STRATEGIC DECISIONS GROUP

READINGS ON

The Principles and Applications of DECISION ANALYSIS

EDITED BY
RONALD A. HOWARD &
JAMES E. MATHESON

VOLUME I: GENERAL COLLECTION

STRATEGIC DECISIONS GROUP

READINGS ON

The Principles and Applications of DECISION ANALYSIS

EDITED BY

RONALD A. HOWARD &

JAMES E. MATHESON

VOLUME I: GENERAL COLLECTION

EDITORS

Ronald A. Howard
James E. Matheson

PRODUCTION STAFF

Rebecca Sullivan, Coordinator
Paul Steinberg, Copy Editor
Leslie Fleming, Graphics

CONTENTS

VOLUME I: GENERAL COLLECTION

VOLUME II: PROFESSIONAL COLLECTION

FOREWORD

Since the term "Decision Analysis" was coined in 1963 (see paper #4), both its theory and practice have developed profusely. Stanford University has been a center for the intellectual development of decision analysis and the catalyst for its extensive application. Consultants associated with Stanford, many of them graduates of the Engineering-Economic Systems Department, have accumulated hundreds of man-years of experience.

This collection is intended to portray the "Stanford School of Decision Analysis," as viewed by the editors. Because the Stanford decision analysis community has the broadest base of practical experience, we believe these papers represent the most successful methods of dealing with decision problems. We have not attempted to represent alternative approaches or to enter into any debate of their relative merits. We have, however, included a few papers from other fields, notably psychology, that have had, and are having, a significant impact on the practice of decision analysis.

In these two volumes, we have collected papers on both the theory and application of decision analysis. Although most of these readings have been published elsewhere, we have added a few unpublished papers to represent recent developments.*

The first volume is designed to be accessible to a general readership and contains introductory papers and descriptions of actual applications. Applications to corporate strategic decisions are necessarily disguised and underrepresented because of their proprietary nature.

The second volume is designed for the professional student of decision analysis. In addition to containing professional and technical papers, it contains some papers discussing recent developments in methodology for approaching health and safety problems. While papers in this volume use technical terminology, many of their ideas will be understandable to anyone.

---

* Where possible, we have indicated authors' current affiliations on the title page of each paper. Affiliation references appearing within the text are taken from the original publication and, therefore, may vary from those on the title pages.

WHAT IS DECISION ANALYSIS?

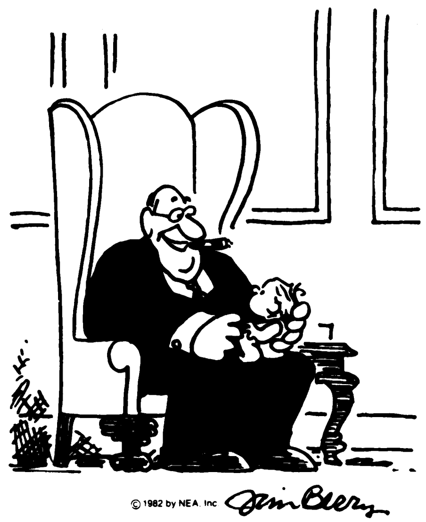

"Today, I'm going to tell you all you'll need to know about 'decision analysis.'"

Reprinted by permission. © 1982 NEA, Inc.

When this nationally syndicated cartoon appeared in 1982, decision analysis had clearly become a common term. In common usage, however, the term has lost precision. By decision analysis, we mean a discipline comprising the philosophy, theory, methodology, and professional practice necessary to formalize the analysis of important decisions. Decision analysis includes procedures and methodology for assessing the real nature of a situation in which a decision might be made, for capturing the essence of that situation in a formal but transparent manner, for formally "solving" the decision problem, and for providing insight and motivation to the decision-makers and implementers.

Confusing the tools of decision analysis with decision analysis itself has contributed to the loss of precision. Because uncertainty is at the heart of most perplexing decision problems, decision analysts frequently use specialized tools, such as decision tree techniques, to evaluate uncertain situations. Unfortunately, many people, some of them educators, have confused decision analysis with decision trees. This is like confusing surgery with the scalpel. Although decision tree techniques are extremely useful in solving problems where uncertainty is critical, in a real decision analysis, most of the effort and creativity is focused on finding and formulating the correct problem and on interpreting the results rather than on performing computations.

INTRODUCTION AND OVERVIEW

Preface

These papers describe the philosophy and methodology of decision analysis.

"The Evolution of Decision Analysis" was written especially for this collection to show the progress in the field. It describes the continuing development of the decision analysis cycle and of the process for capturing the three elements of any decision problem -- values, alternatives, and information -- in formal, but practical, decision models.

"An Introduction to Decision Analysis" extensively discusses the principles and practice of decision analysis and describes the original decision analysis cycle, which was updated in the previous paper.

"Decision Analysis in Systems Engineering," originally presented as a lecture, provides a non-technical discussion of the basic principles and techniques developed from them. The paper discusses the nature of decisions, the relation of rational decision-making to mental health, and its applications to medical and social decisions. It also includes a transcript of a question period with some interesting exchanges.

"Decision Analysis: Applied Decision Theory" introduced the term decision analysis when the paper was presented at a conference in 1965. It describes the earliest version of the decision analysis cycle and one of its first extensive applications to a major problem.

"A Tutorial Introduction to Decision Theory" presents an entertaining example of the theory for treating decisions in the face of uncertainty, which is a cornerstone of decision analysis, and focuses on decision theory as a way of formalizing common sense.

"A Tutorial in Decision Analysis" illustrates the principles and practice of decision analysis by discussing an analysis of a major capital investment decision. It shows how to treat ecological and regulatory issues and how to use value of perfect information calculations.

"The Science of Decision-Making" is an approachable statement of the logical foundations of decision analysis. The paper discusses the barriers to logical thought that had to be surmounted, and the developments required, to create a science of rationality.

"An Assessment of Decision Analysis" is a fairly recent critique of the usefulness and limitations of decision analysis, which, in particular, questions the ethics of using decision analysis in social situations where individuals are involuntarily subjected to the result.

1

THE EVOLUTION OF DECISION ANALYSIS

Ronald A. Howard

Department of Engineering-Economic Systems
Stanford University

Although decision analysis has developed significantly over the last two decades, the basic principles of the field have served well. They are unlikely to change because they are based on simple logic. In the first part of this paper, we summarize the original, fundamental disciplines of decision analysis; in the second part, we show how the discipline has evolved.

PART I: A BRIEF DESCRIPTION OF DECISION ANALYSIS

Making important decisions often requires treating major uncertainty, long time horizons, and complex value issues. To deal with such problems, the discipline of decision analysis was developed. The discipline comprises the philosophy, theory, methodology, and professional practice necessary to formalize the analysis of important decisions.

Overview of Decision Analysis

Decision analysis is the latest step in a sequence of quantitative advances in the operations research/management science field. Specifically, decision analysis results from combining the fields of systems analysis and statistical decision theory. Systems analysis, which grew as a branch of engineering, was good at capturing the interactions and dynamic behavior of complex situations. Statistical decision theory was concerned with logical decisions in simple, uncertain situations. The merger of these concepts creates a methodology for making logical decisions in complex, dynamic, and uncertain situations.

Decision analysis specifies the alternatives, information, and preferences of the decision-maker and then finds the logically implied decision.

Decision-making requires choosing between alternatives, mutually exclusive resource allocations that will produce outcomes of different desirabilities with different likelihoods. While the range of alternatives to be considered is set by the decision-maker, the decision analyst may be able to suggest new alternatives as the analysis progresses.

Since uncertainty is at the heart of most significant decision problems, decision-making requires specifying the amount of uncertainty that exists given available information. Many decision problems become relatively trivial if uncertainty is removed. For example, consider how easily a decision-maker could make a critical decision in launching a new commercial product if he could predict with certainty production and sales costs, price-demand relationships, and governmental decisions. Decision analysis treats uncertainty effectively by encoding informed judgment in the form of probability assignments to events and variables.

Decision-making also requires assigning values on the outcomes of interest to the decision-maker. These outcomes may be as customary as profit or as troubling as pain. Decision analysis determines the decision-maker's trade-offs between monetary and non-monetary outcomes and also establishes in quantitative terms his preferences for outcomes that are risky or distributed over time.

One of the most basic concepts in decision analysis is the distinction between a good decision and a good outcome. A good decision is a logical decision -- one based on the information, values, and preferences of the decision-maker. A good outcome is one that is profitable, or otherwise highly valued. In short, a good outcome is one that we wish would happen. By making good decisions in all situations that face us, we hope to ensure as high a percentage of good outcomes as possible. We may be disappointed to find that a good decision has produced a bad outcome, or dismayed to learn that someone who has made what we consider to be a bad decision has achieved a good outcome. Short of having a clairvoyant, however, making good decisions is the best way to pursue good outcomes.

An important benefit of decision analysis is that it provides a formal, unequivocal language for communication among the people included in the decision-making process. During the analysis, the basis for a decision becomes evident, not just the decision itself. A disagreement about whether to adopt an alternative may occur because individuals possess different relevant information or because they place different values on the consequences. The formal logic of decision analysis subjects these component elements of the decision process to scrutiny. Information gaps can be uncovered and filled, and differences in values can be openly examined. Revealing the sources of disagreement usually opens the door to cooperative resolution.

The formalism of decision analysis is also valuable for vertical communication in a management hierarchy. The organizational value structure determined by policymakers must be wedded to the detailed information that the line manager, staff analyst, or research worker possesses. By providing a structure for delegating decision-making to lower levels of authority and for synthesizing information from diverse areas for decision-making at high levels, decision analysis accomplishes this union.

Methodology

The application of decision analysis often takes the form of an iterative procedure called the Decision Analysis Cycle (see Figure 1). Although this procedure is not an inviolable method of attacking the problem, it is a means of ensuring that essential steps have been considered.

The procedure is divided into three phases. In the first (deterministic) phase, the variables affecting the decision are defined and related, values are assigned, and the importance of the variables is measured without any consideration of uncertainty.

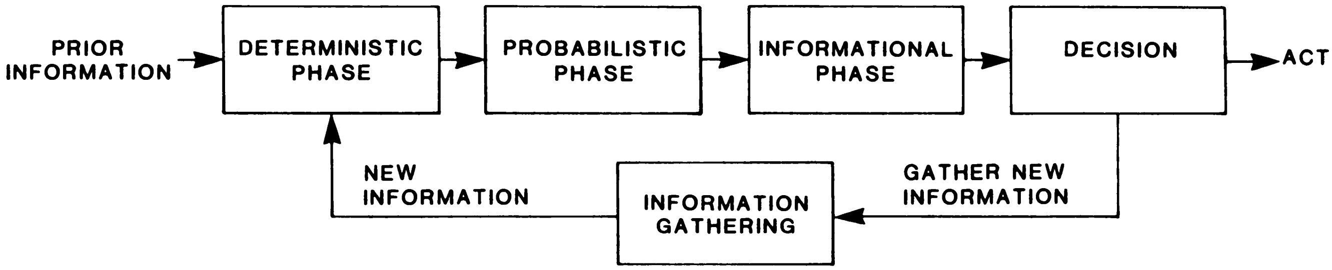

Figure 1: The Decision Analysis Cycle

The second (probabilistic) phase starts with the encoding of probability on the important variables; then, the associated probability assignments on values are derived. This phase also introduces the assessment of risk preference, which defines the best solution in the face of uncertainty.

In the third (informational) phase, the results of the first two phases are reviewed to determine the economic value of eliminating uncertainty in each of the important variables in the problem. In some ways, this is the most important phase because it shows just what it would be worth in dollars and cents to have perfect information. Comparing the value of information with its cost determines whether additional information should be collected.

If there are further profitable sources of information, then the decision should be to gather the information rather than to make the primary decision at this time. The design and execution of the information-gathering program follows.

Since new information generally requires revisions in the original analysis, the original three phases must be performed once more. However, the additional work required to incorporate the modifications is usually slight, and the evaluation, rapid. At the decision point, it may again be profitable to gather new information and repeat the cycle, or it may be more advisable to act. Eventually, the decision to act will be made because the value of new analysis and information-gathering will be less than its cost.

Applying the above procedure ensures that the total effort is responsive to changes in information -- the approach is adaptive. Identifying the crucial areas of uncertainty can also aid in generating new alternatives for future analysis.

Model Sequence

Typically, a decision analysis is performed not with one, but with a sequence of progressively more realistic models. These models generally will be in the form of computer programs. The first model in the sequence is the pilot model, an extremely simplified representation of the problem useful only for determining the most important relationships. Although the pilot model looks very little like the desired final product, it is indispensable in achieving that goal.

The next model in the sequence is the prototype model, a quite detailed representation of the problem that may, however, still be lacking a few important attributes. Although it will generally have objectionable features that must be eliminated, it does demonstrate how the final version will appear and perform.

The final model in the sequence is the production model; it is the most accurate representation of reality that decision analysis can produce. It should function well even though it may retain features that are treated in a less than ideal way.

Starting with the pilot model, sensitivity analyses are used throughout each phase to guide its further evolution. If decisions are insensitive to changes in some aspect of the model, there is no need to model that particular aspect in more detail. The goal of a good modeler is to model in detail only those aspects of the problem that have an impact on the decisions, while keeping the costs of this modeling commensurate with the level of the overall analysis.

Important aids in determining whether further modeling is economically justifiable are the calculations of the value of information. Some variables may be uncertain partially because detailed models have not been constructed. If the analyst can calculate the value of perfect information about these variables, he will have a standard to use in comparing the costs of any additional modeling. If the cost of modeling is greater than the value of perfect information, the modeling is clearly not economically justifiable.

Using a combination of sensitivity analysis and calculations of the value of information, the analyst continually directs the development of the model in an economically efficient way. An analysis conducted in this way provides not only answers, but also often insights for creating new alternatives. When completed, the model should be able to withstand the test of any good engineering design: additional modeling resources could be utilized with equal effectiveness in any part of the model. There is no such thing as a final or complete analysis; there is only an economic analysis given the resources available.

PART II: REFINEMENTS AND NEW DEVELOPMENTS IN DECISION ANALYSIS

Having seen the basic concepts of decision analysis and the main points of its professional practice, let us now examine some of the evolutionary changes in the field over the last two decades.

The Decision Basis

It has become useful to have a name for the formal description of a decision problem; we call it the decision basis. The decision basis consists of a quantitative specification of the three elements of the basis: the alternatives, the information, and the preferences of the decision-maker. We can then think of two essential steps in any decision analysis: the development and the evaluation of the decision basis.

Basis Development

To develop the decision basis, the decision analyst must elicit each of the three elements from the decision-maker or from his delegates. For example, in a medical problem, the ultimate decision-maker should be the patient. The patient would provide the element of preference in the basis, probably in a series of interviews with the decision analyst. In most cases, however, the patient will delegate the alternative and information elements to doctors who, in turn, would be interviewed by the decision analyst. The analyst should be able to certify that the decision basis accurately represents the alternatives, information, and preferences provided directly or indirectly by the decision-maker. We should note here that the alternatives must include alternatives of information-gathering, such as tests, experimental programs, surveys, or pilot plants.

One key issue is the extent to which the decision analyst can provide substantive portions of the decision basis by acting as an expert. In many circumstances, the analyst cannot be an expert because he has only a lay knowledge of the decision field. Even when the analyst does have substantial knowledge of the subject area, he should make clear to the decision-maker when he has changed from the role of decision analyst to that of substantive expert. Playing the role of expert can also force the analyst to defend his views against those of others; to this extent, he would be less of a "fair witness" in the subsequent analysis. Nevertheless, this possible loss of impartiality and fresh viewpoint must be balanced against the communication advantages of dealing with an analyst familiar with the decision field.

Basis Evaluation

Once the basis is developed, the next step is to evaluate it using the sensitivity analysis and value of information calculations described earlier. However, casting the problem as a decision basis shows that value-of-information calculations, important as they are, focus on only one element of the basis -- information.

Using the concept of the basis, we can also compute the value of a new alternative, which we might call the value of control. Such a calculation might well motivate the search for an alternative with certain characteristics and perhaps even the development of such an alternative.

One can perform a similar sensitivity analysis to preference with the intention not of changing preference, but of ensuring that preferences have been accurately assessed. A large change in value resulting from a small change in preference would indicate the need for more interviews about preference.

A Revised Cycle

Using the concept of the basis, we may wish to restructure the decision analysis cycle in the four-phase form shown in Figure 2. Here, the information gathering that must precede analysis or augment subsequent analyses has been included in a basis development phase. The deterministic and probabilistic phases are essentially unchanged, but the informational phase -- renamed "basis appraisal" -- is expanded to include the examination of all three basis elements.

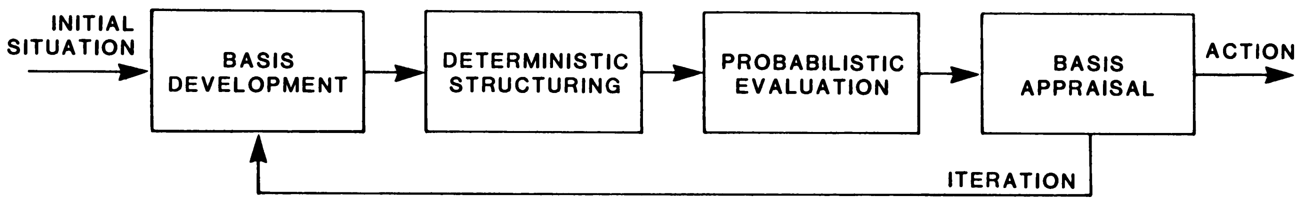

Figure 2: The Revised Decision Analysis Cycle

A Refined Analysis Sequence

As a problem is analyzed, the analysis may progress through the decision analysis cycle several times in increasing levels of detail. The basic distinction is between the pilot and full-scale analysis. The pilot analysis is a simplified, approximate, but comprehensive, analysis of a decision problem. The dictionary defines pilot as "serving as a tentative model for future experiment or development." The full-scale analysis is an increasingly realistic, accurate, and justifiable analysis of a decision problem, where full-scale is defined as "employing all resources, not limited or partial." To understand these distinctions, we must explain in more detail what constitutes a pilot or full-scale analysis.

The purpose of a pilot analysis is to provide understanding and establish effective communication about the nature of the decision and the major issues surrounding it. The content of the pilot analysis is a simplified decision model, a tentative preference structure, and a rough characterization of uncertainty. From a pilot analysis, the decision-maker should expect preliminary recommendations for the decision and the analyst should expect guidance in conducting the full-scale analysis.

The purpose of the full-scale analysis is to find the most desirable action, given the fully developed decision basis. The full-scale analysis consists of a balanced and realistic decision model, preferences that have been certified by the decision-maker, and a careful representation of important uncertainties. From the full-scale analysis, the decision-maker should expect a recommended course of action.

While most analyses progress from pilot to full-scale, some are so complex that valuable distinctions may be made between different stages of full-scale analysis.

The first stage of full-scale analysis is the prototypical stage, which is intended to reveal weaknesses and excesses in the full-scale analysis that are worthy of correction. A prototype is defined as "an original type, form, or instance that serves as a model on which later stages are based or judged."

