are "natural;" this would depend on the formulation of the fundamental objectives. In any case, they are clearly more directly related to the fundamental objective of protecting the health and safety of the public. Their recent application in *risk-informed* regulation has increased over the years not because the fundamental objectives of regulation have changed, but because of improvements in our ability to model these metrics (and perhaps the perception of our ability to model them).

Arguably, some of the perceived problems with today's regulations stem from the particular formulation of the proxy measures, combined with an overemphasis on them. Use of an objectives hierarchy and careful consideration of performance measures can lead to a more effective suite of performance measures. This would be true even for today's reactors; in the case of new designs, for which the regulatory criteria remain to be developed, careful consideration of performance measures will be very important.

III.3 FORMULATION OF ALTERNATIVES

A major theme of Keeney's work (Ref. III-6) is the arguably misplaced focus of many decision processes. It appears that many decision processes begin with recognition of a decision problem, and then proceed immediately to formulation of alternatives. Ultimately, the alternatives are ranked according to their standing with respect to metrics that may have been clarified only after the alternatives themselves had been formulated. Ref. III-6 argues that once a problem has been identified, it is preferable first to consider carefully the fundamental objectives, and only then undertake to formulate the best possible alternatives, and only then to undertake modeling. This latter approach is called "value-focused thinking" (the title of Ref. III-6), and the former is called "alternative-focused thinking."

A related idea is the distinction between decision opportunities and decision problems. In Keeney's terminology, a decision *problem* is an emergent need to make a decision to resolve a specific issue confronting the decision-maker: "problems" are precipitated by external parties or events. A decision *opportunity*, on the other hand, has to be recognized. Recognition of a decision opportunity may occur through a reflective evaluation of the current situation. Value-focused thinking lends itself to identification of decision opportunities; alternative-focused thinking, almost by definition, does not.

The following illustration of "problem" vs. "opportunity" is loosely based on past experience in advanced reactors. Some new designs use so-called "passive" safety systems that minimize or eliminate the role of active components, such as pumps. The failure probability of such systems may well be dominated by phenomena other than failure of active components. Partly because of regulatory history, methods for evaluating such systems are less generally accepted, although methods for quantifying the failure probability of active systems (which are far from perfect) are accepted (up to a point). The emergent need to evaluate a specific "passive" system, as part of a review process that is already underway, is a decision *problem*; however, *prior* recognition of the pervasiveness of the underlying issues (bearing on other new designs and on evaluation of engineering issues in *existing* designs), and looking for something broadly useful to do about them, represents creation of a decision *opportunity.*

Other kinds of decision problems can readily be identified. Once a design review is underway, many aspects pose decision "problems:" whether redundancy in a particular area is adequate, whether margin is adequate, and so on.

A handbook was developed at the Asian Development Bank Institute as a resource for capability building programs, using cases as learning and teaching tools (Ref. III-7). Morrison provides an overview of the decision making process in Part 8 of this handbook. In this section it is suggested that not all alternatives are the same, nor suited to all types of decision situations. It is suggested that the key to formulating well and meaningful alternatives involve:

1. examining options from the perspective of how different people might view them,
2. reversing the objective i.e., visualizing a situation that is the opposite in which the decision is set,
3. shifting one's perspective from the problem to the context in which the problem exists, and
4. using visualization capacities to determine what might go wrong in the future.

In his "Study Guides and Strategies Web Site," Landsberger explains how alternatives should be identified and formulated in decision-making situations (Ref. III-8). The study guide suggests brainstorming as an effective and excellent technique of discovering new alternatives. According to this guide, only those alternatives should be evaluated further in detail that are innovative and need more information, can be combined or eliminated, have the ability to meet opposition and seem promising or exciting.

The SFF (Suitability, Feasibility & Flexibility) Matrix is also described by Landsberger. This matrix is used to rate each alternative on a scale of 1-3 for its suitability, feasibility and flexibility.

Suitability refers to whether the alternative is ethical or practical and appropriate in scale or importance. Feasibility refers to how many resources are needed to solve the problem (i.e. examining the affordability). Flexibility refers to the ability to respond to unintended consequences, or openness to new possibilities. This matrix is then used to add the score for each alternative, comparing the alternatives against each other and then prioritizing the alternatives based on the results obtained.

III.3.1 Allocation of Performance

Given a top-level performance objective (such as CDF, availability, dose limits, or frequency of overexposure), the "allocation" of performance over barriers, SSCs, etc. is the assignment of performance targets to those elements, such that the top-level performance objective is met. Designers need to allocate in their formulation of design intent, in order to be able to carry out their design activities; regulators need to think in terms of an allocation of safety performance in order to reason appropriately about high-level issues such as the balance between prevention and mitigation, and about lower-level issues such as system reliability. Designers and regulators need not work from the same allocation; the designer's plant performance goals will ideally be more stringent than the regulator's allocation (Ref. III-2).

Before selecting performance measures, it is logically necessary to determine what kind of performance and what level of performance is needed. For example, in the reactor arena, each of the cornerstone areas (initiating events, mitigating systems, barrier integrity, emergency preparedness) receives attention. Performance is expected in each cornerstone area. Strong performance in all areas provides an important defense-in-depth component, because to some extent, performance in one area can compensate for lack of performance in another. For example, at most commercial reactors, an increase in initiating events frequency will not typically be a safety issue, if the mitigating systems' performance is satisfactory.

A choice among possible allocations must usually be made, because it is usually formally possible to meet a safety objective in more than one way. However, cost and practicality significantly restrict the possibilities that need to be considered. For reasons discussed in NUREG/BR-0303 (Ref. III-4), it is desirable to specify and monitor performance targets as high on the objectives hierarchy as possible. Allocating performance too far down on the hierarchy reduces licensee flexibility. Arriving at an implementation that maintains safety, while appropriately balancing licensee flexibility with the need for regulatory assurance of ongoing performance, will require some iteration with the allocation step.

III.3.1.1 Example

In order to understand what it means to allocate performance, and why we need to allocate performance, it is useful to consider a specific example. The quantitative analysis presented in the Reactor Safety Study showed that there was an imbalance in the way in which resources were allocated. Much attention had previously been focused on plant response conditional on a design-basis loss-of-coolant accident, while somewhat less attention had been focused on decay heat removal in scenarios initiated by transients or small breaks. This was addressed in a TMI requirement (Ref. III-9), as a result of which auxliliary feedwater systems (AFWS) in PWRs licensed during the 1980's were expected to demonstrate low unreliability "in the range of 1E-4 to 1E-5 per demand" (Ref. III-10). Assuming reasonable initiating event frequencies, satisfaction of this requirement would mean that a certain class of accident sequences was being controlled reasonably well.

This is a specific example of an allocation. Performance in the initiating events area was tacitly allocated at then-current levels, while target performance in AFWS reliability was set at a level somewhat better than was then being achieved at certain plants. Note that this was a design and licensing issue, not a "performance" issue.

This example serves to illustrate the previous comments regarding licensee flexibility. Refer again to Figure III.3. Some flexibility was made available in the application of this requirement to

AFWS systems, provided that alternative methods of core cooling were shown to be available (e.g., feed and bleed). Thus, the real allocation was pitched at the functional level (decay heat removal), with alternative allocations available at the systems level (allocate all performance to AFWS, or alternatively, allocate some to AFWS and some to high-pressure injection and primary depressurization).

Carried down to lower levels, these requirements have implications for such areas as in-service testing and inspection, and technical specifications.

III.3.1.2 Process of Allocation

The starting point is an allocation of performance at the upper levels of the objectives hierarchy. This allocation should appropriately balance prevention and mitigation, while satisfying agency safety objectives. This top-level allocation step is frequently done in the reactor arena beginning with the Quantitative Health Objectives, going on to CDF and LERF objectives, and then inferring practical limits on the frequencies of various accident sequence families from the heuristic guideline that no single sequence type should be allowed to dominate risk. This sort of insight can inform performance allocation over initiating event types, mitigating systems performance, and emergency preparedness.