After the indicated corrections have been made, the analyst has an integrated stage of full-scale analysis that provides the decision-maker with confidence in having a unified, balanced, and economic analysis as a basis for decision. To integrate is "to make into a whole by bringing all parts together: unify." If a decision-maker is making a personal decision that will not require the support or approval of others, then the integrated stage of full-scale analysis is all that is required. However, if the decision-maker must convince others of the wisdom of the chosen course of action or even defend that course against hostile elements, then an additional stage of full-scale analysis will be necessary -- the defensible stage.

The defensible stage of full-scale analysis is intended to demonstrate to supportive, doubtful, and possibly hostile audiences that the analysis provides an appropriate basis for decision. Defensible means "capable of being defended, protected, or justified." Typically, defensible analyses are necessary for important decisions in the public arena; however, even private enterprises may wish to conduct defensible analyses to win the support of workers, financial institutions, or venture partners. Defensible analyses are very demanding because they must show not only that the basis used is reasonable, but also that other possible bases that would lead to different decisions are not reasonable.

Contributions from Psychological Research

One of the most significant factors influencing the practice of decision analysis in recent years has been new knowledge about cognitive processes from the field of psychology. This research, centering on the contributions of Kahneman and Tversky, has had two major effects. First, the research on cognitive biases [10] has shown the need for subtlety and careful procedure in eliciting the probabilistic judgments on which decision analysis depends. Second, and perhaps even more important, the descriptive research on how people actually make decisions [6,11] shows that man is considerably less skilled in decision-making than expected. The main thrust of this research shows that people violate the rules of probabilistic logic in even quite simple settings. When we say that people violate certain rules, we mean that when they are made aware of the implications of their choices, they often wish they had made another choice: that is, they realize they have made a mistake. While these mistakes can be produced in analyzing simple decision settings, they become almost unavoidable when the problem is complex.

These findings may change our interpretation of the logical axioms that are the foundations of decision analysis. We have always considered these axioms as normative: they must be satisfied if our decisions are to have many properties that we would regard as desirable. If a particular individual did not satisfy the axioms, then he would be simply making mistakes in the view of those who followed the axioms. While this interpretation is still possible, a more appropriate way to look at the axioms is that they describe what any person would do if faced with a situation as simple as the one described by the axioms. In other words, the axioms are descriptive of human behavior for simple situations. If, however, the situation becomes more complex, more "opaque" as opposed to "transparent," the axioms are no longer descriptive because the person may unintentionally violate the axiom systems.

We may now think of the job of the decision analyst as that of making "opaque" situations "transparent," so that the person clearly sees what to do. This interpretation of the work may not make it any easier, but it is far more humane than the view that the analyst is trying to impose logic on a willfully illogical world.

Influence Diagrams

The influence diagram is one of the most useful concepts developed in decision analysis [3]. The analyst has always faced the problem of how to reduce the multifaceted knowledge in people's heads to a form that could meet the rigid tests of explicitness and consistency required by a computer. The influence diagram is a major aid in this transformation because it crosses the border between the graphic view of relationships that is very convenient for human beings and the explicit equations and numbers that are the province of present computers. To find a device that can readily be sketched by a layman and yet be so carefully defined that useful theorems concerning it can be proved by formal methods is rare. Although there is a danger that people who do not thoroughly understand influence diagrams may abuse them and be misled, there is an even greater promise that the influence diagram will be an important bridge between analyst and decision-maker.

Valuing Extreme Outcomes

One of the problems perplexing early users of decision analysis was how to treat outcomes so extreme that they seemed to be beyond analysis. For example, the question of how a person's death as the result of medical treatment can be balanced with other medical outcomes, like paralysis or even purely economic outcomes, was especially demanding. These problems appear to raise both ethical dilemmas and technical difficulties. One ethical dilemma centered on who had the right to value lives. A technical difficulty was revealed when an economist testifying in court on the value of a life was asked whether he would be willing to allow himself to be killed if he were given that amount of money. Nevertheless, once the ethical issue is clarified by acknowledging that a person may properly place a value on his own life, then the technical question of how to do it can be addressed quite satisfactorily, especially in the case of exposure to the many small risks present in modern life [4,5]. The results have major implications for many decisions affecting health and safety.

The development of ways to think about the unthinkable has shown that no decision problem lies beyond the realm of decision analysis. That is very satisfying, for were you faced with medical decisions about a loved one, would you want to use second-rate logic any more than a second-rate doctor?

Conclusion

When decision analysis was first developed, a common comment was, "If this is such a great idea, why doesn't [insert name of large, famous company] use it?" Today, it is difficult to find a major corporation that has not employed decision analysis in some form. There are some factors that should lead to even greater use. For example, decision analysis procedures are now more efficiently executable because the increased power of modern computers has reduced the costs of even very complex analyses to an affordable level. The problems that can be successfully attacked now run the gamut of all important decision problems. Increasing uncertainties and rapid change require fresh solutions rather than tested "rules of thumb." Some day, decision analysis of important decisions will perhaps become recognized as so necessary for conducting a provident life that it will be taught in grade school rather than in graduate school.

REFERENCES

1. Ronald A. Howard, "Decision Analysis: Applied Decision Theory," Proceedings of the Fourth International Conference on Operational Research, Wiley-Interscience, New York, 1966, pp. 55-71.

2. Ronald A. Howard, "The Foundations of Decision Analysis," IEEE Transactions on Systems Science and Cybernetics, SSC-4, No. 3, (September 1968): 211-19.

3. Ronald A. Howard and James E. Matheson, "Influence Diagrams," Department of Engineering-Economic Systems, Stanford University, July 1979.

4. Ronald A. Howard, "On Making Life and Death Decisions," Societal Risk Assessment, How Safe Is Safe Enough?, Edited by Richard C. Schwing and Walter A. Albers, Jr., General Motors Research Laboratories, Plenum Press, New York, 1980.

5. Ronald A. Howard, James E. Matheson, and Daniel L. Owen, "The Value of Life and Nuclear Design," Proceedings of the American Nuclear Society Topical Meeting on Probabilistic Analysis of Nuclear Reactor Safety, American Nuclear Society, May 8-10, 1978.

6. Daniel Kahneman and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," Econometrica, 47, No. 2 (March 1979): 263-291.

7. D. Warner North, "A Tutorial Introduction to Decision Theory," IEEE Transactions on Systems Science and Cybernetics, SSC-4, No. 3, (September 1968): 200-10.

8. Howard Riaffa, Decision Analysis: Introductory Lectures on Choices under Uncertainty, Addison-Wesley, 1968.

9. Myron Tribus, Rational Descriptions, Decisions, and Designs, Pergamon Press, 1969.

10. Amos Tversky and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," Science, 185 (Sept. 27, 1974): 1124-1131.

11. Amos Tversky and Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice," Science, 211 (Jan. 30, 1981): 453-458.

2

AN INTRODUCTION TO DECISION ANALYSIS

J. E. Matheson and R. A. Howard

Department of Engineering-Economic Systems
Stanford University

Contents

INTRODUCTION

Decision analysis is a term used to describe a body of knowledge and professional practice for the logical illumination of decision problems. It is the latest link in a long chain of quantitative advances in management that have emerged from the operations research/ management science heritage. It is the result of combining aspects of systems analysis and statistical decision theory. Systems analysis grew as a branch of engineering whose strength was consideration of the interactions and dynamic behavior of complex situations. Statistical decision theory was concerned with how to be logical in simple uncertain situations. When their concepts are merged, they can reveal how to be logical in complex, dynamic, and uncertain situations; this is the province of decision analysis.

Thus, decision analysis focuses logical power to reduce confusing and worrisome problems to their elemental form. It does this not only by capturing structure, but by providing conceptual and practical methods for measuring and using whatever knowledge regarding uncertainty is available, no matter how vague. When all available knowledge has been applied, the problem is reduced to one of preference; thus the best alternative will depend on the desires of the decision-maker. Here again, decision analysis provides conceptual and practical methods for measuring preferences. The problem may require expressing the relative desirability of various outcomes, the effect on desirability of changes in timing, and the tolerance for uncertainty in receiving outcomes. In particular, the impact of uncertainty upon the decision can be measured and interpreted — not left to intuition.

BACKGROUND

History of Quantitative Decision-Making

Operations Research

Operations research was the first organized activity in the scientific analysis of decision-making. It originated in the application of scientific methods to the study of air defense during the Battle of Britain. The development of operations research continued in the U.S. in the Navy's study of antisubmarine and fleet protection problems. After World War II, many of the scientists experienced in operations research decided to apply their new tools to the problems of management.

However, an examination of the transition of operations research from military to civilian problems shows that the limitations inherent in the military applications carried over to the civilian work. Many of the operations researchers trained in the military environment had become used to working only on operationally repetitive problems. In these constantly recurring problems, the impact of the formal analysis became evident to even the most skeptical observers. Some of the researchers, however, concluded that only this type of problem was susceptible to scientific analysis—that is they limited operations research to the study of repetitive processes.

Since repetitive decisions are also important to the civilian world, operations research made substantial headway in its new environment. Yet, the insistence on repetition confined the efforts of operations researchers within the province of lower and middle management, such as inventory control, production scheduling, and tactical marketing. Seldom did the analysts study decision problems relevant to the top executive.

Management Science

In the mid-1950s, operations research spawned an offshoot—management science. This discipline developed in response to a deep concern that the special problems of management were not receiving sufficient attention in operations research circles. This new field grew to emphasize science more than management, however. Management scientists have been accused of having more interest in those problems that are subject to elegant mathematical treatment than in those of the top executive, which are generally less easily quantified.

Although many students of business have considered the problems of top management, they have not generally had the scientific and mathematical training necessary to give substance to their ideas and to allow their application in new situations. When the top manager sought help on a problem, he often had to choose between a mathematician who was more concerned with the idiosyncrasies of the situation than with its essence and an experienced "expert" who might be tempted to apply an old solution to a radically new problem. Thus, the early promise of scientific aids for the executive was slow in materializing.

Decision Analysis

In the last few years, a new discipline, called "decision analysis," has developed from these predecessors. It seeks to apply logical, mathematical, and scientific procedures to the decision problems of top management that are characterized by the following:

- *Uniqueness.* Each is one of a kind, perhaps similar to—but never identical with—previous situations.
- *Importance.* A significant portion of the organization's resources is in question.
- *Uncertainty.* Many of the key factors that must be taken into account are imperfectly known.
- *Long run implications.* The enterprise will be forced to live with the results of the situation for many years, perhaps even beyond the lifetimes of all individuals involved.
- *Complex preferences.* The task of incorporating the decision-maker's preferences about time and risk assumes great importance.

Decision analysis provides a logical framework for balancing all these considerations. It permits mathematical modeling of the decision, computational implementation of the model, and quantitative evaluation of the various courses of action. This report describes and delineates the potential of decision analysis as an aid to top management.

The Timeliness of Decision Analysis

An appropriate question is why decision analysis has only recently emerged as a discipline capable of treating the complexities of significant decision problems. The answer is found in the combination of three factors: historical circumstance, development of complementary capabilities, and the need for increased formalism.

The Computer Revolution

Despite the elaborateness of its logical foundations, decision analysis would be merely an intellectual curiosity rather than a powerful tool if the means were not available to build models and to manipulate them economically. The rapid development of the electronic computer in the past two decades has made feasible what would have been impossible only a quarter of a century ago. The availability of electronic computation is an essential condition for the growth of the decision analysis field.

The Tyranny of the Computer

A powerful tool is always subject to misuse. The widespread use of computers has led some managers to feel that they are losing rather than gaining control over the operations of their organizations. These feelings can lead to a defensive attitude toward the suggestion that computers should be included in the decision-making process.

Decision analysis can play a major role in providing the focus that management requires to control application of computers to management activities. When examined through decision analysis, the problem is not one of management information systems, but one of providing management with structured decision alternatives in which management experience, judgment, and preference have already been incorporated. Since properly applied decision analysis produces insight as well as answers, it places control in, rather than out of, the hands of the decision-maker.

The Need for Formalism

A final force in the current development of decision analysis is the trend toward professional management in present organizations. The one-man show is giving way to committees and boards, and the individual entrepreneur is becoming relatively less important. A concomitant of this change is the need for new professional managers to present evidence of more carefully reasoned and documented decisions. Even the good intuitive decision-maker will have to convince others of the logic of his decisions.

However, the need for more formalism may also be imposed from outside the organization. The nature of competition will mean that when one company in an industry capitalizes on the efficacy of decision analysis, the others will be under pressure to become more orderly in their own decision-making. To an increasing extent, good outcomes resulting from intuitive decisions will be regarded in the same light as winnings at the races—that is, as the result of luck rather than of prudent managerial practice.

The Essence of Decision Analysis

Definition of Decision

In describing decision analysis, the first step is to define a decision. In this report, a decision is considered an irrevocable allocation of resources, in the sense that it would take additional resources, perhaps prohibitive in amount, to change the allocation. Some decisions are inherently irrevocable, such as whether or not to amputate a pianist's hand; others are essentially irrevocable, such as the decision by a major company to enter a new field of endeavor.

Clearly, no one can make a decision unless he has resources to allocate. For example, a manufacturer may be concerned about whether his competition will cut prices, but unless he can change something about the way he does business, he has no decisions to make. Concern without the ability to make decisions is simply "worry." It is not unusual in practice to encounter decision problems that are really worries. Exposing a decision problem as a worry may be very helpful if it allows the resources of the decision-maker to be devoted more profitably to other concerns.

Another common phenomenon is the study, which is an investigation that does not focus on a decision. Until a decision must be made, how can the economic balance of the study be determined? For example, suppose someone requested a study of the automobile in his particular community. The person conducting the study might survey cars' weight, horsepower, displacement, braking ability, seating capacity, make, type, color, age, origin, and on and on. However, if a decision were required concerning the size of stalls in a parking facility, or the length of a highway acceleration lane, the pertinent characteristics would become clear. Further, decision analysis could even determine how extensive a survey, if any, would be economic. Thus, concentrating on a decision to be made provides a direct focus to the analysis that is achievable in no other way. Studies, like worries, are not our concern: decisions are.

The next step is to define a decision-maker: an individual who has the power to commit the resources of the organization. In some cases, the decision-maker may be an organizational entity, such as an executive committee. It is important, however, to distinguish advisory individuals or bodies from those with the power to commit the organization. Study upon study may be performed within an organization advocating or decrying a certain course of action, but until resources are committed, no decision has been made. The first step in any decision analysis is the identification of the responsible party.

The Distinction Between a Good Decision and a Good Outcome

Before there can be a formal discussion of decision analysis, the distinction between a good decision and a good outcome must be understood. A good decision is one based on the information, values, and preferences of a decision-maker. A good outcome is one that is favorably regarded by a decision-maker. It is possible to have good decisions produce either good or bad outcomes. Most persons follow logical decision procedures because they believe that these procedures, speaking loosely, produce the best chance of obtaining good outcomes.

To illustrate this point, suppose that we had agreed to serve as decision analysis consultants to a person who said that he would engage only in gambles that were weighted in his favor. Then this person informed us that he had purchased a ticket in a lottery. There were 100 tickets in the lottery, the prize was \$100, and he paid \$10 for the ticket. We demonstrate to him that with 1 chance in 100 of winning the \$100, his expected income from the ticket is only 1/100 of \$100 or \$1, so that having paid \$10 for the ticket, his expected loss on the entire prospect is \$9. Consequently, in view of this person's expressed desire to avoid unfavorable gambles, we say that he has made a bad decision.

However, the next day he receives a check for \$100 as a consequence of having won the lottery; everyone agrees that this is a good outcome for him. Yet we must report that his decision was bad in spite of the good outcome, or, perhaps better, that his outcome was good in spite of the bad decision. This would be a proper situation to be described as "lucky."

Suppose, however, that the person had paid only 10 cents for his ticket. In this case, his expected income is still \$1, but because he spent only 10 cents for the ticket, his net expected earnings are 90 cents. Consequently, we would compliment him on his good decision. Yet if no winnings check appears on the next day, the client has now experienced a bad outcome from his good decision.

The distinction between good outcomes and good decisions is especially important in maintaining a detached, professional attitude toward decision problems. Recriminations based on hindsight in the form of "Why didn't it work?" are pointless unless they reveal that available information was not used, that logic was faulty, or that the preferences of the decision-maker were not properly encoded. The proper framework for discussing the quality of decisions and outcomes is a major aid in using hindsight effectively.

Decision Analysis as a Language and a Philosophy

The decision analysis formalism serves both as a language for describing decision problems and as a philosophical guide to their solution. The existence of the language permits precision in specifying the many factors that influence a decision.

The most important feature of the language is its ability to represent the uncertainty that inevitably permeates a decision problem. The language of probability theory is used with only minor changes in terminology that reflect a subjective interpretation of probabilistic measurement. We regard probability as a state of mind rather than of things. The operational justification for this interpretation can be as simple as noting the changing odds on a sporting contest posted by gamblers as information about the event changes. As new information arrives, a new probability assignment is made. Decision analysis uses the same subjective view of probability. By so doing, statements regarding uncertainty can be much more precise. Rather than saying, "There is some chance that a bad result is likely," or an equivalent ambiguous statement, we shall be able to speak directly of the probability of a bad result. There is no need for vagueness in the language that describes uncertainty. Putting what is not known on the record is the first step to new knowledge.

Decision analysis can also make a major contribution to the understanding of decision problems by providing a language and philosophy for treating values and preferences. "Values" mean the desirability of each out-outcome; "preferences" refer to the attitudes of the decision-maker toward postponement or uncertainty in the outcomes he receives. Placing values and preferences in unambiguous terms is as unusual in current decision-making as is the use of direct probability assignments. Yet both must be done if the procedure is to be used to full advantage.

Later sections of this report describe the theory and practice of assigning probabilities, values, and preferences, but the impact of thinking in such terms can be indicated here. A most important consequence of formal thought is the spontaneous resolution of individual differences that often occurs when the protagonists can deal in unambiguous terms. Two people who differ over the best alternative may find their disagreements in the areas of probability assignment, value, or preference. Thus, two men who are equally willing to take a risk may disagree because they assign different probabilities to various outcomes; or two men who assign the same probability to the outcomes may differ in their aversion to risk. It is unlikely that the nature of the disagreement will emerge without the formal language. More likely, epithets such as "foolhardy" or "rock-bound conservative," will prevent any communication at all.

The decision analyst must play a detached role in illuminating the decision problem if he is to resolve differences. He must be impartial, never committing himself to any alternative, but rather showing how new information or changes in preference affect the desirability of available alternatives. The effectiveness of the decision analyst depends as much on his emotional detachment as on his knowledge of formal tools.

Decision analysis is a normative, rather than a descriptive, approach to decision problems. The decision analyst is not particularly interested in describing how decision-makers currently make decisions; rather he is trying to show how a person subscribing to certain logical rules would make these decisions in order to maximize attainment of his objectives. The decision procedures are derived from logic and from the desires of the decision-maker and are in this sense prescriptive.

Decision analysis is more than a language and a philosophy, but the experience of its users justifies it on this basis alone. By focusing on central issues, the approach often illuminates the best course of action in a way that makes discord evaporate.

Decision Analysis as a Logical and Quantitative Procedure

Decision analysis provides not only the philosophical foundations, but also a logical and quantitative procedure for decision-making. Since decision analysis encodes information, values, and preferences numerically, it permits quantitative evaluation of the various courses of action. Further, it documents the state of information at any stage of the problem and determines whether the gathering of further information is economically justifiable. The actual implementation of decision analysis models is typically a computer program that enables the many facets of the problem to be examined together. Most of this report will describe how the philosophy of decision analysis carries over into practice.

Delegation of Responsibility

Decision analysis provides both philosophical and operational guidelines for delegating responsibility in an organization. If we want someone to make a good decision, we must provide that individual not only with the information but also with the values and preferences that are relevant to the decision. The key principle is that the delegator must supply a subordinate decision-maker with whatever information, values, and preferences required for him to reach the same decision that the delegating individual would have reached in the same situation. While few organizations currently use decision analysis principles in handling the problem of delegation, these principles are available when needed. It is rare that an organization performs a decision analysis on one of its major decisions without simultaneously obtaining new insight into its organizational structure.

THE DECISION ANALYSIS CYCLE

Decision analysis as a procedure for analyzing a decision is described below. This procedure is not an inviolable method of attacking the problem, but is a means of ensuring that essential steps have been consciously considered.

The figure describes decision analysis in the broadest terms. The procedure is iterative and comprises three phases. The first is a deterministic phase, in which the variables affecting the decision are defined and related, values are assigned, and the importance of the variables is measured without any consideration of uncertainty.

Fig. 1—The Decision Analysis Cycle

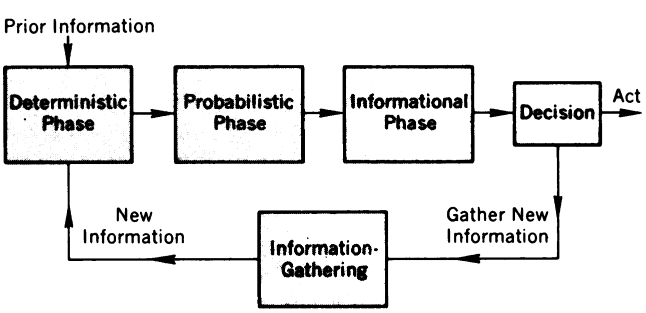

The second, or probabilistic, phase introduces probability assignments on the important variables and derives associated probability assignments on values. This phase also introduces the assignment of risk preference, which provides the best solution in the face of uncertainty.

The third, or informational, phase reviews the results of the last two phases to determine the economic value of eliminating uncertainty in each of the important variables in the problem. In some ways, this is the most important phase because it shows just what it could cost in dollars and cents not to have perfect information. A comparison of the value of information with its cost determines whether additional information should be collected.

If there are profitable further sources of information, then the decision should be to gather the information rather than to make the primary decision at this time. Thereupon will follow the design and execution of the information-gathering program, whether it be a market survey, a laboratory test, or military field trials.

The information that results from this program may change the model and the probability assignments on important variables. Therefore, the original three phases must be performed once more. However, the additional work required to incorporate the modifications should be slight and the evaluation rapid. At the decision point, it may again be profitable to gather new information and repeat the cycle or it may be more advisable to act. Eventually, the value of new analysis and information-gathering will be less than its cost, and the decision to act will then be made.

This procedure will apply to a variety of decision situations: in the commercial area, to the introduction of a new product or the change in design of an old one; in the military area, to the acquisition of a new weapon or the best defense against that of a potential enemy; in the medical area, to the selection of a medical or surgical procedure for a patient; in the social area, to the regulation and operation of public utilities; and finally, in the personal area to selection of a new car, home or career. In short, the procedure can be applied to any decision susceptible to logical analysis.

The Deterministic Phase

Descriptions of the various phases of the procedure follow beginning with the deterministic phase. The deterministic phase is essentially a systems analysis of the problem. Within this phase, efforts devoted to modeling are distinguished from efforts devoted to analysis. The elements of the phase appear in Figure 2.

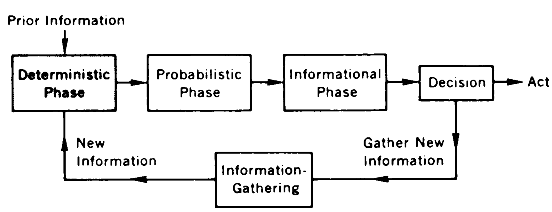

Fig. 2
The
Deterministic
Phase

MODELING:
- Bound Decision
- Identify Alternatives
- Establish Outcomes
- Select System Variables
- Create Structural Model
- Create Value Model
- Create Time Preference Model

ANALYSIS:
- Measure Sensitivity
  - to Decision Variables
  - to State Variables

Modeling

Modeling is the process of representing the various relationships of the problem in formal, mathematical terms. The first step in modeling is to bound the decision, to specify precisely just what decision must be made. This requires listing in detail the perceived alternatives. Identification of the alternatives will separate an actual decision problem from a worry.

The next step—finding new alternatives—is the most creative part of decision analysis. New alternatives can spring from radically new concepts; more often they may be careful combinations of existing alternatives. Discovering a new alternative can never make the problem less attractive to the decision-maker; it can only enhance it or leave it unchanged. Often the difficulty of a decision problem disappears when a new alternative is generated.

The next step is to specify the various outcomes that the set of alternatives could produce. These outcomes are the subsequent events that will determine the ultimate desirability of the whole issue. In a new product introduction, for example, the outcomes might be specified by sales levels and costs of production or even more simply by yearly profits. Thus, there is a certain amount of arbitrariness in what to call an outcome. For decision analysis, however, an outcome is whatever the decision-maker would like to know in retrospect to determine how the problem came out. In a military problem, the outcome could be a complicated list of casualties, destruction, and armament expenditures; in a medical problem, it could be as simple as whether or not the patient dies.

Now comes the challenging process of selecting the system variables for the analysis, which are all those variables on which the outcomes depend. We can identify the system variables by imagining that we have a crystal ball that will answer any numerical questions relative to the decision problem, except, of course, which alternative to select. We could ask it questions about the outcome variables directly, thereby making them the only system variables in the problem. But typically outcome variables are difficult to think about in advance in the real world, and so we might choose to relate the outcome variables to others that are easier to comprehend. For example, we might like to know the sales level of a new product. Or in lieu of this, we might attempt to relate the sales to our own price and quality and the competitors' price and quality, factors that we might regard as more accessible. These factors would then become system variables in the analysis.

The selection of system variables is therefore a process of successive refinement, wherein the generation of new system variables is curtailed by considering the importance of the problem and the contributions of the variables. Clearly, allocation of the national budget can economically justify the use of many more system variables than can the selection of a new car.

Once we have decided on the system variables to use in the problem, each one must be distinguished either as a variable under the decision-maker's control or as a variable determined by the environment of the problem. System variables that are under the decision-maker's control are called decision variables. The selection of an alternative in a decision problem is really the specification of the setting of the decision variables. For example, in the new product introduction problem, the product price and the size of production facilities would both be decision variables.

System variables in the problem that are determined by the environment are known as state variables. Although state variables may have a drastic effect on the outcomes, they are autonomous, beyond the control of the decision-maker. For example, in the new product introduction, the cost of a crucial raw material or the competitor's advertising level might be state variables.

We shall want to examine the effect of fluctuations in all system variables, whether decision variables or state variables. To aid in this task, the decision-maker or his surrogate must specify for each system variable a nominal value and a range of values that the variable may take on. In the case of a decision variable, the nominal value and range are determined by the decision-maker's preconceptions regarding the interesting alternatives. In the case of state variables, the nominal value and range reflect the uncertainty assigned to the variables. For convenience, we can often think of the nominal value of a state variable as its expected value in the mathematical sense and of the range as the 10th percentile and 90th percentile points of its probability distribution.

Selecting system variables and setting nominal values and ranges require extensive consultation between the decision-maker and the decision analyst. At this stage, it is better to err by including a variable that will later prove to be unimportant than it is to eliminate a variable prematurely.

The next step is to specify the relationships among the system variables. This is the heart of the modeling process—i.e., creating a structural model that captures the essential interdependencies of the problem. This model should be expressed in the language of logic—mathematics—typically by a set of equations relating the system variables. In most decisions of professional interest, these equations will form the basis for a computer program to represent the model. The program provides rapid evaluation of model characteristics at modest cost.

Constructing a model of this type requires a certain sophistication in the process of orderly description and a facility for careful simplification. The procedure is elementary, but not trivial; straightforward, but not pedestrian.

Now the decision-maker must assign values to outcomes. Just as there was difficulty in defining an outcome, so there may be some question about the distinction between an outcome and its value. For example, in a business problem, the decision-maker may think of his future profit as both the outcome and the value associated with it. However, maintaining the generality of the formulation requires creating a distinction between the two.

To illustrate the necessity for this, consider a medical question involving the amputation of an arm. The outcomes of interest might be complete recovery, partial recovery, or death, each with or without the operation. These outcomes would describe the results but would not reveal their value. For example, if the patient were a lawyer, he might consider death by far the most serious outcome and be willing to undergo the amputation if it sufficiently reduced the probability of death. These feelings might be based on the observation that an arm is not essential to his career. To a concert pianist, however, amputation might be worse than death itself, since life without being able to play might be unbearable. Consequently, he would be rational in refusing the amputation even if this choice made his death more likely.

Although in some cases the decision can be reached as a result of ordering outcomes in terms of desirability, most problems of practical interest require a numerical (cardinal) ranking system. Therefore, assigning a value means assigning a numerical value to an outcome. Though there may be many elements of value in the outcome, the final value assignment is a single number associated with that outcome.

In commercial situations, the value assigned to an outcome will typically be some form of profit. In social and military problems, however, the value assignment is more difficult because it requires measuring the value of a human life, or a cultured life, or a healthy life in dollars and cents terms. Though these questions of evaluation may be difficult, logic demands that they be approached directly in monetary terms if monetary resources are to be allocated.

The final step in creating the deterministic model is to specify the time preference of the decision-maker. Time preference is the term used to describe the human phenomenon of impatience. Everyone wants good things to happen to him sooner rather than later. This impatience is reflected in a willingness to consume less now rather than postpone the consumption. The payment of interest on savings accounts and the collection of interest on loans are mere reflections of this phenomenon. Consequently, representing the desires of a decision-maker requires a realistic mechanism for describing his time preference, a mechanism that reduces any time stream of value to a single number called worth.

For a corporate financial decision, worth will often be simply the discounted difference between future income and expenditures using an interest rate that depends upon the relationship of the corporation to its financial environment. In the military or medical fields, worth may be more difficult to establish.

The modeling part of the deterministic phase thus progresses from the original statement of the decision problem to a formal description suitable for detailed examination by logical and computational analysis. The decision-maker's value assignments and his time preference permit rating any outcome that appears as a time stream first as a set of values in time and then as an equivalent worth.

Analysis

Analysis based on the deterministic phase centers on observing how changes in the variables affect worth. Experimentation of this type is known as sensitivity analysis; it is highly effective in refining the formulation of the problem.

The first sensitivity analysis we perform is associated with the decision variables. First, fixing all other state variables in the problem at their nominal values, we then allow one of the decision variables to traverse its assigned range and observe how worth changes. Of course, these observations are usually carried out by computer program. If we find that a particular decision variable has a major effect, then we know that we were correct in including it in the original formulation. But if a decision variable has little or no effect, we are justified in considering its removal as a decision variable. If reflection reveals that the latter is the case, we would say that we have eliminated an impotent decision variable. For example, the time of introduction of a new product might seem to be a decision variable of major importance, but because of the combined effects of competitive reaction and the gaining of production experience, it might turn out to have very little effect. The timing of entry would then be an impotent variable.

Next, we perform sensitivity analyses on the state variables, which are uncertain and over which the decision-maker has no control. With all other system variables at their nominal values, we observe the change in worth while sweeping one state variable over its range. If a state variable has a major effect, then the uncertainty in the variable deserves special attention. Such variables are called aleatory variables to emphasize their uncertainty.

If, however, varying a state variable over its range produces only a minor change in worth, then that variable might well be fixed at its nominal value. In this case, we say that the state variable has become a fixated variable. A state variable may become fixated either because it has an important influence on the worth per unit of its range, but an extremely small range, or because it has little influence on the worth per unit of its range, even though it has a broad range.

There is no reason to conclude that a fixated variable is unimportant in an absolute sense. For example, the corporate tax rate may be a fixated variable in a problem because no change in it is anticipated within the time period under consideration. Yet it is possible that an unforeseen large change in this rate could change a favorable venture into an unfavorable one.

Although sensitivity analysis has been described as if it concerns only changes in one variable at a time, some of the most interesting sensitivity results are often observed when there are simultaneous changes in state variables. Since the possibilities of changing state variables jointly grows rapidly with the number of state variables, an important matter of judgment for the decision analyst is to determine the amount of simultaneous sensitivity analysis that is economic.

The Probabilistic Phase

The net result of the deterministic sensitivity analysis on the autonomous state variables is to divide them into aleatory and fixated classes. The probabilistic phase determines the uncertainty in value and worth due to the aleatory variables. The phase will be divided into steps of modeling and analysis; Figure 3 illustrates its internal structure.

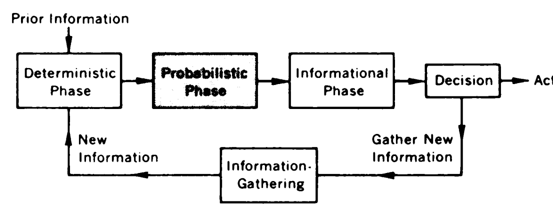

Fig. 3
The
Probabilistic
Phase

MODELING:
- **Encode Uncertainty on Aleatory Variables**
- **Encode Risk Preference**

ANALYSIS:
- **Develop Worth Lotteries and Certainty Equivalent**
- **Measure Stochastic Sensitivity**
- **Measure Risk Sensitivity**

Modeling Probability Distributions

The first modeling step in the probabilistic phase is the assignment of probability distributions to the aleatory variables. Either the decision-maker or someone he designates must assign the probability that each aleatory variable will exceed any given value. If any set of aleatory variables is dependent, in the sense that knowledge of one would provide information about the others, then the probability assignments on any one variable must be conditional on the values of the others. Gathering these assignments amounts to asking such questions as, "What are the odds that sales will exceed 10 million units in the first year?" (See section entitled "Encoding Knowledge and Preferences.") Strange as such questions may be in the current business world, they could be the standard executive language of tomorrow.

Analysis

With knowledge from the deterministic phase of how the worth depends on the state variables and assigned probability distributions on the aleatory variables, it is a straightforward calculation to determine the probability distribution of worth for any setting of the decision variables; this probability distribution is the "worth lottery." The worth lottery describes the uncertainty in worth that results from the probability assignments to the aleatory variables for any given alternative (setting of decision variables.) Of course, the values of the fixated variables are never changed.

To select a course of action, the analyst could generate a worth lottery for each alternative and then select the one that is more desirable. But how would he know which worth lottery is most desirable to the decision-maker?

One important principle that allows judging one worth lottery as being better than another is that of stochastic dominance, which is illustrated in Figure 4. Part *A* of this figure shows the worth lottery for two alternatives in both probability densities and excess probability distribution forms. The excess probability distribution, or excess distribution, is the probability that the variable will exceed any given value plotted as a function of that value. Its height at any point is the area under the probability density function to the right of that point. Comparison of the excess distributions for the two alternatives reveals that, for any value of X, there is a higher probability that alternative 2 will produce a worth in excess of that X than will alternative 1. Consequently, a decision-maker preferring more worth to less would prefer alternative 2. If alternative A has an excess distribution that is at least as great as that of alternative B at any point and greater than B at at least one point, alternative A stochastically dominates alternative B. If stochastic dominance exists between two competing alternatives, there is no need to inquire into the risk preference of the decision-maker, who rationally must rule out the stochastically dominated alternatives.

Fig. 4

Part A—Stochastic Dominance

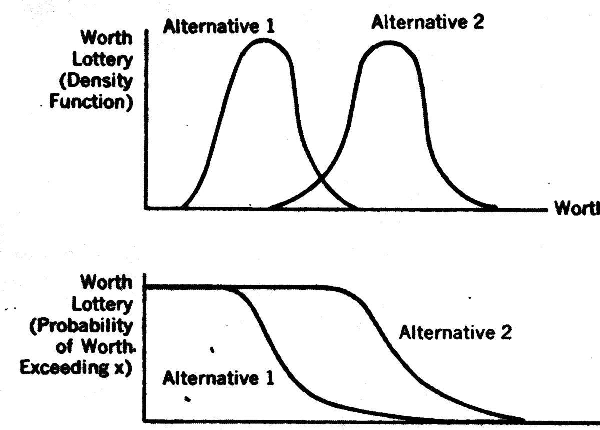

Part B—Lack of Stochastic Dominance

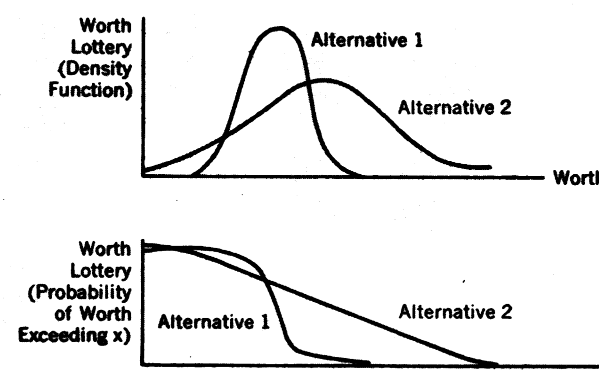

Part B of Figure 4 illustrates a case in which stochastic dominance does not exist. The excess distributions on worth for the two alternatives cross. If the decision-maker wants to maximize his chance of receiving at least a small amount of worth, he would prefer alternative 1; if he wants to maximize his chance of receiving at least a large amount of worth, he would prefer alternative 2. In situations like this, where stochastic dominance does not apply, the risk preference of the decision-maker must be encoded formally, as shown below.

Just because alternative A stochastically dominates alternative B does not mean that the decision-maker will necessarily achieve a higher worth by following alternative A. For example, if alternative A produces worths of five to 15 with equal probability and alternative B produces worths of zero and ten with equal probability, then A stochastically dominates B. Yet it is possible that A will produce a worth of five while B will produce a worth of ten. However, not knowing how the lottery will turn out, the rational man would prefer alternative A.

Modeling Risk Preference

If stochastic dominance has not determined the best alternative, the analyst must turn to the question of risk preference. To demonstrate that most individuals are averse to risk, it is only necessary to note that few, if any, are willing to toss a coin, double or nothing, for a year's salary. Organizations typically act in the same way. A realistic analysis of decisions requires capturing this aversion to risk in the formal model.

Fortunately, if the decision-maker agrees to a set of axioms about risk taking (to be described in the following section), his risk preference can be represented by a utility curve like that shown in Figure 5. This curve assigns a utility to any value of worth. As a consequence of the risk preference axioms, the decision-maker's rating of any worth lottery can be computed by multiplying the utility of any possible worth in the lottery by the probability of that worth and then summing over all possible worths. This rating is called the expected utility of the worth lottery.

If one worth lottery has a higher expected

Fig. 5
A Typical Utility Curve

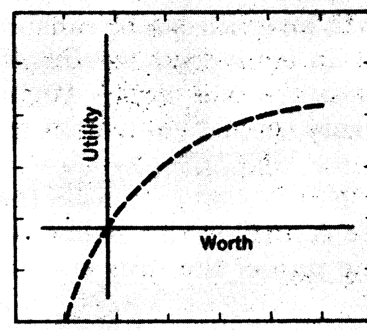

utility than another, then it must be preferred by the decision-maker if he is to remain consistent with the axioms. The analyst is not telling the decision-maker which worth lottery he should prefer but only pointing out to him a way to be consistent with a very reasonable set of properties he would like his preferences to enjoy.