From this point, one develops the allocation by decomposing each high-level performance objective into sub-objectives corresponding to the performance elements on the adjacent lower level. The collective satisfaction of these sub-objectives assures performance at the higher level. This entails making choices, because this decomposition is not unique. (The Analytic Hierarchy Process (AHP) is discussed below as a possible way of making such choices.) This process is continued on down the objectives hierarchy, until a point is reached at which it is no longer appropriate to articulate objectives, because there would be no net benefit to measuring performance at that level. The operational character of this guidance means that intelligent allocation must be done iteratively with the selection of performance measures and the implementation.

In the reactor arena, because of variations in plant design, it will be difficult to perform this kind of allocation generically below the functional level. As illustrated in the AFWS example cited above, the requirement in that case was actually formulated at the functional level, with the expectation that either the AFWS alone would meet it, or additional systems would be credited in a more flexible evaluation. On the other hand, system-level, train-level, and component-level requirements clearly derive from an implicit allocation of performance that is not numerically precise, but derives from implicit consideration of safety objectives. Improving the alignment of regulatory requirements with real safety objectives is the essence of "risk-informing" regulatory practice. Significant progress can be made in this area without overly detailed developments of objectives hierarchies or allocations.

Because allocation requires choices, it is appropriate to consider allocation itself as a decision problem. Given a spectrum of allocations that could nominally satisfy top-level *safety* objectives, which one should be implemented? Broadly speaking, many allocations can be rejected on cost grounds, because they call for unachievable (infinitely expensive) perfection in one or more areas. This still leaves many possibilities to consider. (Again, see below for discussion of the AHP.)

In the mid-1980's, some work was done (Ref. III-11) to explore the potential of allocation for NRC purposes. Four high-level objectives were used: core damage frequency, expected latent fatalities, expected acute fatalities, and cost. These objectives were used in a sample application to establish whether it is feasible to establish reliability targets for elements of a PRA model. The authors concluded that it is, in fact, feasible, but that in today's operating fleet, the allocation essentially has to be plant-specific. It also ultimately needs to be informed by what was then being called a "preference assessment." The report discusses stakeholder commentary on an earlier draft, some of which takes issue with the idea that allocation ought to be applied by NRC, at least at the rather detailed level explored in that report. However, it is acknowledged by the stakeholders that at least some of the issues that argue against allocation of today's individualized plants are inapplicable to "standardized" designs. It was concluded by some that allocation over elements of a standard design might usefully inform requirements aimed at its corresponding group of plants.

While the detailed allocation of the sort contemplated in NUREG/CR-4048 has not been applied widely, numerous instances of higher-level allocation have been applied, at least informally. For example, the idea that a single class of accidents ought not by itself to contribute to risk at the level of the quantitative health objectives has appeared in the technical rationale for selected backfits. It also appears in the presentation of the "Framework for Risk-Informed Changes to the Technical Requirements of 10 CFR 50" (Attachment 1 to SECY 00-198, Ref. III-12). In this Framework, the allocation began with the Quantitative Health Objectives, and proceeded to recommend (in its Figure 3.1) goal values for metrics related to core damage frequency, large early release frequency, and conditional containment failure probability, for initiating events having different frequencies (and potentially different characteristics). Those values were meant to be used in a process of evaluation of regulatory requirements.

For present purposes, the basis for that allocation is especially pertinent, because defense in depth plays an important role in it (Section III.2.2 of the Framework). According to the Framework, part of defense in depth is avoiding over-reliance on any one performance area in satisfying the top-level safety objectives. Therefore, defense-in-depth considerations favor allocations that spread out performance over redundant and diverse barriers and SSCs. This is a property of the Framework's recommendations, and is an example of the kind of "preference" mentioned (but not explored) in NUREG/CR-4048 (Ref. 11). The discussion in Section III.2.2 of the Framework also notes that where uncertainty is significant, tighter requirements may be applied in some areas to compensate for that uncertainty. This theme was taken up at greater length in Part II of this report.