Thus, the utility curve provides a practical method of incorporating risk preference into the model. When faced with a choice between two alternatives whose worth lotteries do not exhibit stochastic dominance, the analyst computes the expected utility of each and chooses the one with the higher expected utility.

Although the expected utility rating does serve to make the choice between alternatives, its numerical value has no particular intuitive meaning. Therefore, after computing the expected utility of a worth lottery, the analyst often returns to the utility curve to see what worth corresponds to this expected utility; we call this quantity the certain equivalent worth of the worth lottery. The name arises as follows: if another worth lottery produced the certain equivalent worth with probability one, then it and the original lottery would have the same expected utilities and hence would be equally preferred by the decision-maker. Consequently, the certain equivalent worth of any worth lottery is the amount of worth received for certain, so that the decision-maker would be indifferent between receiving this worth and participating in the lottery. Since almost all utility curves show that utility increases as worth increases, worth lotteries can be ranked in terms of their certain equivalent worths. The best alternative is the one whose worth lottery has the highest certain equivalent worth.

Analysis

In returning to the analysis of the probabilistic phase, the first step is to compute the certain equivalent worth of each of the alternatives. Since the best decision would be the alternative with the highest certain equivalent worth, the decision probably could be considered solved at this point. The careful analyst, however, will examine the properties of the model to establish its validity and so would not stop here. The introduction of risk preference is another point at which to check the sensitivity of the problem. For example, by setting all decision variables but one to their nominal values and then sweeping this one decision variable through its range, the analyst may find that although this variation changes the worth lottery it does not significantly change the certain equivalent worth. This result would indicate that the decision variable could be fixed at its nominal value.

Aleatory variables receive the same sensitivity analysis by setting one of them equal to a trial value within the range and then allowing the others to have the appropriate conditional joint probability distribution. When the decision variables are given their nominal values, the program will produce a worth lottery and hence a certain equivalent worth for the trial value. Sweeping the trial value from one end of its range to the other shows how much certain equivalent worth is changed. If the change is small, there is evidence that the particular aleatory variable may be changed to a fixated variable. We call this procedure measurement of the stochastic sensitivity of a variable. It is possible that an aleatory variable showing a large deterministic sensitivity could reveal only a small stochastic sensitivity and vice versa. Consequently, any decisions to remove variables from aleatory status on the basis of deterministic sensitivity might well be reviewed at this time by measurement of stochastic sensitivity.

As in the case of deterministic sensitivity, we can measure the stochastic sensitivity of many variables, simultaneously. Once more, the decision analyst must judge how far it is profitable to proceed. Measurement of stochastic sensitivity is a powerful tool for locating the important variables of the problem.

There is one other form of sensitivity analysis available at this point: risk sensitivity. In some cases, it is possible to characterize the utility curve by a single number—the risk aversion constant (just when this is possible will be discussed later). However, when the risk aversion constant is applicable we can interpret it as a direct measure of a decision-maker's willingness to accept a risk. An individual with a small risk aversion constant is quite willing to engage in a fair gamble; he has a tolerant attitude toward risk. As his risk aversion constant increases, he becomes more and more unwilling to participate. If two men share responsibility for a decision problem, the less risk tolerant will assign a lower certain equivalent worth for any given worth lottery than will the other. Perhaps, however, when the certain equivalent worths are computed for all alternatives for both men, the ranking of certain equivalent worths might be the same for both, or at least the same alternative would appear at the top of both lists. Then there would hardly be any point in their arguing over the desirable extent of risk aversion and a possible source of controversy would have been eliminated.

The measurement of risk sensitivity determines how the certain equivalent worths of the most favorable alternatives depend on the risk aversion constant. The issue of risk aversion can often be quickly resolved.

The problem structure, the set of alternatives generated, the probability assignment to aleatory variables, the value assessments, the statement of time preference, and the specification of risk preference combine to indicate

Fig. 6—The Decision Analysis Hierarchy

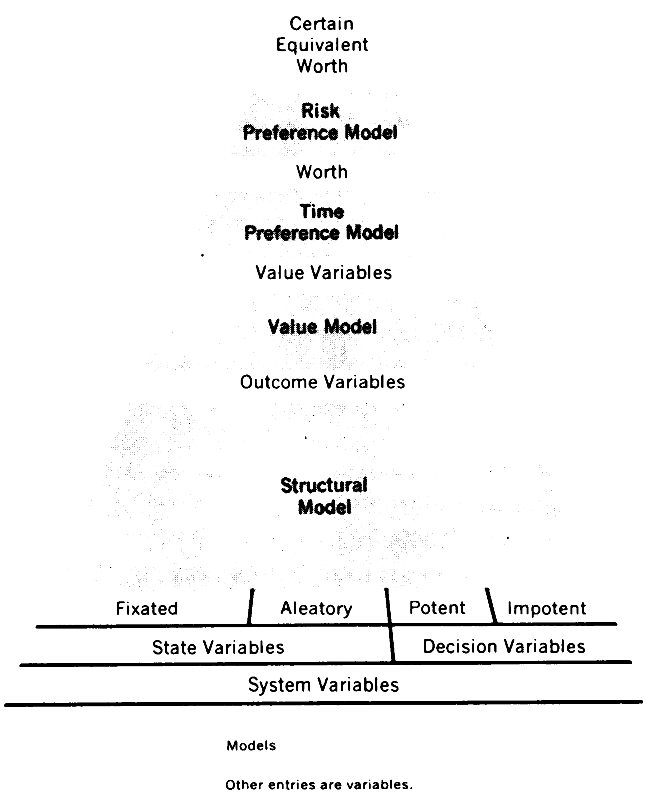

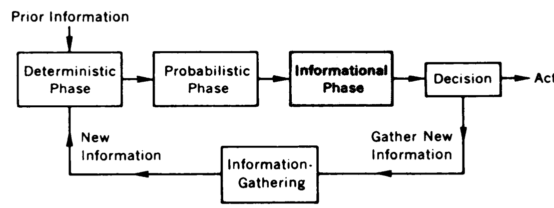

Fig. 7
The Informational Phase

ANALYSIS:

- **Measure Economic Sensitivity (Determine Value of Eliminating Uncertainty in Aleatory Variables)**

MODELING:

- **Explore Feasibility of Information Gathering**

the best alternative in the problem. The overall procedure is illustrated by the decision analysis pyramid in Figure 6. However, it still may be best to obtain more information rather than to act. This determination is made in the third phase, as described below.

The Informational Phase

The informational phase is devoted to finding out whether it is worthwhile to engage in a possibly expensive information-gathering activity before making a decision. It is, in the broadest sense, an experimental design procedure from which one very possible result is the decision to perform no experiment at all. Figure 7 shows the steps in the phase.

Analysis

The fundamental idea in the informational phase is that of placing a monetary value on additional information. A key concept in approaching this value is that of clairvoyance. Suppose someone exists who knows in advance just what value a particular aleatory variable would assume in the decision problem—a clairvoyant. How much should the decision-maker be willing to pay him for his services?

To answer this question, recall that the discussion of stochastic sensitivity described how to compute the certain equivalent worth given that an aleatory variable took on a value s. In that procedure, the decision variables were set equal to their best values from the probabilistic phase. Suppose now that we engage the clairvoyant at a cost k, and then he tells us that the aleatory variable will take on the value s. First, we would set the decision variables to take best advantage of this information. However, since the other aleatory variables are still uncertain, they would be described by the appropriate distributions, given the available information. The computer program would then determine the expected utility of the entire decision problem including the payment to the clairvoyant, all conditional on his reporting s.

Before engaging the clairvoyant, however, the probability to be assigned to his reporting s as the value of the particular aleatory variable is described by the probability distribution showing the current state of knowledge on this variable. Consequently, we obtain the expected utility of purchasing his information on the variable at a cost k by multiplying the expected utility of the information given that he reports s and costs k, by the current probability that he will report s and then summing over all values of s. The analyst uses the current probability in this calculation because if the clairvoyant is reliable, the chance of his reporting that the variable falls in any range is just the chance that it will fall in that range.

Knowing the expected utility of purchasing the information from the clairvoyant at a cost of k, we can gradually increase k from zero until the expected utility of purchasing the information is just equal to the expected utility of proceeding with the decision without clairvoyant information. The value of k that establishes this equivalence is the value of clairvoyance on the aleatory variable.

The value of clairvoyance on an aleatory variable represents an upper bound on the payment for any experimental program designed to provide information on this variable, for no such program could be worth more than clairvoyance. The actual existence of a clairvoyant is not material to this discussion; he is merely a construct to guide our thinking.

We call the process of measuring the value of clairvoyance the measurement of economic sensitivity. If any aleatory variable exhibits high economic sensitivity, it is a prime candidate for an information-gathering program. It is possible, however, for a variable to have a high stochastic sensitivity and a low economic sensitivity because the available alternatives cannot take advantage of the information received about the variable. To determine the importance of joint information, the analyst can measure the value of clairvoyance on more than one variable at a time.

The actual information-gathering programs available will seldom provide perfect information, so they will be less valuable than clairvoyance. Extension of the discussion of clairvoyance shows how their value can be measured. Whereas the clairvoyant reported a particular value s for an aleatory variable, a typical experimental program will provide only a new probability distribution for the aleatory variable. The analyst would then determine the best decision, given this new information, and compute the expected utility of the decision problem. He would next multiply the expected utility by the probability that the exerimental program would come out in this way and then sum over all possible outcomes of the experimental program. The result would be the expected utility of the experimental program at a given cost. The cost that would make the expected utility just equal to the expected utility of the problem without the experimental program would be the value of the experimental program. If the value is positive, it represents the maximum that one should pay for the program. If the value is negative, it means that the experimental program is expected to be unprofitable. Consequently, even though it would provide useful information, it would not be conducted.

Modeling

At this stage, the decision-maker and the analyst must identify the relevant information-gathering alternatives, from surveys to laboratory programs, and find which, if any, are expected to make a profitable contribution to the decision problem. In considering alternatives, they must take into account any deleterious effect of delay in making the primary decision. When the preferred information-gathering program is performed, it will lead, at least, to new probability assignments on the aleatory variables; it might also result in changing the basic structure of the model. When all changes that have been implied by the outcome of the experimental program are incorporated into the model, the deterministic and probabilistic phases are repeated to check sensitivities. Finally, the informational phase determines whether further information-gathering is profitable. At some point, further information will cost more than it is worth, and the alternative that currently has the highest certainty equivalent will be selected for implementation.

The iterative decision analysis described above is not intended to fit any particular situation exactly but, rather, all situations conceptually. A discussion follows on two procedures required to carry out the analysis: encoding knowledge and preferences.

ENCODING KNOWLEDGE AND PREFERENCES

Encoding Knowledge as Probability Distributions

Perhaps the single most unusual aspect of decision analysis is its treatment of uncertainty. Since uncertainty is the central problem in decision-making, it is essential to understand the conceptual and logical foundations of the approach to this issue.

The Importance of Uncertainty

The importance of uncertainty is revealed by the realization that decisions in situations where there is no random element can usually be made with little difficulty. Only when uncertainty exists about which outcome will occur is there a real decision problem.

For example, suppose that we are planning to take a trip tomorrow and that bad weather is forecast. We have the choice of flying or of taking a train. If a clairvoyant told us the consequences of each of these acts, then our decision would be very simple. Thus, if he said that the train would depart at 9:13 A.M. and arrive at 5:43 P.M. and if he described in detail the nature of the train accommodations, the dining car, and the people whom we would meet as traveling companions, then we would have a very clear idea of what taking the train implied. If he further specified that the plane would leave 2 hours late and arrive $2\frac{1}{2}$ hours late, stated that the flight would be especially bumpy during a certain portion of the trip, and described the meals that would be served and the acquaintances we would meet, then the flying alternative would be described as well.

Most of us would have little trouble in making a decision about our means of travel when we considered these carefully specified outcomes in terms of our tastes and desires. The decision problem is difficult because of the uncertainty of departure and arrival times and, in the case of the plane, even whether the trip would be possible at all. The factors of personal convenience and pleasure will be more or less important depending upon the urgency of the trip and, consequently, so will the uncertainties in these factors. Thus we cannot make a meaningful study of decision-making unless we understand how to deal with uncertainty. Of course, in the problems that are of major practical interest to the decision analyst, the treatment of uncertainty is even more pressing.

It is possible to show that the only consistent theory of uncertainty is the theory of probability invented 300 years ago and studied seriously by mathematicians the world over. This theory of probability is the only one that has the following important property: the likelihood of any event's following the presentation of a sequence of points of data does not depend upon the order in which those data are presented. So fundamental is this property that many would use it as a defining basis for the theory.

The Subjective Interpretation of Probability

A reasonable question is: If probability is so essential to decision-making, why hasn't its importance been more widely appreciated until now? The answer is that many users of probability theory (but certainly not the original developers) considered probabilities to be physical parameters of objects, such as weight, volume, or hardness. For example, there was much mention of "fair" coins and "fair" dice, with the underlying notion that the probability of events associated with these objects could be measured in the real world.

For the past 15 years, however, an important minority of experts on the subject have been advancing the view that probabilities measure a person's state of knowledge about phenomena rather than the phenomena themselves. They would say, for example, that when someone describes a coin as "fair" he really means that on the basis of all evidence presented to him he has no reason for asserting that the coin is more likely to fall heads than tails. This view is modern, but not a product of modern times. It was studied clearly and convincingly 200 years ago but remained buried for a long time.

An example illustrating this view of probability follows: An astronaut is about to be fired into space on a globe-circling mission. As he is strapping himself into his capsule on top of a gleaming rocket, he asks the launch supervisor, "By the way, what's the reliability of this rocket?" The launch supervisor replies "Ninety nine percent—we expect only one rocket in one hundred to fail." The astronaut is reassured but still has some doubts about the success of his mission. He asks, "Are these rockets around the edge of the field the same type as the one I'm sitting on?" The supervisor replies, "They're identical." The astronaut suggests, "Let's shoot up a few just to give me some courage."

The rocket is fitted with a dummy payload, prepared for launching, and fired. It falls in the ocean, a complete failure. The supervisor comments, "Unlucky break, let's try another." Unfortunately, that one also fails by exploding in mid-air. A third is tried with disastrous results as it disintegrates on its pad. By this time, the astronaut has probably handed in his resignation and headed home. Nothing could convince him that the reliability of his rocket is still 99%.

But, in reality, what has changed? His rocket is physically unaffected by the failure of the other rockets. Its guidance system, rocket engine, and life support system are all exactly the same as they were before the other tests. If probability were a state of things, then the reliability of his rocket should still be 0.99. But, of course, it is not. After observing the failure of the first rocket, he might have evaluated the reliability of his rocket at, say, 0.90; after the second failure, at 0.70; and finally after the third failure, at perhaps 0.30. What happened was that his state of knowledge of his own rocket was influenced by what happened to its sister ships, and therefore his estimate of its reliability must decrease. His final view of its reliability is so low that he does not choose to risk his life.

The view of probability as a state of things is just not tenable. Probability should be considered as the reading of a kind of mental thermometer that measures uncertainty rather than temperature. The reading goes up if, as data accumulate, it tends to increase the likelihood of the event under consideration. The reading of 1 corresponds to certainty that the event will occur, the reading of 0 to certainty that it will not occur. The inferential theory of probability is concerned with the question of how the reading ought to fluctuate in the face of new data.

Encoding Experience

Most persons would agree that it would be unwise to make a decision without considering all available knowledge before acting. If someone were offered an opportunity to participate in a game of chance by his best friend, by a tramp, and by a business associate, he would generally have different feelings about the fairness of the game in each case. A major problem is how to encode the knowledge he has in a usable form. This problem is solved by the observation that probability is the appropriate way to measure his uncertainty.

All prior experience must be used in assessing probabilities. The difficulty in encoding prior knowledge as probability is that the prior information available may range in form from a strong belief that results from many years of experience to a vague feeling that arises from a few haphazard observations. Yet there is probably not a person who had no information about an event that was important to him. People who start out saying that they have no idea about what is going to happen can always, when pressed, provide probability assignments that show considerable information about the event in question. The problem of those who would aid decision-makers is to make the process of assigning probabilities as simple, efficient, and accurate as possible.

The Practical Encoding of Knowledge

In the probabilistic phase of decision analysis, we face the problem of encoding the uncertainty in each of the aleatory variables. In organizational decision-making, prior probability distributions (or priors) should be assigned by the people within the organization who are most knowledgeable about each state variable. Thus, the priors on engineering variables will typically be assigned by the engineering department; on marketing variables, by the marketing department; and so on. However, since each case is an attempt to encode a probability distribution that reflects a state of mind and since most individuals have real difficulty in thinking about uncertainty, the method of extracting the priors is extremely important. As people participate in the prior-gathering process, their attitudes are indicated successively by: "This is ridiculous." "It can't be done." "I have told you what you want to know, but it doesn't mean anything." "Yes, it seems to reflect the way I feel." And "Why doesn't everybody do this?" In gathering the information, the analyst must be careful to overcome the defenses the individual develops as a result of being asked for estimates that are often a combination of targets, wishful thinking, and expectations. The biggest difficulty is in conveying to the man that the analyst is interested in his state of knowledge and not in measuring him or setting a goal for him.

If the subject has some experience with probability, he often attempts to make all his priors look like normal distributions, a characteristic known as "bell-shaped" thinking. Although normal distributions are appropriate priors in some circumstances, they should not become foregone conclusions.

Experience has shown certain procedures to be effective in this almost psychoanalytic process of prior measurement. One procedure is to make the measurement in a private interview to eliminate group pressure and to overcome the vague notions that most people exhibit about probabilistic matters. Unless the subjects are already experienced in decision analysis, the distribution of forms on which they are supposed to draw their priors has proved worse than useless.

The interview begins with such questions as "What are the chances that x will exceed ten?" This approach is taken because people seem much more comfortable in assigning probabilities to events than they are in sketching a probability density function. The interviewer also skips around, asking the probability that x will be "greater than 50," "less than ten," "greater than 30," often asking the same question again later in the interview. The replies are recorded out of the view of the subject so as to frustrate any attempt at forced consistency on his part. As the interview proceeds, the subject often considers the questions with greater and greater care, so that his answers toward the end of the interview may represent his feelings much better than did his initial answers.

The interviewer can change the form of the questions by asking the subject to divide the possible values of an aleatory variable into n intervals of equal probability. The answers to all these questions enable the analyst to draw the excess probability distribution for the aleatory variable, a form of representation that seems easy to convey to people without formal probabilistic training.

The result of the interview must be a prior that the subject is willing to live with, regardless of whether it will describe a lottery on who buys coffee or on the disposal of his life savings. The analyst can test the prior by comparing it with known probabilistic mechanisms. For example, if the subject says that some aleatory variable x is equally likely to be less or greater than a, then he should be indifferent about whether he is paid \$100 if x exceeds a or if he can call the toss of a coin. If he is not indifferent, then he must change a until he is. The end result of such questions is to produce a prior that the subject is not tempted to change in any way. Although the prior-gathering process is not cheap, the analyst need perform it only on the aleatory variables.

In cases where the interview procedure is not appropriate, the analyst can often obtain a satisfactory prior by drawing one himself and then letting the subject change it until the subject is satisfied. This technique may also be useful as an educational device in preparation for the interview.