The context of the "Framework" document was a staff initiative (Option 3) to review regulatory requirements for possible modification, and the high-level allocation was intended for that purpose. It can be seen as part of a value-focused process intended to steer the creation of decision alternatives for modifying regulations.

The Analytic Hierarchy Process, discussed below, has also been applied to allocation, and could be considered as a tool for making tradeoffs of the kind treated in the Framework document.

III.3.2 Formulation of an Appropriate Implementation: Selecting Performance Measures and Performance Goals

Deciding what functional performance to expect from various elements of an advanced design is part of the task. Another part of the task is to decide how that performance is to be achieved and assured. This calls for a suitable combination of regulatory requirements, inspections, and performance indicators, collectively aimed at assuring performance at (or better than) the stated target levels. These elements then need to be captured within appropriate parts of the regulatory framework.

The allocated performance level is the level of performance that is needed in order to satisfy safety objectives. In general, normal performance will be significantly better than this level. In order for a performance-based approach to be viable, the performance criterion needs to be set at a level having significant margin to the allocated level, in order for problems to be detected and addressed before a safety issue arises. In order for this approach to be practical, there must also be some margin between the normal performance level and the criterion.

It is useful to illustrate this point by continuing with the AFWS example introduced above, in which a target range for unreliability was given. Once that objective was fixed, very different implementations could have been chosen. In the implementation actually selected, considerable flexibility was allowed in design, and this flexibility is reflected in the variety of system topologies to be found in the US operating fleet. However, the implementation was not fully performance-based. Instead, the implementation followed the traditional regulatory approach: system reliability was achieved through numerous prescriptive ASME code requirements on testing and inspection, and Technical Specifications addressed allowed outage time in the usual way. Essentially, the effect of the top-level requirement on unreliability was primarily to drive performance to a level better than that of a typical, active two-train fluid system with no backup, and secondarily, to provide improved guidance on the formulation of the traditional prescriptive requirements.

Currently, formulation of regulatory approaches is supposed to be performance-based to the extent appropriate (Ref. III-4). Therefore, given the objectives hierarchy and the allocation for an advanced reactor, one would use the objectives hierarchy to select performance areas for attention within the regulatory approach, and try to work directly with performance areas as high up on the hierarchy as possible. If it is impractical to assess performance at a particular level, the next level down is tried. In the AFWS example, it is readily established that it is impractical to measure unreliability at the system level, because too few system demands occur to support a useful measurement. One must measure or prescribe at a level lower than that of the AFWS, and the complement of measures and prescriptive requirements must collectively provide the needed assurance of performance.

Allocation is necessary, but not always sufficient. When moving down the hierarchy, looking for a level at which an explicit allocation can be addressed either by monitoring or by prescriptive requirements, it is important to recognize that it is possible to underperform at a higher level, even if redundant elements at lower levels nominally succeed. To see why this is true, return once again to the AFWS example. Even if all trains are highly reliable and available, performance at the system level can suffer if (for example) the trains' unavailabilities are sometimes concurrent, or the contributors to train unreliability include shared hardware or common cause contributions affecting more than one train. Concurrent unavailability is addressed through technical specifications, and the possibility of common cause mechanisms is one reason for prescriptive requirements on testing and inspection. It is formally necessary to consider cross-cutting issues at each level in the hierarchy.

At this level, too, the AHP (discussed below) suggests itself as a tool for prioritizing performance measures.

III.3.3 Long-Term Priorities

In setting priorities for advanced reactor research, it is necessary but not sufficient to reach technically defensible conclusions. The conclusions must also be reached within a reasonable time. This is complicated by the circumstance that over time, as new information becomes available, priorities may shift; as some uncertainties are resolved, or designs are modified, the significance of other uncertainties may change. This circumstance affects many research program areas, not just advanced reactors (indeed, not just reactors). A comparatively recent idea presented by Keeney and McDaniels [Ref. III-13, hereafter KM] suggests a way of addressing this within formal decision theory.