If two or more aleatory variables are dependent, then the procedure requires priors that reflect the dependencies. The technique of prior gathering is generally the same but somewhat more involved. Since the treating of joint variables is a source of expense, the analyst should formulate the problem so as to avoid them whenever possible.

An Actual Probability Assessment

Figure 8 illustrates prior-gathering. The decision in a major problem was thought to depend primarily on the average lifetime of a new material. Since the material had never been made and test results would not be available until three years after the decision was required, it was necessary to encode how much knowledge the company now had con-

Fig. 8—Priors on Material Lifetime

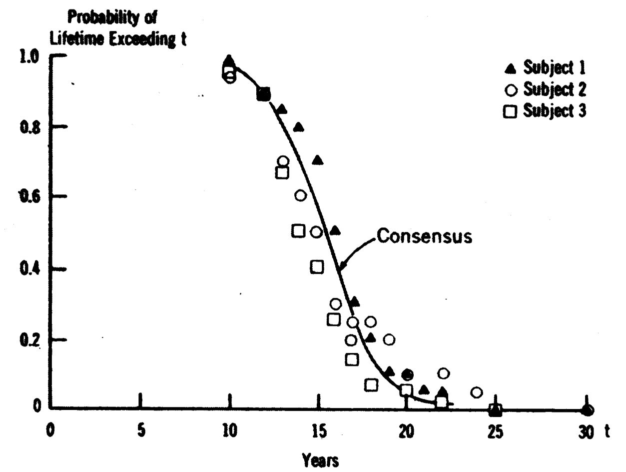

cerning the life of the material. This knowledge resided in three professional metallurgists who were experts in that field of technology. These men were interviewed separately according to the principles described. They produced the points labeled "Subjects 1, 2, and 3" in the figure. These results have several interesting features. For example, for $t = 17$, Subject 2 assigned probabilities of 0.2 and 0.25 at various points in the interview. On the whole, however, the subjects were remarkably consistent in their assignments. Subject 3 was more pessimistic about the lifetime than was Subject 1.

Upon conclusion of the interviews, the three subjects were brought together, shown the results, and a vigorous discussion took place. Subjects 1 and 3 each brought forth information of which the other two members of the group were unaware. As the result of this information exchange, the three subjects drew the consensus curve—each said that this curve represented the state of information about the material's life at the end of the meeting. Later, their supervisor said he understood their position on the new material for the first time.

It has been suggested that the proper way to reconcile divergent priors is to assign weights to each, multiply, and add, but this experiment is convincing evidence that any such mechanistic procedure misses the point. Divergent priors are an excellent indicator of divergent states of information. The experience just described not only produced the company's present encoding of uncertainty about the material's lifetime, but at the same time encouraged and effected the exchange of information within the group.

Encoding New Information

Following the encoding of the original information about an aleatory variable by means of a prior probability distribution, or about an event by the assignment of a probability, the question naturally arises as to how these probability assignments should be changed in the light of new information. The answer to this question was provided by Bayes in 1763; it is most easily introduced by considering the case of an event. Suppose that we have assigned some probability $p(A)$ to an event A's occurring and that another event B is statistically related to A. We describe this relationship by a conditional probability of B given A, $p(B|A)$, the probability of B if A occurs; assign this probability also. Now we are told that B has, in fact, occurred. How does this change the probability that A has occurred; in other words, what is the probability of A given B,$p(A|B)$?

Bayes showed that to be logical in this situation, the probability of A given B,$p(A|B)$, must be proportional to the probability of A, $p(A)$, and the probability of B given A, $p(B|A)$. This relationship is expressed as $p(A|B)$ is proportional to $p(A)$ times $p(B|A)$.

The important thing to remember is that any posterior (after new information) probability assignment to an event is proportional to the product of the prior probability assignment and the probability of the new information given that the event in question occurred. The same idea carries over in the much more complicated situations encountered in practice.

Thus, Bayes' interpretation shows how new information must be logically combined with original feelings. Subjective probability assignments are required both in describing the prior information and also in specifying how the new information is related to it. In fact, as already mentioned, Bayes' interpretation is the only method of data processing that ensures that the final state of information will be the same regardless of the order of data presentation.

Encoding Values and Preferences

The other subjective issue that arises in decision analysis is the encoding of values and preferences. It seems just as difficult to obtain an accurate measurement of desires as of information.

The value issue penetrates the core of the decision problem. Whether personal or organizational, the decision will ultimately depend on how values are assigned. If each alternative could produce only a single outcome, it would only be necessary to rank the outcomes in value and then choose the alternative whose outcome was highest in value. However, typically each alternative can produce many possible outcomes, outcomes that are distributed in time and also subject to uncertainty. Consequently, most real decision problems require numerical measures of value and of time and risk preference.

Measuring Value

The application of logic to any decision problem requires as one of its fundamental steps the construction of a value function, a scale of values that specifies the preference of the decision-maker for one outcome compared with another. We can think of the problem as analogous to the one we face if we have someone buy a car for us: We must tell our agent what features of the car are important to us and to what extent. How do we value performance relative to comfort, appearance relative to economy of operation, or other ratings?

To construct a value function in the car purchase problem, we can tell our agent the dollar value we assign to each component of a car's value. We might say, for example, that given our usage characteristics, a car that runs 18 miles to a gallon of gas is worth $40 a year more to us than a car that runs only 15 miles and that foam rubber seats are worth $50 more to us than ordinary seats. When we had similarly specified the dollar value of all the possible features of a car, including those whose values might not be additive, our agent would be able to go into the marketplace, determine the value and price of every offered car, and return with the most profitable car for us (which might, of course, be no car at all). In following this philosophy, we do not care if, in fact, there are any cars for sale that have all or any part of the features that we have valued. The establishment of the value function depends remotely, if at all, on the spectrum of cars available.

The main role of the value function is to serve as a framework of discussion for preferences. The value function encodes preferences consistently; it does not assign them. Consequently, the decision-maker or decision analyst can insert alternative value specifications to determine sensitivity of decisions to changes in value function. The process of assigning values will naturally be iterative, with components of value being added or eliminated as understanding of the problem grows.

A question that arises is, "Who should set the values?" In a corporate problem, to what extent do the values derive from management, stockholders, employees, customers, and the public? The process of constructing a value function brings into the open questions that have been avoided since the development of the corporate structure.

Establishing Time Preference

The general tendency of people and organizations is to value outcomes received sooner more highly than outcomes received later. In an organization, this phenomenon usually occurs in connection with a time stream of profit. Time streams that show a greater share of their returns in earlier time periods are generally preferred.

A number of concepts have arisen to cope with time preference in corporations. To illustrate these concepts, let $x(n)$ be the cash flow in year n in the future, positive or negative, where $n = 0$ is the beginning of the present year, $n = 1$ next year, and so on. A positive cash flow indicates that income exceeds expenditures, a negative cash flow implies the reverse. Negative cash flows will usually occur in the early years of the project.

The most elementary approach, the payback period method, rests on the assumption that the cash flow will be negative in early periods and will then become and remain positive for the balance of the project. The payback period is the number of the period in which cumulative cash flow becomes positive.

The payback period came into common use when projects were typically investments in capital equipment, investments characterized by a high initial outlay gradually returned in the course of time. However, only a few modern investments have such a simple structure. The project may contain several interspersed periods of investment and return. There would seem to be little justification for use of the payback period in modern corporate decision-making.

The idea of internal rate of return was introduced as a more sophisticated time preference measure. The internal rate of return is derived from the present value of the project, defined by

$$PV(i) = x(0) + x(1)\left(\frac{1}{1+i}\right) + x(2)\left(\frac{1}{1+i}\right)^2 + \cdots$$

where i is interpreted as an annual interest rate for funds connected with the project. The internal rate of return is the value of i that makes the present value equal to zero; in other words, the solution of the equation

$$PV(i) = 0.$$

A justification offered for the use of internal rate of return is that application of the method to an investment that pays a fixed interest rate, like a bond or a bank deposit, produces an internal rate of return equal to the actual interest rate. Although this property is satisfying, it turns out to be insufficient justification for the method. One defect, for example, is that more than one interest rate may satisfy the equation; that is, it is possible for an investment to have two internal rates of return, such as 8% and 10%. In fact, it can have as many as the number of cash flows in the project minus one. A further criticism of the method is that it purports to provide a measure of the desirability of an investment that is independent of other opportunities and of the financial environment of the firm. Although meticulous use of internal rate of return methods can lead to appropriate time preference orderings, computing the present value of projects establishes the same ordering directly, without the disadvantages of internal rate of return. Furthermore, present value provides a measure of an investment such that the bigger the number, the better the investment. The question that arises is what interest rate i to use in the computation.

Much misunderstanding exists about the implications of choosing an interest rate. Some firms use interest rates like 20% or 25% in the belief that this will maintain profitability. Yet at the same time they find that they are actually investing most of their available capital in bank accounts. The overall earnings on capital investment will therefore be rather low. The general question of selecting i is too complicated to treat here, but the fundamental consideration is the relationship of the firm to its financial environment.

There is a cogent logical argument for the use of present value. If a decision-maker believes certain axioms regarding time streams —axioms that capture such human characteristics as greediness and impatience—then the time preference of the decision-maker for cash streams that are certain must be characterized by the present value corresponding to some interest rate. Furthermore, if a bank is willing to receive and disburse money at some interest rate, then, for consistency, the decision-maker must use this bank interest rate as his own interest rate in the calculation. Present value is therefore a well-founded criterion for time preference.

In this discussion of time preference, there has been no uncertainty in the value of cash streams. Undoubtedly, it was the existence of uncertainty that made payback periods and artificially high interest rate criteria seem more logical than they in fact are. Such procedures confuse the issues of time and risk preference by attempting to describe risk preference as a requirement for even greater rapidity of return. Decision analysis requires a clear distinction between the time and risk preference aspects of decision-making.

Establishing Risk Preference

The phenomenon of risk preference was discussed in connection with the proposition of tossing a coin, double or nothing, for next year's salary: most people will not play. However, suppose they were offered some fraction of next year's salary as an inducement to play. If this fraction is zero, there is no inducement, and they will refuse. If the fraction is one, they have nothing to lose by playing and they have a .5 probability of ending up with three times next year's salary; clearly, only those with strange motivations would refuse. In experiments on groups of professional men, the fraction required to induce them to play varies from about 60% to 99%, depending on their financial obligations. Obviously, the foot-loose bachelor has a different attitude than does the married man with serious illness in the family.

The characteristic measured in this experiment is risk aversion. Few persons are indifferent to risk—i.e., willing to engage in a fair gamble. Fewer still prefer risk—i.e., willing to engage in the kind of gambles that are unfair, such as those offered at professional gambling establishments. When considering sums that are significant with respect to their financial strength, most individuals and corporations are risk-averse.

A risk-averse decision-maker is willing to forego some expected value in order to be protected from the possibilities of poor outcomes. For example, a man buys life, accident, and liability insurance because he is risk-averse. These policies are unfair in the sense that they have a negative expected value computed as the difference between the premium and the expected loss. It is just this negative expected value that becomes the insurance company's profit from operations. Customers are willing to pay for this service because of their extreme aversion to large losses.

A logical way to treat the problem of risk aversion is to begin with the idea of a lottery. A lottery is a technical term that refers to a set of prizes or prospects with probabilities attached. Thus, tossing a coin for next year's salary is a lottery and so is buying a life insurance policy. The axioms that the decision-maker must satisfy to use the theory are:

- Given any two prizes in a lottery, he must be able to state which he prefers or whether he is indifferent between them. His preferences must be transitive: if he prefers prize A to B and prize B to C, he must also prefer A to C.
- If he prefers A to B and B to C, he must be indifferent to receiving B for certain or participating in a lottery with A and C as prizes for some probability of winning A.
- If he prefers A to B, then given the choice of two lotteries that both have prizes A and B, he will prefer the one with the higher probability of winning A.
- He treats as equivalent all lotteries with the same probabilities of achieving the same prizes, regardless of whether the prizes are won in one drawing, or as the result of several drawings that take place at the same time.

It is possible to show that an individual who wants to act in accordance with these axioms possesses a utility function that has two important properties. First, he can compute his utility for any lottery by computing the utility of each prize, multiplying by the probability of that prize, and then summing over all prizes. Second, if he prefers one lottery to another, then his utility for it will be higher.

If the prizes in a lottery are all measured in the same commodity, then, as discussed previously, the certain equivalent of the lottery is the amount of the commodity that has the same utility as the lottery. The concepts of utility and certain equivalent play a central role in understanding risk preference.

In the practical question of measuring risk preference, one approach is to present an individual with a lottery and to ask him his certain equivalent. Or, we can provide the certain equivalent and all prizes but one and let him adjust the remaining prize until the certain equivalent is correct in his view. Finally, we can fix the certain equivalent and prizes and let him adjust the probabilities. All these questions permit us to establish the relationships between points on his utility curve and, ultimately, the curve itself. The interviewing in which the curve is measured is similar to that used for generating priors: the same need for education exists. The same types of inconsistency appear.

Although useful utility curves for individuals and organizations can be found in this manner, most decision-makers prefer to have some guidance in the selection of utility curves. The decision analyst can often provide this guidance by asking whether the decision-makers will accept additional axioms. One such axiom is: if all the prizes in the lottery are increased by some amount Δ, then the certain equivalent of the lottery will increase by Δ. The argument for the reasonableness of the axiom is very simple. The additional amount Δ is money in the bank, no matter which prize in the lottery is won. Therefore, the new lottery should be worth more than the original lottery. The counter argument is that having Δ in the bank changes the psychological orientation to the original lottery.

If this Δ axiom is added to the original set, then it is possible to show not just that a utility curve exists but that it must have a special form called the exponential form. A useful property of this exponential form is that it is described by a single number. This means that the analyst can characterize the utility curve of any individual or organization that wants to subscribe to these axioms by a single number—the risk aversion constant.

It is far easier to demonstrate to a decision-maker the consequences of his having different risk aversion coefficients and to measure his coefficient than it is to attempt to find a complete utility curve that is not of the exponential form. Encoding risk aversion in a single number permits measuring the sensitivity to risk aversion, as discussed earlier. In most practical problems, the entire question of risk aversion appears to be adequately treated by using the exponential form with a risk aversion constant appropriate to the decision-maker.

A cautionary note on the problem of practical measurement of risk aversion: experiments have revealed that the certain equivalents offered by subjects in hypothetical situations differ markedly from those offered when the situations are made real. This difficulty shows that the analyst must treat risk preference phenomena with great care.

Joint Time and Risk Preference

In most problems, both time and risk preference measures are necessary to establish the best alternative. Typically each outcome is represented by a time sequence of dependent uncertain values.

The question of how to describe preferences in such problems is fundamentally related to the way in which information on successive outcomes is revealed and to the extent to which it can help in making future decisions.

Two approaches illustrate the nature of the problem, each of which is appropriate under certain conditions. The first—that used in the original discussion of the probabilistic phase—is to compute the worth lottery implied by the model and then use the current utility function to develop the certain equivalent worth of the lottery. This approach is appropriate when there is no opportunity to utilize the information about outcomes as it is revealed, and thus where the prime interest is in the position occupied after all outcomes have been revealed.

Another approach is to imagine dealing with two agents. The first is a banker who will always pay immediately the amount specified by a particular company's time preference function applied to any time stream of values that is known with certainty. The other is a risk broker who will always pay the company's certain equivalent for any lottery. When faced with an uncertain stream of income, the company alternately deals with the risk broker to exchange lotteries for certain equivalents and with the banker to convert fixed future payments into present payments. The result of this alternating procedure is ultimately a single equivalent sum to represent the entire future process. Although appealing, the method may lead to the conclusion that the decision-maker should be willing to pay for "peace of mind" even when it has no effect on his financial future.

Thus the time-risk preference question ultimately depends on the decision-maker's tastes and options. The decision analyst can provide guidance in selecting from the many available approaches the one whose implications are best suited to the particular situation.

APPLICATIONS

In brief form, two examples illustrate the accomplishments and potential of decision analysis. In each case, the focus is on the key decision to be made and on the problems peculiar to the analysis.

New Product Introduction

A recent decision analysis was concerned with whether to develop and produce a new product. Although the actual problem was from another industry we shall suppose that it was concerned with aircraft. There were two major alternatives: to develop and sell a new aircraft (A_2) or to continue manufacturing and selling the present product (A_1). The decision was to be based on worth computed as the present value of future expected profits at a discount rate of 10% per year over a 22-year period. Initially, the decision was supposed to rest on the lifetime of the material for which the prior probability distribution, or priors, were obtained (Figure 8); however, a complete decision analysis was desired. Since several hundred million dollars in present value of profits were at stake, the decision analysis was well justified.

In the general scheme of the analysis, the first step was to construct a model for the business, as shown in Figure 9, which was primarily a model of the market. The profit associated with each alternative was described in terms of the price of the product, its operating costs, its capital costs, the behavior of competitors, and the natural characteristics of customers. Suspicion grew that this model did not adequately capture the regional nature of demand. Consequently, a new model was constructed that included the market characteristics region by region and customer by customer. Moving to the more detailed basis affected the predictions so much that the additional refinement was clearly justified. However, other attempts at refinement did not affect the results sufficiently to justify a still more refined model.

Fig. 9
New Product Introduction Decision Analysis

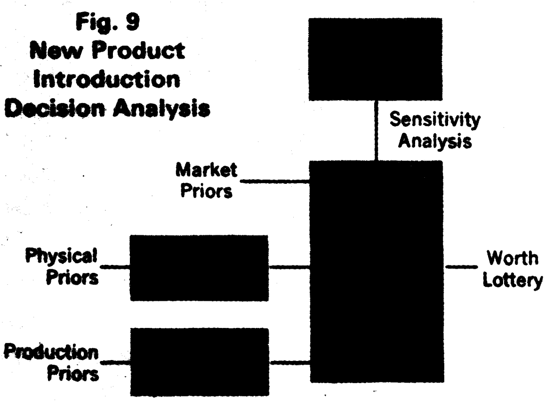

Next, a sensitivity analysis was performed to determine the aleatory variables. These turned out to be operating cost, capital cost, and a few market parameters. Because of the complexity of the original business model, an approximation was constructed showing how worth depended on these aleatory variables in the area of interest. The coefficients of the approximate business model were established by runs on the complete model.

The market priors were directly assigned with little trouble. However, because the operating and the capital costs were the two most important in the problem, their priors were assigned according to a more detailed procedure. First, the operating cost was related to various physical features of the design by the engineering department; this relationship was called the operating cost function. One of the many input physical variables was the average lifetime of the material whose prior appears in Figure 8. All but two of the 12 physical input variables were independent. The priors on the whole set were gathered and used together with the operating cost function in a Monte Carlo simulation that produced a prior for the operating cost of the product.

The engineering department also developed the capital cost function, which was much simpler in form. The aleatory variables in this case were the production costs for various parts of the product. A simulation produced a prior on capital cost.

With priors established on all inputs to the approximate business model, numerical analysis determined the worth lottery for each alternative. The worth lotteries for the two alternatives closely resembled those in Figure 4, Part A. The new product alternative A_2 stochastically dominated the alternative A_1 (continuing to manufacture the present product). The result showed two interesting aspects of the problem. First, it had been expected that the worth lottery for the new product alternative would be considerably broader than it was for the old product. The image was that of a profitable and risky new venture compared with a less profitable, but less risky, standard venture. In fact the results revealed that the uncertainties in profit were about the same for both alternatives, thus showing how initial impressions may be misleading.