The paper by KM is explicitly concerned with "thinking and analysis regarding climate change policies." The time scale of climate change is significantly longer than that of advanced-reactor research, but both areas affect future generations as well as the present generation, and neither arena can be addressed optimally by a static decision model aimed at resolving all program issues now, in the face of all of today's uncertainties.

In setting up their discussion, KM develop a preliminary list of fundamental objectives for climate change policy decisions (their Table 1). KM's *preliminary* top-level objectives are

> Minimize the net adverse impacts of climate change,
> Minimize net direct costs of climate change policies,
> Maximize equity.

Each of these is developed in more detail in KM. In principle, one could write down a decision model that addresses these objectives, and models the consequences of various alternatives, reflecting very substantial uncertainties in the consequences of those alternatives. Using this model, one could select an alternative despite the uncertainties, and go about implementing it. KM's main point is that there is a better approach than trying to settle policy now for all time in light of substantial uncertainties. A key excerpt from KM is provided below.

> Rather than attempt to model the long-term consequences of current decisions, analysts should use near-term proxy measures to describe the government's ability to deal effectively with future decisions when they are made. ... [T]he framework outlined here ... argues for an adaptive approach that focuses on selecting policies based on near-term consequences, *and the learning they will provide to place governments in better positions to address climate change decisions in the future* [emphasis added].

Later on in the paper, KM speak of two kinds of objectives: "fundamental objectives for which performance of alternatives can be predicted in the short term, and ... other objectives that are proxies for achieving better long-term performance on the same fundamental objectives.... The two kinds of objectives provide a practical basis for making current decisions that have both short-term and long-term implications."

In other words, for purposes of the near term, one supplements the preliminary set of fundamental objectives with the proxy objective "to position ourselves better to address these same fundamental objectives later on (KM Figure 1)." This includes fostering intellectual capital, fostering institutional capital, developing an improved basis for evaluating alternatives and an improved basis for formulating better ones in the first place, and so on (KM Figure 2).

This implies a program of research to position ourselves better in the future, and it declines to presume that current programmatic assumptions should foreclose certain kinds of research.

Proxy Objectives for Near-Term Decision-making

Following are proxy objectives that suggest themselves to research planners. That is, in applying a simple prioritization scheme to decide which of a proposed set of programs to fund, these objectives might beneficially supplement objectives that reflect the value of uncertainty reduction in specific safety issues. They are all aspects of what might be called stewardship.

- Develop / Maintain Intellectual Capital:

One important dimension of a research program is maintaining (even recovering) expertise in sub-fields in which expertise is being lost.

- Maintain Critical Facilities:

A substantial investment exists in certain kinds of facilities. This investment will be lost if these facilities are decommissioned. A snapshot of strictly near-term fundamental objectives might not provide a basis for maintaining them, but a broader consideration of keeping certain options open might lead to a different conclusion.

- Improve the Knowledge Base

  There are several dimensions to this. One is that more knowledge supports better evaluation of alternatives. Another is to recognize that better alternatives might be forthcoming from an improved knowledge base.

III.4 ANALYTIC HIERARCHY PROCESS

The analytic hierarchy process (AHP) (Ref. III-14) is widely used for ranking or prioritizing a list of items. According to its originator (Saaty), these items can be very different kinds of things. For example, Part I of this report summarizes a recent application of prioritization methodology, aimed at developing a tool for prioritizing the analysis of recent operating events for purposes of informing the current operations of a facility. In that application, each operating event receives a score based on its characteristics and on the relative weights of those characteristics. In that application, the AHP was used to establish the relative weights.

The method is based on a series of pairwise comparisons of the items in question. For each pair of items in the list, the decision-maker decides which of the two is more important, and how much more important it is. This way of eliciting values and judgments leaves many factors implicit, which some would consider a drawback. The method's primary author seems to believe that many complicating factors are actually best dealt with using just such an approach, because it is difficult or impossible to capture and model all such factors explicitly.