Second, the average lifetime of the material whose priors appear in Figure 8 was actually of little consequence in the decision. It was true enough that profits were critically dependent on this lifetime if the design were fixed. But leaving the design flexible to accommodate to different average material lifetimes was not an expensive alternative. The flexible design reduced sensitivity to material lifetime so much that its uncertainty ceased to be a major concern.

The problem did not yield as easily as this, however. Figure 10 shows the present value of profits through each number of years t for each alternative. Note that if returns beyond year 7 are ignored, the old product has a higher present value; but in considering returns over the entire 22-year period, the relationship reverses. When managers saw these results they were considerably disturbed. The division in question had been under heavy pressure to show a profit in the near future, and alternative A_2 would not meet that requirement. Thus, the question of time preference that had been quickly passed off as one of present value at 10% per year became the central issue in the decision. The question was whether the division was interested in the quick kill or the long pull.

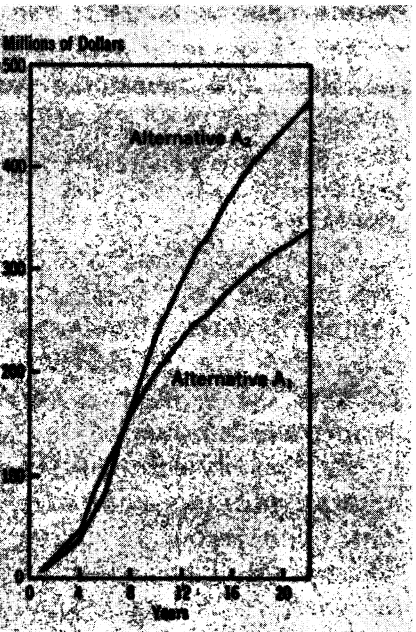

Fig. 10 Expected Present Value of Profit

This problem clearly illustrates the use of decision analysis in clarifying the issues surrounding a decision. A decision that might have been made on the basis of a material lifetime was shown to depend more fundamentally on the question of time preference for profit. The extensive effort devoted to this analysis was considered well spent by the company, which is now interested in instituting decision analysis procedures at several organizational levels.

Space Program Planning

A more recent application in a quite different area concerned planning a major space program. The problem was to determine the sequence of designs of rockets and payloads that should be used to pursue the goal of exploring Mars. It was considered desirable to place orbiters about Mars as well as to land vehicles on the planet to collect scientific data.

The project manager had to define the design for each mission—that is, the type and number of launch vehicles, orbiters, and landers. The choice of design for the first mission could not logically be made without considering the overall project objectives and the feasible alternatives. Key features of the problem were the time for the development of new orbiting and landing vehicles, cost of each mission, and chances of achieving objectives.

Approach to Solution

To apply decision analysis to the problem posed, a two-phase program was adopted. The first or pilot phase consisted of defining a simplified version of the decision. To the maximum extent possible, however, the essential features of the problem were accurately represented and only the complexity was reduced. This smaller problem allowed easier development of the modeling approach, and exercising of the model provided insight into the level of detail required in structuring the inputs to the decision. The second phase consisted of developing the more realistic and complex model required to decide on an actual mission.

The Pilot Phase

To begin the decision analysis, four possible designs were postulated to represent increasing levels of sophistication. Figure 11 shows these designs and their potential accomplishments. The questions were: what design should be selected for the first opportunity, and what sequence of designs should be planned to follow the first choice? Should the project manager, for example, elect to provide the ultimate level of capability in the initial design in the face of uncertainties in the Martian environment and difficulties in developing complex equipment to survive the prelaunch sterilization environment? Or should he choose a much simpler design that could obtain some information about the Martian environment to be used in developing subsequent, more complex, vehicles.

Decision Trees

The heart of the model used in analyzing the decision was a decision tree that represented the structure of all possible sequences of decisions and outcomes and provided for cost, value, and probability inputs. Such trees contain two types of nodes (decision nodes and chance nodes) and two types of branches (alternative branches and outcome branches), as illustrated in Figure 12. Emanating from each decision node is a set of alternative branches, each branch representing one of the alternatives available for selection at that point of decision. Each chance node is fol-

Fig. 11—Configurations and Performance

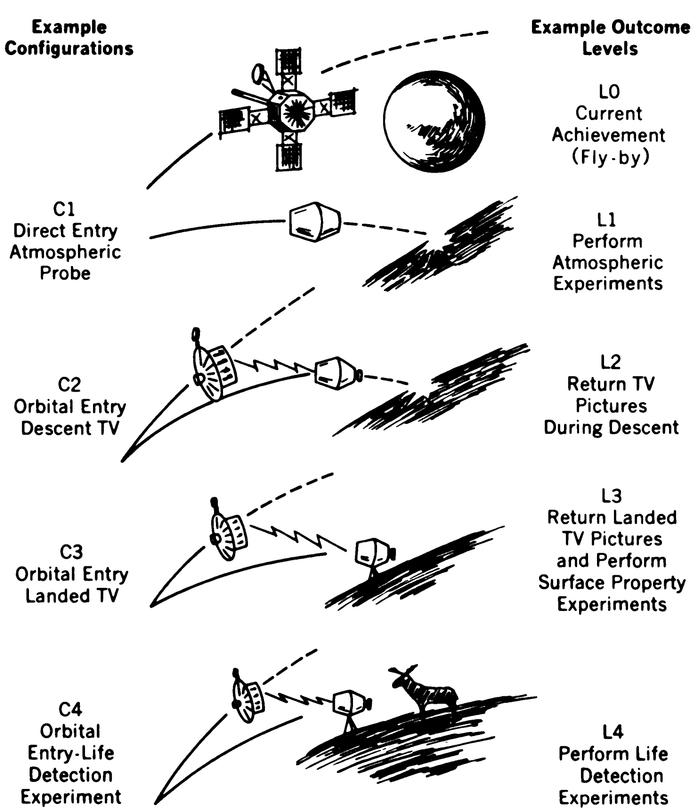

Fig. 12—Tree Relationships

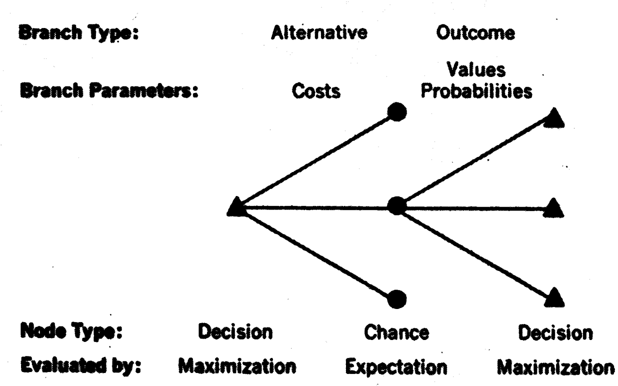

lowed by a set of outcome branches, one branch for each outcome that may be achieved following that chance node. Probabilities of occurrence and values are assigned to each of these outcomes; costs are assigned to each decision alternative.

Two fundamental operations, expectation and maximization, are used to determine the most economic decision from the tree. At each chance node, the expected profit is computed by summing the probabilities of each outcome, multiplied by the value of that outcome plus expected profit of the node following that outcome. At each decision node, the expected profit of each alternative is calculated as the expected profit of the following node ("successor node") less the cost of the alternative. The optimum decision is found by maximization of these values over the set of possible alternatives, i.e., by selecting the alternative of highest expected profit.

Order of Events

The particular sequence of mission decisions and outcomes was a significant feature of the pilot analysis. As illustrated in Figure 13, the initial event of significance was the selection of the 1973 mission configuration. However, since lead time considerations re-

Fig. 13 ORDER OF EVENTS

	1968	1969	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	• • •
First Flight	S					L	O						
Second Flight					S			L	O				
Third Flight							S			L	O		
Fourth Flight									S			L	O
Fifth Flight											S		• • •

S = Select L = Launch O = Outcome

quired that the 1975 configuration decision be made in 1972, the second mission decision had to be made prior to obtaining the first mission results. Similarly, the 1977 decision had to be made before obtaining the results of the 1975 mission, although after the 1973 mission results. In general, then, a mission configuration was made in ignorance of the results of the previous mission.

Tree Example

A complete decision tree for the pilot project, with the additional assumption that L2 is the highest level of success, is presented in Figure 14. The model that produces the numerical probabilities, values, and costs used in the example will be discussed later. Node 1, at the left side of the tree is the initial decision to select either a $C1$ or a $C2$ for the first launch opportunity. The box designated LO above this node indicates that the state at this node is the current level of achievement. Suppose a $C1$ is selected. The cost of that $C1$ is $850 million, indicated by the "−850" that is written under that branch. As a result of this choice, the next node is decision node 2. The box designated LO, $C1$ above this node indicates that the state of this node is the current level of achievement and a $C1$ is being constructed for the first launch. Now either a $C1$ or $C2$ must be selected for the second launch. If a $C1$ is selected, the cost is $575 million, and the next node is chance node 7. The two branches following this node represent the possible outcomes of the first launch. The LO' outcome which would be failure to better LO on the first try, occurs with probability 0.1 whereas the $L1$ outcome occurs with probability 0.9. The value of the LO' outcome is zero, whereas the value of the LO outcome is 1224. Now follow the case of the $L1$ outcome to decision node 34. The state $L1$, $C1$ at this node, means that the highest level of success is $L1$ and that a $C1$ is being constructed for the next launch. Since $L1$ has already been achieved at this point in the tree, a $C2$ is the only design that may be launched in the third opportunity, at a cost of $740 million. This leads to decision node 35, where the state is $L1$, $C2$.

Node 35 in the example tree illustrates coalescence of nodes, a feature vital to maintaining a manageable tree size. Node 35 on the upper path through the tree can be reached from four other paths through the tree as indicated in the exhibit. If the coalescence did not occur, the portion of the tree following node 35 would have to be repeated four additional times. In the full pilot tree, coalescence results in a reduction of the number of branches in the tree by a factor of 30.

Along the path 1-2-7-34-35, at decision node 35, a $C2$ must be selected for the fourth opportunity. At chance node 36, the outcome of the third launch is either an $L1'$ (failure to better $L1$ with one attempt, which leads to node 38), or an $L2$ (which achieves a value of 1714 and successfully completes the program). These outcomes occur with probability 0.3 and 0.7, respectively. If $L1'$ is the outcome, chance node 38 is reached, where the outcome of the fourth launch is represented. The probability of $L1''$ is 0.24, and the probability of $L2$ is 0.76. Note that the probability of 12 has increased over that of node 36 (0.7 to 0.76) because of the experience gained previously.

One can similarly follow and interpret many other paths through the tree. A policy is a complete selection of particular alternatives at all decision nodes. This limits the set of all possible paths to a smaller subset. (It is not possible, for example, to reach node 26 if a $C1$ is chosen at node 1.) The probabilities, values, and cost of these paths then determine the characteristics of the decision policy.

The most economic policy, given the input data specifications, is defined as the policy that maximizes the expected profit of the project, i.e., expected value less expected cost.

Fig. 14
Example of a Decision Tree

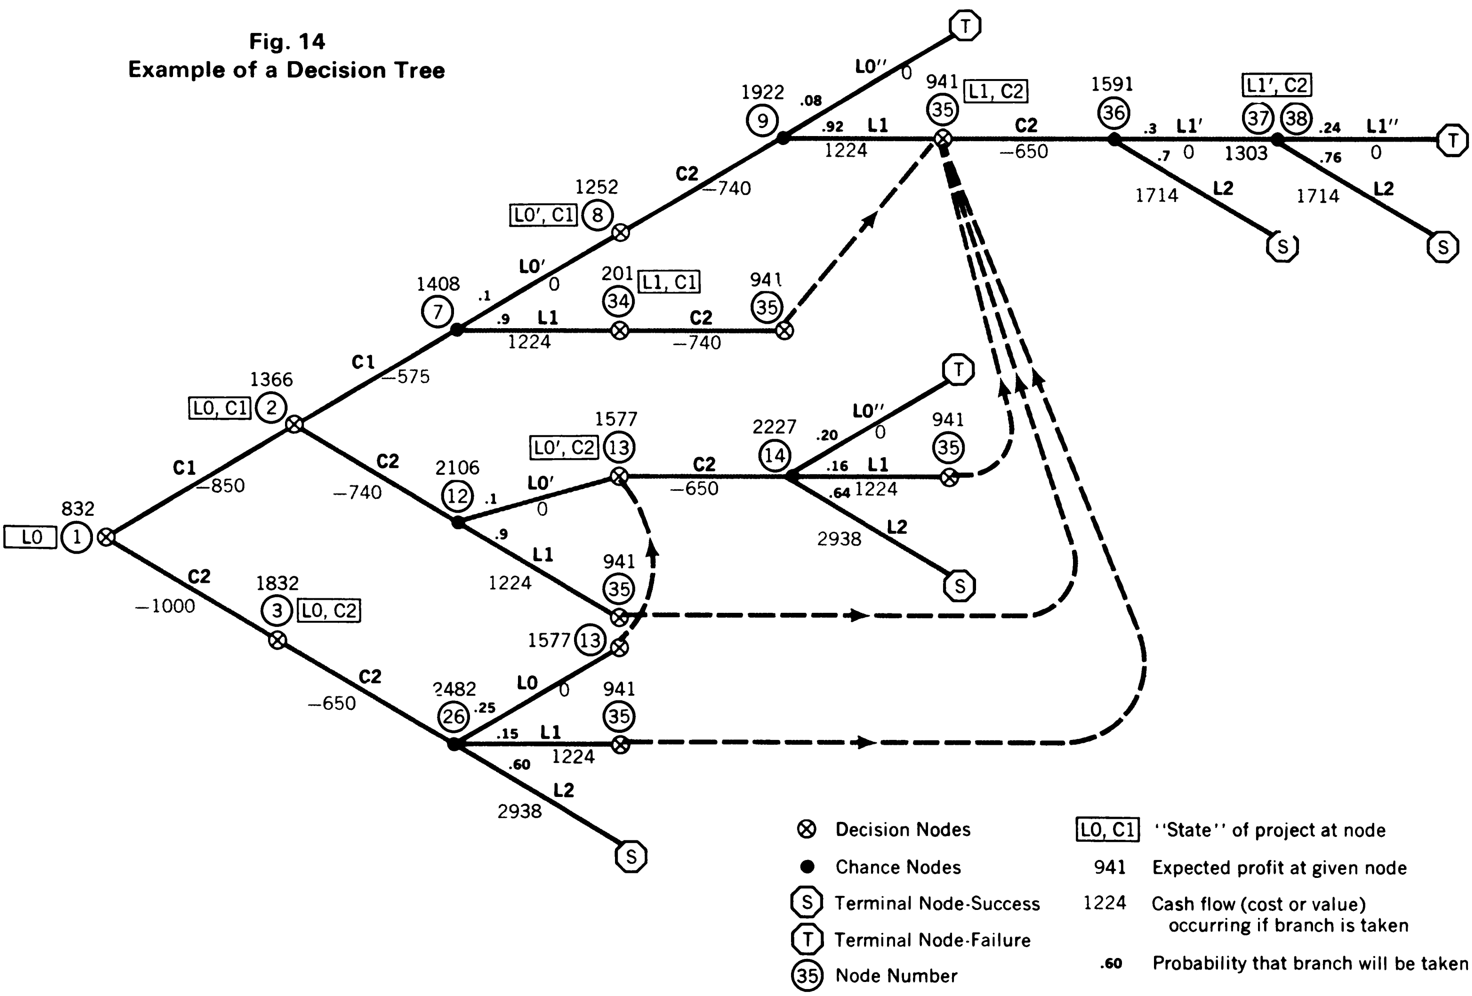

The technique illustrated here eliminates many of the nonoptimum policies from explicit consideration; it is the "roll back" technique that starts from the right side of the tree and progresses left to the beginning of the tree, making all decisions and calculations in *reverse* chronological order. Thus, when each decision is made, only policies that optimize decisions for the following decision nodes are considered.

Consider node 38 in Figure 14. At this chance node the probability of achieving $L1''$, which is worth nothing, is 0.24, and the probability of achieving $L2$, which is worth 1714, is 0.76. Thus, the expected profit of node 38 is: $0.24(0) + 0.76(1714) = 1303$. This number is written near node 38.

The calculations are carried out in this manner backwards through the tree. The first decision node with more than one choice is node 2. If a $C1$ is selected, it costs \$575 million (-575) and leads to node 7 with an expected profit of 1408, which yields $-575 + 1408 = 833$. If a $C2$ is selected, it costs \$740 million (-740) and leads to node 12 with an expected profit of 2106, which yields $-740 + 2106 = 1366$. Since 1366 is greater than 833, the most economic decision is to select a $C2$ at node 2, which results in an expected profit of 1366.

Finally, the first decision is a choice between a $C1$ with an expected profit of 516 or a $C2$ with an expected profit of 832. Maximum expected profit is achieved by the choice of a $C2$ resulting in an expected profit of 832. This

Fig. 15
Decision Tree

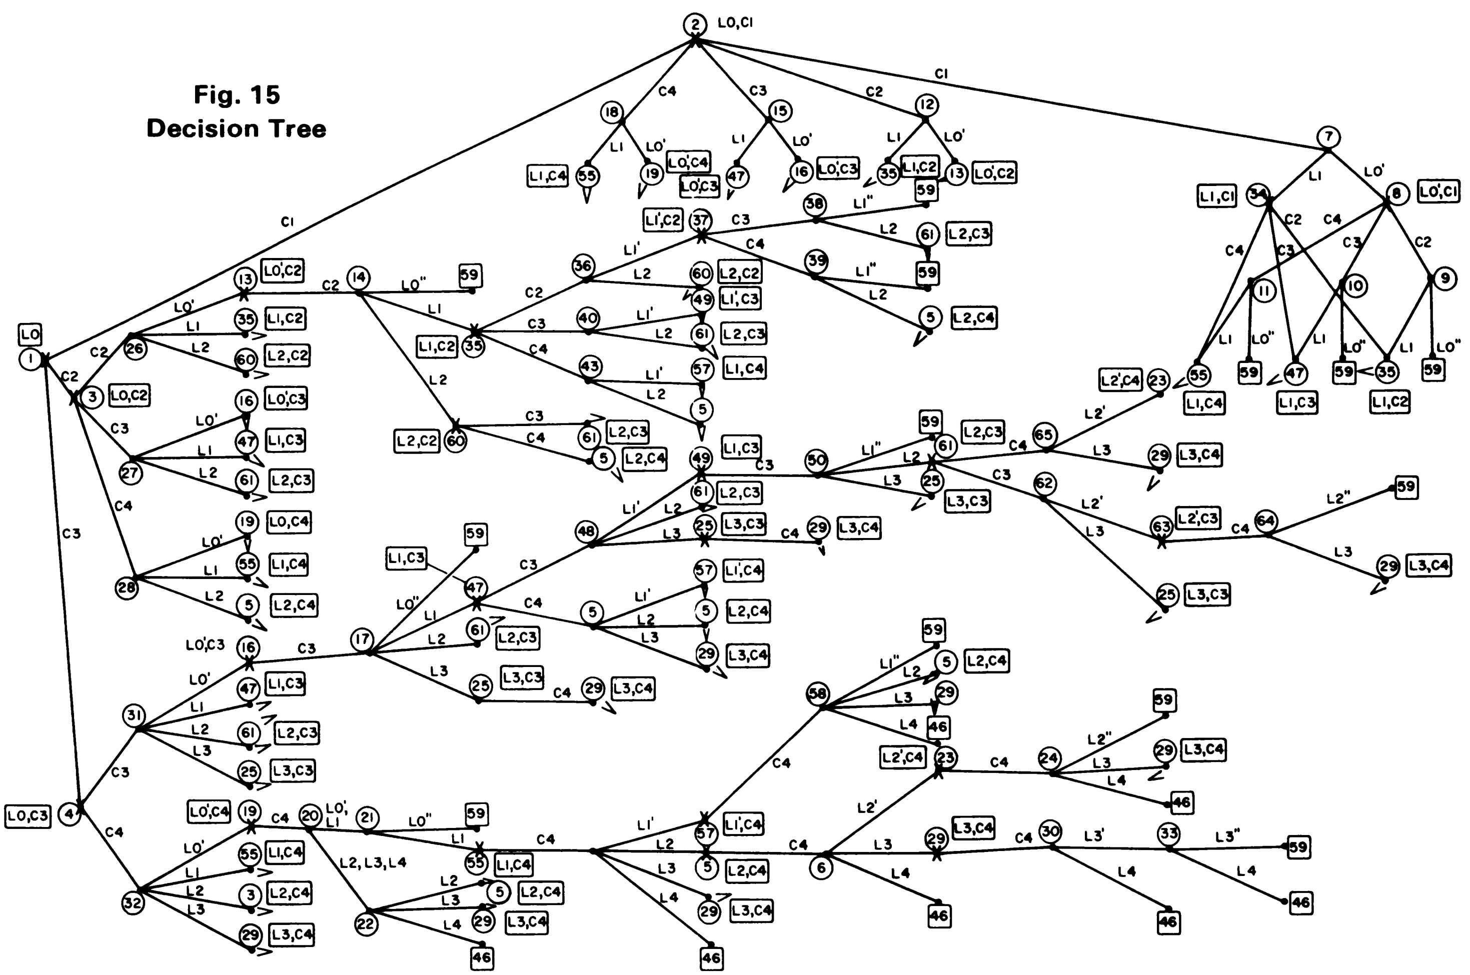

Note: Nodes 46 and 59 are the terminal nodes. Node 46 corresponds to L4 and is reached by achieving a totally successful project. Node 59 is reached when two successive failures force termination of the project prior to achieving L4.
↙ indicates direction of coalesced node bearing same number.

is the expected profit of the entire project at the time the first decision is made.

Figure 15 illustrates the complexity of the completed decision tree for the pilot phase of the analysis.

Value Assignment

A particularly important part of this study was the specification of the value to be attached to the outcomes of the program. Since the decision-makers were reluctant to state values in dollar terms, a tree of point values was employed. The value tree is simply a convenient way of showing how the total value of the project is to be broken down into its component outcomes. Figure 16 shows a value tree for the pilot analysis. The points assigned to each tip of the tree are the fraction of total program value assigned to this accomplishment; the values accumulate as the program progresses. A total dollar value assigned to a perfect program therefore determines the dollar values used in the decision tree.

Fig. 16—The Value Tree

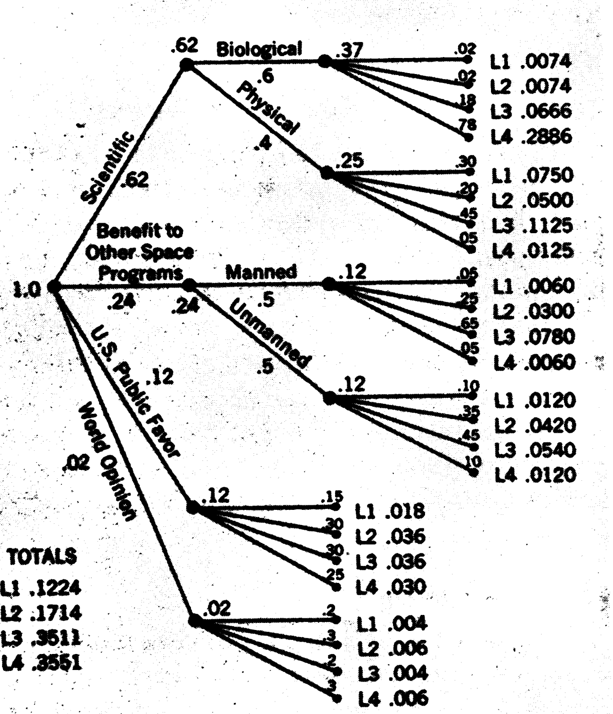

To derive a value measure, a value tree is constructed by considering first the major components of value and then the subcategories of each type, which are identified in more and more detail until no further distinction is necessary. Then each tip of the tree (constructed as above) is subdivided into four categories, each corresponding to the contribution of one of the four levels of achievement within the value subcategory represented by that tip.

The number 1.0 attached to the node at the extreme left of the value tree for the pilot analysis represents the total value of all the objectives of the pilot project (thus, the value of achieving *L*1, *L*2, *L*3, and *L*4). The four branches emanating from this node represent the four major categories of value recognized by the pilot model. The figure 0.62 attached to the upper branch represents the fraction of total value assigned to science. Two branches emanate from the science node, and 60% of the science value falls into the category of biological science. The 0.37 attached to the biological science node represents the fraction of *total* value attached to biological science, and is obtained by taking 60% of 0.62 (the fraction of total value attached to all science). Finally, the bottom branch following the biological science node indicates that 78% of the biological science value is achieved by jumping from *L*3 to *L*4.

The final step in value modeling is to obtain the fraction of total value to be attached to achieving each of the four levels. If all the contributions to achieving *L*1 (e.g., contributions to world opinion, U.S. public favor, physical science) are added, the result is the fraction of value that should be attached to achieving *L*1. The same process is followed for reaching *L*2 from *L*1, *L*3 from *L*2, and *L*4 from *L*3. The results of such a calculation are presented in the lower left corner of Figure 16.

Summary

On the basis of the promising results of working with the pilot model, a more complete model was developed to encompass nearly all of the factors involved in selecting the actual mission. It provided a more precise structure for assigning initial values, probabilities, and costs, and for updating probabilities and costs based on results achieved. The following tabulation shows a summary comparison of the complexity of the pilot model with the more complete model.

DECISION TREE COMPARISON TABLE

Pilot	Feature	Full Scale
4	Mission Designs	14
5	Outcomes	56
56	Decision Tree Nodes	3153
1592	Paths Through Tree	354,671,693

Clearly, the full-scale decision tree could not be represented graphically. The tree was constructed and evaluated by computer program specially developed for this application.

A model such as the one described here can be a valuable tool throughout the life of a project. As the project progresses, the knowledge of costs, probabilities, and values will improve as a result of development programs and flights. Improved knowledge can be used in the decision process each time a design must be selected for the next opportunity.

An important additional benefit of this analysis is that it provides a language for communicating the structure of the space project and the data factors relevant to the project decisions. It provides a valuable mechanism for discourse and interchange of information, as well as a means of delegating the responsibility for determining these factors.

FUTURE TRENDS

Decision analysis should show major growth, both in its scope of applications and in its effect on organizational procedures. This section presents various speculations about the future.

Applications

Market Strategy Planning

The importance of decision-making in a competitive environment has stimulated the use of decision analysis in both strategic and tactical marketing planning. The strategic problems are typically more significant because they affect the operations of the enterprise over many years. Strategic analysis entails building models of the company and of its competitors and customers, analyzing their interactions, and selecting strategies that will fare well in the face of competitive activities. Since most of this work is of a highly confidential nature, little has appeared in the public literature; nevertheless, there is reason to believe that many large U.S. corporations are performing work of this kind, however rudimentary it may be. The competitive analyses of a few quite sophisticated companies might rival those conducted in military circles.

Resource Exploration and Development

Resource exploration by mineral industries is a most natural application for decision analysis. Here the uncertainty is high, costs are great, and the potential benefits extremely handsome. At all levels of exploration—from conducting aerial surveys, through obtaining options on drill-test locations, to bidding and site development—decision analysis can make an important contribution. Organizations approaching these problems on a logical, quantitative basis should attain a major competitive advantage.

Capital Budgeting

In a sense, all strategic decision problems of a corporation are capital budgeting problems, for its ultimate success depends upon how it allocates its resources. Decision analysis should play an increasingly important role in the selection of projects and in objective comparisons among them. Problems in spending for research and development programs, investment in new facilities, and acquisitions of other businesses will all receive the logical scrutiny of decision analysis. The methodology for treating these problems already exists; it now remains for it to be appreciated and implemented.

Portfolio Management

The quantitative treatment of portfolio management has already begun but it will receive even more formal treatment in the hands of decision analysts. The desires of the investing individual or organization will be measured quantitatively rather than qualitatively. Information on each alternative investment will be encoded numerically so that the effect of adding each to the portfolio can be determined immediately in terms of the expressed desires. The human will perform the tasks for which he is uniquely qualified: providing information and desires. The formal system will complement these by applying rapid logic.

Social Planning

On the frontiers of decision analysis are the problems of social planning. Difficult as it may be to specify the values and the criteria of the business organization, this problem is minor compared with those encountered in the public arena. Yet if decision-making in the public sector is to be logical, there is no alternative.

The problems to which a contribution can be made even at the current stage of development are virtually endless: in decisions associated with park systems, farm subsidies, transportation facilities, educational policy, taxation, defense, medical care, and foreign aid, the question of values is central in every case.

The time may come when every major public decision is accompanied by a decision analysis on public record, where the executive branch makes the decision using values specified by the people through the legislative branch. The breakdown of a public decision problem into its elements can only serve to focus appropriate concern on the issues that are crucial. For the first time, the public interest could be placed "on file" and proposals measured against it. A democracy governed in this fashion is probably not near at hand, but the idea is most intriguing.

Procedures

The effect of decision analysis on organizational procedures should be as impressive as its new applications. Some of the changes will be obvious, others quite subtle.

Application Procedures

Standardization by type of application will produce special forms of analyses for various types of decisions—for example, marketing strategy, new product introduction, research expenditures. This standardization will mean special computer programs, terminology, and specialization of concepts for each application. It will also mean that the important classes of decisions will receive much more effective attention than they do now.

Analytical Procedures

Certain techniques, such as deterministic, stochastic, and economic sensitivity analyses that may be performed with the same logic regardless of the application will be carried out by general computer programs. In fact, the process of development is well under way at the present time. Soon the logical structure of any decision analysis might be assembled from standard components.

Probabilistic Reporting

The introduction of decision analysis should have a major impact on the way organizational reporting is performed externally and internally. Externally, the organization will be able to illustrate its performance not just historically by means of balance sheets and operating statements, but also projectively by showing management's probability distributions on future value. Since these projections would be the result of a decision analysis, each component could be reviewed by interested parties and modified by them for their own purposes. However, management would have a profitable new tool to justify investments whose payoffs lie far in the future.

Organizational management will acquire new and more effective information systems as a result of decision analysis. Internal reporting will emphasize the encoding of knowledge in quantitative form. Instead of sales forecasts for next year, there will be probability distributions of sales. Thus, the state of information about future events will be clearly distinguished from performance goals.

Delegation by Value Function

An important logical consequence of decision analysis is that delegation of a decision requires only transmission of the delegator's present state of information and desires. Since both of these quantities can be made explicit through decision analysis, there should be an increase in the extent and success of delegation. In the external relationships of the firm, the delegation will no doubt appear as an increased emphasis on incentive contracts, where the incentives reflect the value function of the organization to the contractor. This trend is already evident in defense contracting.

Internally, the use of the value function for delegation should facilitate better coordination of the units of the organization. If explicit and consistent values are placed on the outcomes of production, sales, and engineering departments, then the firm can be sure that decisions in each unit are being made consistently with the best overall interests of the firm. The goal is to surround each component of the organization with a value structure on its outputs that encourages it to make decisions as would the chief decision-maker of the organization if he were closely acquainted with the operations of the component.

Organizational Changes and Management Development

The introduction of decision analysis will cause changes in organizational behavior and structure. A change should take place in the language of management, for the concepts discussed in this report are so relevant to the decision-making process that, once experienced in using them, it is difficult to think in any other terms. The explicit recognition of uncertainty and value questions in management discussions will in itself do much to improve the decision-making process.

Special corporate staffs concerned with the performance of decision analysis are already beginning to appear. These people would be specially trained in decision analysis, probability, economics, modeling, and computer implementation. They would be responsible for ensuring that the highest professional standards of logic and ethics are observed in any decision analysis.

Special training for decision analysts will be accompanied by special training for managers. They will need to know much more than they do now about logical structure and probability if they are to obtain full advantage from the decision analyst and his tools. No doubt much of this training will occur in special courses devoted to introducing decision analysis to management. These courses will be similar to, but more fundamental than, the courses that accompanied the introduction of computers into the U.S. economy.

Management Reward

Encouraging managers to be consistent with organizational objectives in decision-making requires adjusting the basis for their rewards to that objective. If rewarded only for short run outcomes, they will have no incentive to undertake the long range projects that may be in the best interest of the organization. It follows that any incentive structure for management will have to reward the quality of decisions rather than the quality of outcomes. The new financial statements that show probability distributions on future profit would be the key to the reward structure. After these distributions had been "audited" for realism, the manager would receive a reward based upon them in a predetermined way. Thus, the manager who created many new investment opportunities for a company could be rewarded for his efforts even before any were fully realized.

To make this system feasible requires distinguishing between two kinds of managers: the one who looks to the future and prepares for it; and the one who makes sure that today's operations are effective and profitable. The distinction is that between an admiral and a captain, or between the general staff and the field commanders. Specialization of function in corporate management with significant rewards and prestige attached to both planning and execution could be the most important benefit of decision analysis.

CONCLUSION

Although an organization can achieve ultimate success only by enjoying favorable outcomes, it can control only the quality of its decisions. Decision analysis is the most powerful tool yet discovered for ensuring the quality of the decision-making process: its ultimate limit is the desire of the decision-maker to be rational.

3

DECISION ANALYSIS IN SYSTEMS ENGINEERING

Ronald A. Howard

Department of Engineering-Economic Systems
Stanford University

Reprinted from Systems Concepts: Lectures on Contemporary Approaches to Systems (Ralph F. Miles, ed.) New York: John Wiley & Sons, 1973.

Decision Analysis in Systems Engineering

RONALD A. HOWARD

Professor of Engineering-Economic Systems
School of Engineering
Professor of Management Science
Graduate School of Business
Stanford University

Ronald A. Howard is a Professor of Engineering-Economic Systems and a Professor of Management Science at Stanford University. He received his Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology and served on the faculty until 1965. He is the author of *Dynamic Programming and Markov Processes* and *Dynamic Probabilistic Systems.* He has editorial responsibilities for *Management Science, Operations Research, IEEE Transactions on Systems, Man, and Cybernetics, Journal of Optimization and Control,* and *Journal of Socio-Economic Systems.* He is the Series Editor for the Wiley series, Decision and Control. Howard's research and consulting interests are currently concentrated on decision analysis.

The past decade has seen the development of a new profession—decision analysis, a profession concerned with providing a rational basis for decision-making. While it may seem strange that people can make their living by helping other people make decisions, that is just what decision analysts do. So that we can better see the need for this new profession, let us start by taking a look at the kind of decision-making we use in our everyday lives (see Fig. 1).

Descriptive Decision-making

In this descriptive view of decision-making, we first examine the environment of human decisions. The environment can be described by several characteristics. As we go through these characteristics, think of them not only in terms of modern corporate or governmental decisions, but in terms of any decision—personal, romantic, or historical. Perhaps we might even think back to the dawn of history when the caveman was trying to decide whether to take one path or another in order to avoid the saber-toothed tiger. He saw his environment, as we now see ours—uncertain. If there is any one attribute of the environment that gives us the most difficulty in decision-making, it is uncertainty. Furthermore, the environment is complex—we see many different factors interacting in ways we often cannot understand. It is dynamic. It evolves over time. What we do today has effects that may not be evident for years.

Unfortunately, in business, military, or national problems the environment is often competitive. There are hostile intelligences that are trying to make life better for themselves at our expense. Perhaps, most unfortunate of all, our resources are finite—in spite of the exhortations of religious leaders over the centuries, man perceives himself as being a limited creature who has to allocate what he has, rather than to expand it.

The typical human reaction to these characteristics of the environment is confusion or worry, whether it be corporate—and there are

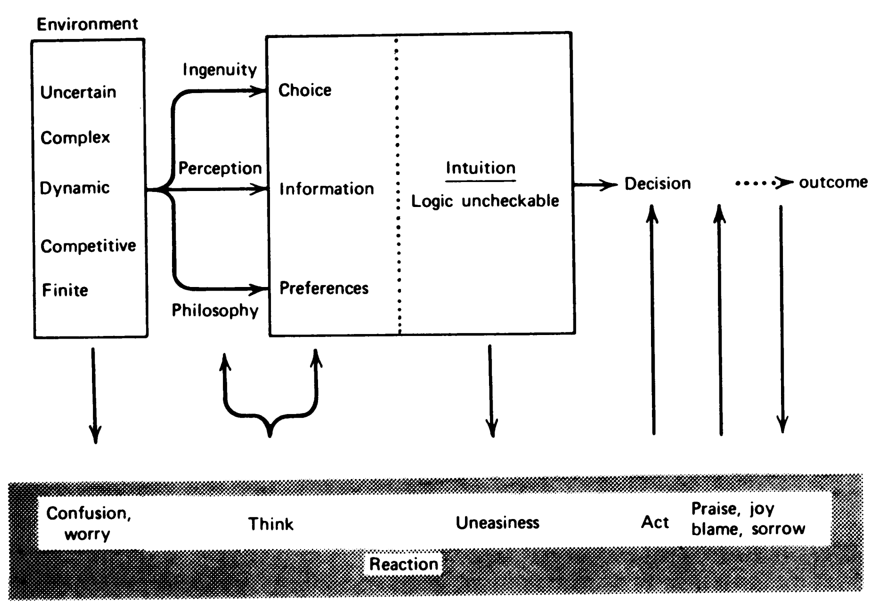

Figure 1 Decision-making (descriptive).

many worried corporate executives—or individual. But man has weapons available in his fight with the environment—his fight to make a decision. We separate these weapons into three types.

First, man has ingenuity. He uses that ingenuity to conceive and formulate different courses of action; that is, he has the potential of choice. Second, he has perception. He can learn from what he sees. He can form judgments about the world. In other words, he can gather information about his environment. Finally, and in some ways perhaps most important, he has a philosophy. He has guiding principles of his life that give him preferences among the various outcomes that he might obtain from his decisions.

Although we have only looked at the personal application of these ideas, corporations, colleges, and governments also have these attributes. We might combine them into the idea of thinking. This is the process of thought—the process that man brings to the problem of making a decision.

At this point let us define a decision. A decision is an allocation of resources that is revocable only at a cost in some resource, such as time or money. For practical purposes we should think of it as an irrevocable allocation of resources. It is not a mental commitment to changing the state of the world or to carrying out some course of action, but rather some actual physical change.