The results of the pairwise comparisons are quantified in a particular way (summarized below) and entered into a matrix. An eigenvector of this matrix (the one corresponding to the largest eigenvalue) then directly yields the relative weights of the items. Perhaps the easiest way to understand the method is to work backwards from the desired end result, namely, the rank ordering of the items. The following discussion is based on Saaty and Vargas. Suppose that we have a list of 4 items $\{A_i, i=1,4\}$ whose relative weights $\{w_i\}$ we already know. Define the matrix

$$\overline{\overline{A}} = \begin{bmatrix} \frac{w_1}{w_1} & \frac{w_1}{w_2} & \frac{w_1}{w_3} & \frac{w_1}{w_4} \\ \frac{w_2}{w_1} & \frac{w_2}{w_2} & \frac{w_2}{w_3} & \frac{w_2}{w_4} \\ \frac{w_3}{w_1} & \frac{w_3}{w_2} & \frac{w_3}{w_3} & \frac{w_3}{w_4} \\ \frac{w_4}{w_1} & \frac{w_4}{w_2} & \frac{w_4}{w_3} & \frac{w_4}{w_4} \end{bmatrix}.$$

Element (i,j) of this matrix is the ratio of the weight of the i^{th} item to that of the j^{th} item. Consider the vector

$$\begin{bmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \end{bmatrix}$$

Clearly, this vector is an eigenvector of $\overline{\overline{A}}$; it is also the list of weights whose determination in real examples is the purpose of this tool. It turns out that this eigenvector is the one corresponding to the largest eigenvalue of the matrix defined above. In the ideal (perfectly consistent) example presented above, this largest eigenvalue will numerically equal the number of items being ranked (in this case, 4).

Correspondingly, in its simplest form, the ranking tool of the AHP method consists essentially of eliciting the matrix elements themselves directly from the decisionmaker(s), entering these elements into the matrix, and finding the eigenvector of the largest eigenvalue.[3] The elements of this eigenvector give the desired ranking of items.

In real applications, the series of pairwise comparisons will not yield a perfectly consistent result. Inconsistency will manifest itself in the value of the largest eigenvalue. The greater the inconsistency is, the greater is the departure of the largest eigenvalue from being equal to "# of items." If the discrepancy becomes too large, as indicated by the eigenvalue (see Saaty and Vargas for details), then it is necessary to revisit the previous steps.

If this tool is being used for group decision-making, each matrix element is taken to be the geometric mean of values elicited from the individuals in the group.

The simple application summarized above is not the whole of the AHP. Rather, the AHP consists of applying the above tool repeatedly at each level of a comprehensively developed hierarchy of items. (This, of course, is where the method gets its name.) However, the AHP *tool* described above is widely applied to rank items in a list, as was done in the "prioritization" example discussed in the main body.

Allocation is discussed above as an important application of objectives hierarchies. The AHP has been used in allocation of resources (see Saaty and Vargas) and suggests itself for this application. Much existing work in the nuclear arena has been aimed at identifying sets of allocations having the property of satisfying higher-level objectives, and then ranking them by cost or other dimensions of preference. The full-blown AHP (applying the tool at several levels of a hierarchy) might complement those approaches, by providing a more natural way of folding in preferences regarding the proper balance between prevention and mitigation.

A virtue of the method is that it organizes group discussion in such a way that consensus regarding the ranking of a fixed list of items can be approached. On the other hand, some (e.g., p. 548 of Ref. III-5) appear to disagree with the idea of dealing only implicitly with the complexities and seeming intangibles that characterize many decision situations. For example, much of the effort in the PrOACT formulation (see Part I of this report) is aimed at flushing out precisely such complexities and "intangibles." The AHP offers a way to avoid getting mired in those complexities. However, the prioritization experience summarized in Part I shows that reconciling the views of diverse stakeholders is not easy even if the AHP is applied.

---

[3] Actually, less than half of the matrix elements need to be elicited; the diagonal elements are 1, and element (i,j) is the reciprocal of element (j,i).

III.5 SUMMARY

This part of the report has discussed the idea of an objectives hierarchy. The objectives hierarchy is vital not only to ranking decision alternatives once they are formulated, but perhaps more importantly to formulating good decision alternatives in the first place. Allocation (of resources or performance) also needs to be informed by a hierarchy of performance elements. Generally, many thought processes that are not always thought of as "decision analysis" can benefit from development of an objectives hierarchy.

Although this topic is basic to decision analysis, much energy has been spent on it even in recent times. One reason for this is that it is easy for decision-makers to default into alternative-focused thinking, in which much energy is spent on consequence modeling and ranking of alternatives that were developed after a relatively modest investment of energy. Countering a perceived tendency towards alternative-focused thinking evidently seems worthwhile to leaders in the field of decision analysis.

One output of the development of an objectives hierarchy is a set of performance measures or attributes. In general decision-making, these are factors in the ranking of decision alternatives; in prioritization, they are factors in the ranking of different items. Even the formulation of performance measures requires some thought; natural, constructed, and proxy measures are defined and compared. "Natural" measures correspond best to fundamental objectives, but may be unattainable for one reason or another. Constructed and proxy measures therefore have their uses. Over-reliance on proxy measures is arguably responsible for some of the regulatory conditions that risk-informed and performance-based regulation are intended to address.

As mentioned in Part I, two distinct decision processes need to be discussed in advanced-reactor prioritization. One process relates to safety decision-making for advanced reactors; this process has an objectives hierarchy that includes safety objectives as fundamental. Technical uncertainties associated with performance elements on this hierarchy are potential research topics. Prioritization of safety research has a distinct set of objectives, relating to institutional priorities, programmatic timing issues, the need to balance anticipatory and confirmatory research, and so on.

A regulatory alternative for addressing a safety objective in nuclear reactor regulation has two aspects: an allocation of performance, stating what levels of performance are required from particular parts of a facility, and an implementation of that allocation, articulating a program aimed at making the allocation "come true." As shown in Part I, the selection and weighting of the performance measures determines which technical uncertainties most limit the utility of the safety decisions being made. For example, technical uncertainties that affect only on-site consequences may cut very differently from technical uncertainties that bear on the likelihood of a large offsite release of radioactive material.

In the formulation and prioritization of research programs, the following objectives need to be addressed:

- technical uncertainties in safety assessments, and the value of reducing those uncertainties;
- cost of research programs;
- timing of available research outputs, and their relation to near-term agency needs;
- actual success probability of proposed research programs;
- the balance between programs promising modest progress at a high success probability and programs offering the potential for significant progress but having significant program uncertainty;
- the need to prepare the ground for future research programs:
  - the need to maintain and develop expertise;
  - the need to maintain and develop facilities.

In near-term research prioritization, the latter objective is essentially a proxy for downstream program needs that are linked to downstream safety decisions, but are not yet known in detail. Coping with such needs through use of such a proxy objective is a recent idea.

A complementary approach to some of the above tasks is the use of the AHP, which in its simplest form is a way of ranking a list of items. One "plus" of the AHP is that because of the way in which it elicits weighting factors, it naturally accounts for many implicit considerations whose elicitation and mapping into an explicit objectives hierarchy would be a daunting and time-consuming task.

III.6 REFERENCES

III.1. SECY 99-007, "Recommendations for Reactor Oversight Process Improvements," January 8, 1999 (USNRC).

III.2. NUREG/CR-5392, "Elements of an Approach to Performance-Based Regulatory Oversight," R. W. Youngblood *et al.* (USNRC, 1998).

III.3. NUREG/CR-6642, "Risk Analysis and Evaluation of Regulatory Options for Nuclear Byproduct Material Systems," E. R. Schmidt *et al.* (USNRC, 2000).

III.4. NUREG/BR-0303, "Guidance for Performance-Based Regulation," N. P. Kadambi (USNRC, 2002).

III.5. Robert T. Clemen, *Making Hard Decisions: An Introduction to Decision Analysis*, Second edition (Duxbury Press, 1996).

III.6. Ralph L. Keeney, Value-Focused Thinking: A Path to Creative Decisionmaking (Harvard University Press, 1992).

III.7. Terrence Morrison, Actionable Learning – A Handbook for Capacity Learning Through Case Based Learning, Asian Development Bank Institute, can be obtained from http://www.adbi.org/TM/Actionable%20Learning.pdf.

III.8. Joseph Landsberger, Study Guide and Strategies Web Site, University of St Thomas, http://www.iss.stthomas.edu/studyguides/problem/problemsolvingV2.htm.

III.9. 10 CFR 50.34 (f) (1) (ii) (TMI Requirements)

III.10. NUREG-0800, Standard Review Plan, Section 10.4.9, Rev. 2 (July 1981).

III.11. N. Z. Cho *et al.*, "A Methodology for Allocating Reliability and Risk," NUREG/CR-4048 (USNRC, 1986).

III.12. SECY 00-198, "Status Report on Study of Risk-Informed Changes to the Technical Requirements of 10 CFR Part 50 (Option 3) and Recommendations on Risk-Informed Changes to 10 CFR 50.44 (Combustible Gas Control)," September 14, 2000 (USNRC). (See especially Attachment 1, "Framework for Risk-Informed Changes to the Technical Requirements of 10 CFR 50")

III.13. Ralph L. Keeney and Timothy L. McDaniels, "A Framework to Guide Thinking and Analysis Regarding Climate Change Policies," Risk Analysis 21 No. 6, pp. 989-1000 (Society for Risk Analysis, 2001).

III.14. T. L. Saaty and L. G. Vargas, *The Logic of Priorities* (Kluwer-Nijhoff Publishing, 1982).

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-6833

2. TITLE AND SUBTITLE

Formal Methods of Decision Analysis Applied to Prioritization of Research and Other Topics

3. DATE REPORT PUBLISHED

MONTH	YEAR
September	2003

4. FIN OR GRANT NUMBER

Y6406

5. AUTHOR(S)

R.W. Youngblood, W.C. Arcieri, S.A. Faridi, N.P. Kadambi*

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.)*

Information Systems Laboratories, Inc. (ISL)
11140 Rockville Pike, Suite 500
Rockville, Maryland 20852

*Division of Systems Analysis and Regulatory Effectiveness
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001

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 Systems Analysis and Regulatory Effectiveness
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, DC 20555-0001

10. SUPPLEMENTARY NOTES

N.P. Kadambi, NRC Project Manager

11. ABSTRACT *(200 words or less)*

This report discusses the application of methods of formal decision analysis to prioritization of research to be carried out in support of licensing of advanced reactor designs. Formal decision methods are useful in this area for two reasons. (1) Prioritization is a special case of decision-making. (2) Prioritization of safety research is closely related to safety decision-making, and formal analysis of safety decisions points the way more clearly to specific research tasks needed to support safety decisions. The report is presented in three main parts. Part I provides an overview of prioritization. Two main themes emerge from this overview: (1) the effect of uncertainty on decisions, and (2) the need to clarify decision objectives and carefully formulate decision alternatives. The first theme is taken up in Part II, and the latter in Part III. Specific topics considered include the development and use of an objectives hierarchy, kinds of performance measures, and the value of information to reduce uncertainty, including a discussion of hypothesis testing and "Receiver Operating Characteristic Curves." Many agency decisions could arguably benefit from application of individual tools discussed here, although application of the full suite of formal decision-analysis tools may not be warranted in all cases.

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

Decision Analysis
Prioritization
Uncertainty
Receiver Operating Characteristic Curves
Formal Methods
Performance-Based Regulation

13. AVAILABILITY STATEMENT

unlimited

14. SECURITY CLASSIFICATION

(This Page)
unclassified

(This Report)
unclassified

15. NUMBER OF PAGES

16. PRICE

Printed on recycled paper

Federal Recycling Program

UNITED STATES
NUCLEAR REGULATORY COMMISSION
WASHINGTON, DC 20555-0001

OFFICIAL BUSINESS