Intuition

The process by which man makes the overwhelming number of decisions in his life (or a corporation in its life) is intuition. He uses an intuitive process to balance the choices, the information, the preferences that he has expressed and to arrive at a course of action—to arrive at a resource allocation. We cannot say much about how intuition operates, but we have all met people who operate intuitively. Indeed, we all make intuitive decisions in our own lives every day: which route to take to get to work, when to get up in the morning, and so forth. We would be foolish to substitute any other principle for intuition in the majority of the decisions we make. But there are some decisions that we, as individuals or as organizations, face that are so important—so crucial to our existence, survival, and gathering of joy—that we must strive for a better way of making them.

The characteristic of intuition that is most bothersome to us is that its logic in uncheckable. If a person were the chairman of the board and made a decision by intuition, he might say, "Well fellows, I've read all the reports and, having thought it over, I think we ought to merge with Company X." While such a decision could be a great idea, we really have no way of evaluating it. There is no way of checking step-by-step to determine whether this decision is the logical consequence of the choices, information, and preferences that were available to the decision maker. It all went on in his mind—behind closed doors, so to speak. While a one-man company may be able to get away with such a decision, in our increasingly interdependent corporate world or in our society it becomes increasingly important for a man to be able to show people why he arrived at a particular decision. It is also important for them to be able to see what changes in factors surrounding that decision might have led to a different decision.

Thus we often find that one result of the intuitive decision-making process is uneasiness on the part of the individual or the organization making the decision. You would be surprised at the number of corporate decision makers who arrive at a decision on intuitive grounds and then, after the fact (after they have made the mental commitment, but before they have written the check), come looking for some better way of making the decision because they are uneasy about whether that decision is consistent with their choices, information, and preferences.

Decision Versus Outcome

One other thing we ought to mention about descriptive decision-making is the unfortunate human tendency to equate the quality of the decision with the quality of the outcome it produces. Each decision is followed by an outcome that is either joyful or sorrowful, for example: the surgeon decides to amputate the arm, and the patient either recovers or dies; the investor decides to buy some new stock, and the stock either makes money or loses money. We tend to say if the stock lost money, or if the patient died, that the decision maker made a bad decision.

Well, logically, that is indefensible, because the only way you can evaluate the quality of a decision is by whether it is consistent with the choices, information, and preferences of the decision maker. While we all prefer a joyful outcome to a sorrowful outcome, only the decision is under our direct control. We must seek aid in exploiting that control to the fullest extent, but we must distinguish the quality of the decision from the quality of the outcome.

Here is a simple illustration. Suppose that a person were offered the opportunity to call the toss of a coin for $100. If he calls that toss correctly he wins $100. If he does not, he wins nothing. There are very few people who wouldn't like to play such a game. Suppose that the offer were for a payment of $5. We can be sure that many people would like to play such a game for $5. Then, picture a line of people waiting for their turns. The first one comes up, we take his $5, we toss the coin, he calls it, and he loses. Now what? What do we say? We say he had a bad outcome. The next one comes up; he also pays his $5, and he wins. He had a good outcome. These people have both made the same decision. It was a good decision, but making that good decision is no guarantee of a good outcome. Speaking loosely, making a good decision is only doing the best we can to increase the chances of a good outcome.

One thing that anyone who deals with decision analysis should keep in mind is the importance of differentiating between the quality of the decision and the quality of the outcome. This distinction is the very beginning of the study of decision-making. It is this transcendence of the intimacies of outcome by conceptualizing the decision-making process that allows us to study formally what "good decision" means.

In most cases we do not really know what is a good decision. We are so used to characterizing the kind of decision that was made by the kind of outcome produced that we really have not until now had a procedure—an engineering analysis, a science, if you like—for recognizing a good decision. One of the "reasons to be" for decision analysis is to formulate the idea of what a good decision is, and to formulate it in quantitative terms that can be conveyed from one person to another, compared from one situation to another.

So much for descriptive decision-making. We will probably be using it a lot in our personal lives and in our organizations' lives, but it has the shortcomings of intuition that we think we can now transcend.

Decision Analysis

In this chapter we examine an alternative to descriptive decision-making, an alternative called "decision analysis." Here is a very brief definition. It is the balancing of the factors that influence a decision and, if we wanted to add another word, a *logical* balancing of the factors that influence a decision. Typically these factors might be technical, economic, environmental, or competitive; but they could also be legal or medical or any other kind of factor that affects whether the decision is a good one.

The Precursors of Decision Analysis

Decision analysis is a term we also use to describe the outgrowth of two earlier fields, namely, decision theory and systems modeling methodology. Decision theory was largely the province of academics until very recently. They treated the question of how to be rational in very simple, but uncertain, situations dealing with balls in urns, coin-tossing, small amounts of money, and the like. But it turned out that there was enough meat to the question of what is a good decision (even in simple cases) that theorists for years—going back to Bernoulli in 1738—have been worried about what really constitutes a good decision. However, decision theory was a theory for very simple decisions and certainly far from application to the complex corporate or even personal decisions we face today.

Over the past 30 years, we have also seen the development of a systems modeling methodology. That systems modeling methodology provided means of treating the complex and dynamic aspects of the environment in a way that had never been contemplated before. Of course, the advent of the computer played a large role. Decision analysis is the child of both of these developments. It is a way to combine the ability to handle complexity and dynamics with the ability to handle decision-making in the face of uncertainty into a single discipline that can treat all three simultaneously.

A Language and a Procedure

This new discipline has two interesting aspects. First, it is a language and philosophy for decision-making. It is a way to talk about the decision-making process even if you never set pencil to paper to do a computation. Indeed, organizations that have begun to think in this way—to use this language and philosophy—can never, it appears, revert to their old ways of thinking. It is a kind of reverse Gresham's law: the specification of language and clarity of concept keeps us from thinking about decisions in ways that might not be fruitful in the making of them.

But more than that, as far as the profession itself is concerned, is the idea that decision analysis is a logical and quantitative procedure. It is not simply a way to talk about decision-making; it is actually a way to make a decision. It is a way to build a model of a decision that permits the same kind of checking and testing and elimination of bugs that we use in the engineering of an automobile or an airplane. If it were not for the fact that "decision engineering" somehow implies a kind of manipulation of the decision-making process rather than an analysis of it, this field might be called decision engineering rather than decision analysis.

The Decision Analysis Formalism

How does decision analysis differ from intuitive decision-making? In some ways, not at all; in other ways, very significantly. First, consider the environment (see Fig. 2). There is "bad news" on that score because the environment is still uncertain, complex, dynamic, competitive, and finite. We shall have to live with it—decision analysis is not a "crystal ball" procedure, much as people wish it were. So we will still be confused and worried when we start out on a decision problem. Furthermore, there is no hope for people who do not like to think, because we must be ingenious, perceptive, and philosophical in order to carry out decision analysis. So it is not much help yet.

Choice. Where we start to get help is now. First, let us go through the three aspects of ingenuity, perception, and philosophy, one by one, and see how they are treated within this new discipline. The idea of choice is spelled out by enumerating specific alternatives that are available in this decision problem. They may be finite alternatives, like amputate or do not amputate; or they may be alternatives described by continuous variables, such as the capacity of a plant, the price of a new product, or even the size of a budget that will be set for

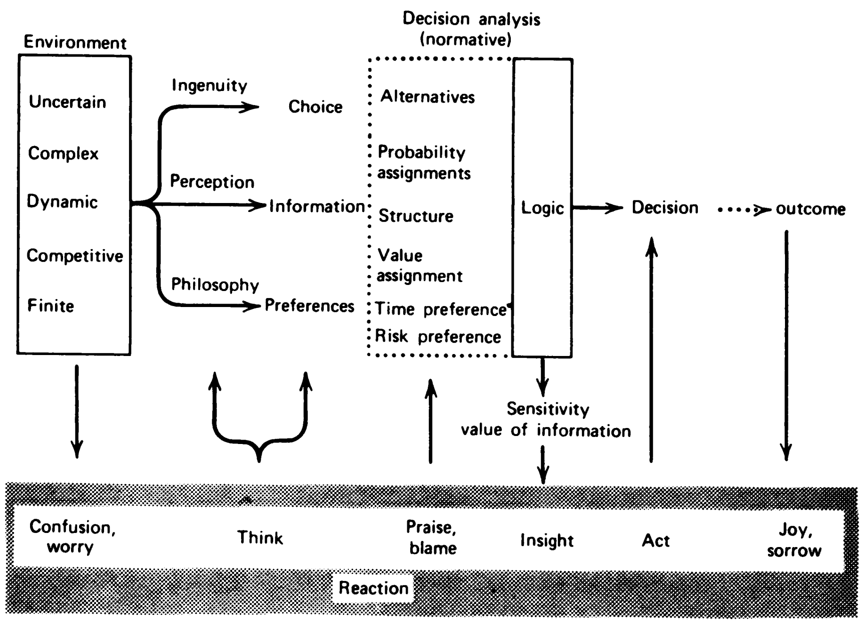

Figure 2 Decision-making using decision analysis.

a lower-level organization. The point is that alternatives are described quantitatively and not in general terms.

Incidentally, the process for developing new alternatives is one that we wish we were more able to comment upon. Ingenuity is required; there is just as much need for creativity in the process, as analyzed here, as there is in intuitive decision-making. The only advantage that decision analysis can bring to the search for new choices or alternatives is the same kind of help that any analytic model provides when brought to bear on a problem. For example, using engineering models, while not directly synthetic in most cases, can lead to insights into existing designs that suggest new alternatives, new ways of solving the engineering problem. The same thing can happen here, but it is not part of the formal structure.

The Encoding of Information. The first new thing comes about when we look at information. We represent information in two ways. We characterize uncertainty by means of probability assignments, and we represent relationships by means of models, that is, by structuring the problem. Let us talk about structuring first.

STRUCTURING. Structuring is the kind of "head bone connected to the neck bone" arrangement we find in physical models. But now we are talking about a decision model: a way of representing the underlying logical relationships of a decision problem—be it national, legal, industrial, or whatever—in a mathematical model that shows what affects what. This is something that is discussed in other parts of this book—the process of modelling and using computers in modelling. Although it is not customary to build formal models of decisions in the same way that we build formal models of other engineered systems, decision analysts do just that. In fact, anyone who is going to be a professional in this area is required to be conversant with modern modelling techniques.

Now let us return to the treatment of uncertainty.

PROBABILITY. Since many readers may not be familiar with the field of decision analysis, there is no reason to examine the long arguments that used to go on as to whether probability was a state of mind or a state of things. Decision analysts believe that it is a state of mind, a way of representing one person's uncertainty about a particular event or variable and that it has no necessary interpretation whatever in terms of real-world long-run frequencies.

The whole idea of describing uncertainty by means of probability assignments has come about only in the last ten or twenty years. Before that, probability was thought of as the province of statisticians—a region that only experts could enter. We may still need an expert to do more complex probabilistic manipulations, but we do not need an expert to think in probabilistic terms, which is what decision analysis requires. In some ways this is the most unique part of the decision analyst's trade, that he is able to deal effectively with the assessment and implications of uncertainty.

Lest the reader think this aspect is being overemphasized, here is a quote from a book that is in an entirely different field. It is a non-technical book called *The Search for Authenticity* by James F. T. Bugental who is a psychoanalyst. He writes:

> Let us pause to examine this quest for certainty. By "certainty" I mean the opposite of contingency. Having survived a disastrous fire in our neighborhood and being concerned about my home, I decide to investigate the likelihood I would not be so fortunate again. I find my odds are 1000 to 1 against the likelihood that my house will burn, but I am not content and so have the brush cleared back some distance. Now the odds are 1500 to 1, I find. Still concerned, I have an automatic sprinkling system installed. Now I'm told my odds are 3500 to 1. However I try, though, I must recognize always that I cannot achieve certainty that the house will not burn. I may do much, but I can't be sure that a nuclear firestorm will not make my efforts vain. I may build my house underground but still I can't be sure but that the earth might be drawn closer to the sun and the whole world thus be ignited.
>
> Now these are ridiculous extremes, of course, but the point remains: there is no true certainty to be had. So it is with any issue. Nevertheless, we seek that certainty constantly. We buy insurance, seat belts, medicines, door locks, education, and much else to try to protect ourselves against tragedy, to secure good outcomes. So long as we recognize we are dealing in probabilities, such choices can be useful. But every therapist has seen the pathology of seeking for certainty instead of better probabilities.[1]

One could take that last sentence, replace "therapist" by "decision analyst" and say that every decision analyst has seen the pathology of seeking for certainty instead of better probabilities. Every corporation the author has ever encountered believes that the secret of security comes from making things certain. In a course at Stanford University where students actually go out and do decision analysis as part of the course work, one of the presentations concerned the idea of "proven" reserves of mining ore. There we see the intent to make something certain when it is not certain, because a proven reserve is not a proven

[1]James F. T. Bugental, *The Search for Authenticity,* Holt, Rinehart and Winston, 1965, pp. 74–75.

reserve—it is something with probabilities attached to it, probabilities in terms of the amount, the type, the cost of extraction, and so forth.

Another example from class concerned how to treat people who have angina. Do you give them special new surgical procedures or do you give them conventional medication? Here the point came up again that there were some procedures that were "proven" medical treatments and others that were unproven. In other words, the world had again been divided into things that were OK and certain and approved by society, and others that were not OK, and there was a very small area in between. In these examples, as we all see in our own lives, we are continually trying to get the uncertainty out of the way because it is so painful with which to deal (as Bugental, the psychoanalyst, says).

We see the same thing in corporations. They attempt to set budgets, goals, and growth rates in an endeavor to ascertain what is basically uncertain. I claim that the way to corporate health is not to try to make the world certain, but to live with it in its present uncertain state, to act in the best possible way given the kind of world we live in. Bugental also sees that as the key to mental health, so I guess we agree even though we are in different fields.

We'll return later (see the question period at the end of this chapter) to how we go about making probability assignments. For the moment, let us just say that there is a way to do it and that such assignments become one of the two parts of the total encoding of information (the other being structure) that lead finally to putting into the model what we know.

So if we had to characterize the inputs to decision analysis, we would say choice is what we can do, and information is what we know. Now we come to the third: preferences—what we want.

The Establishment of Preferences. It turns out that because of our previous inability, or perhaps a better word would be reluctance, to deal with uncertainty, we have never gotten in most decision problems to the question of what we really want. It is a very interesting exercise to take a guy who has a tough decision because there is a lot of uncertainty in it, and ask him, "Well, suppose I eliminated all the uncertainty, suppose I told you for sure what was going to happen here and here and here, then, what would you like?" He often does not know. Think about it in terms of new possible states-of-being for the United States. If we could snap our fingers and have any state we wanted, which would we want?

Decision analysis separates uncertainty from the establishment of preferences. Once we tell a decision maker, "Look, let me worry about the uncertainty, that's my business." "We just encode that and do the best thing we can with the uncertainty that exists, given the structure and alternatives." "What do you really want?" "How much more is this outcome worth than that?" Then he has a problem he can work on constructively.

We break the idea of preferences up into three categories: The first kind of preference we call value assignment; the second, time preference; and the third, risk preference. Value assignment is concerned with situations where you have different outcomes and you say, "How much more is this outcome worth than that?" "What's the relative value of the two?"

VALUE ASSIGNMENT. Here is an example that may drive this topic home. Suppose we consider a medical case with which we can all identify. You walk into the doctor's office and he says, "You've got acute something-or-other and we're going to have to do this to you in the hospital, and you're going to be there for a day with severe pain, and then you'll be all right." And you say, "Well, what's severe pain?" And he says, "It's like pulling a wisdom tooth without an anesthetic." Each of us has his own opinion on whether this is a suitable torment of Hell, but at least you can think about what that outcome would mean to you—how joyful or sorrowful it would be. We have to allow for all tastes.

However, we have another alternative: to take a magical drug that will produce an instant painless cure for your malady. You have a choice—either a day in the hospital with pain and then cure, or the instant cure with the magical drug with no side effects (see Fig. 3). How much more would you pay for the instant cure via the magical drug compared to a day in the hospital with pain? Magical drugs are expensive, so let us see how much you would pay in addition to what the hospital trip would cost in order to obtain the drug. Would you pay a dollar? Sure, you'd pay a dollar. $10? Sure, that wisdom tooth is

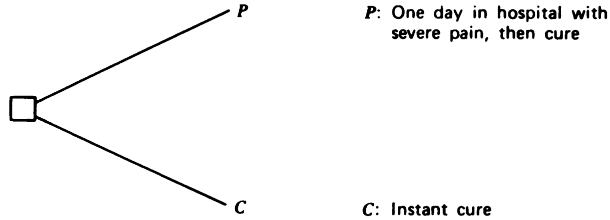

Figure 3 Value assignment.

pretty painful. How about $10,000? Your reply would probably be: "You must think I'm made of money." "Where am I going to get $10,000?"

So now we are bounding you. We may have a few millionaire readers who laugh at $10,000; we may have struggling students who say, "Give me ten bucks and take out all four wisdom teeth for all I care." But each of us in his own financial situation can say how much more he would pay for one than the other. Notice that we are making a decision out of the value question. We are saying, if you could go down each route, how much more would we have to make the instant cure cost before you would be indifferent. So one of the key ideas is to use the idea of comparison and adding positive increments to one side or the other until you say, "O.K., I cannot tell the difference." That is the value question. Given an outcome that occurs now with no uncertainty, how much do you like it?

TIME PREFERENCE.[4] The next question we face is time preference. **Time preference concerns the worth we place on values that are distributed over time.** This involves what we call the "greed-impatience trade-off." We are usually willing to accept less if we can get it sooner. Establishing the time preference of an individual or a corporation is not simple, but we can demonstrate that it is very important.

One case we worked on involved a person who had to choose between having an operation for a kidney transplant or being put on dialysis indefinitely with a kidney machine. It turned out that the whole question boiled down to time preference for him and, on further investigation, it developed that what was important to him was to live until his children got through college. So his time preference had an interesting structure. He placed a high value on being alive until some point x years in the future, and after that not so much. That is an unusual kind of time preference, but one any complete theory has to be able to accommodate.

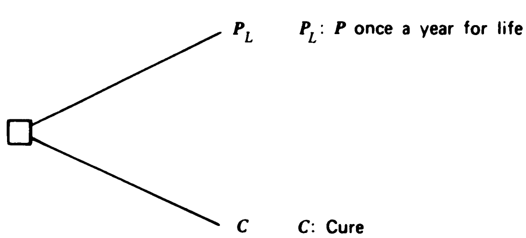

Figure 4 Time preference.

How can we demonstrate time preference in the medical example? In the medical case, suppose we think about a new event P_L, which is getting that one day in the hospital with pain once a year for life, as opposed to an instant cure now (see Fig. 4). Every year you have to go back to the hospital for one day and undergo the painful treatment or, on the other hand, you can get an immediate cure. It is clear that the instant cure is now worth more than it was before, since you were already going to be in for the first year anyway, but how much more depends on your attitude toward the future. If you say that anything that happens more than 30 days out you really do not care about, you will have one answer to your question. But, if you are very much concerned about retirement income, you will have another attitude. The way you answer this question will give us a lot of insight into how you value time.

Of course companies and nations face time preference questions. If we are thinking about setting up a national park system or other long-run investment, we are going to have to think about how much benefits in the future are worth relative to benefits today.

RISK PREFERENCE: Although value and time are certainly important, probably the most unusual aspect of the profession we are talking about is its ability to handle the third of the three, namely, risk preference. It is easy to demonstrate risk preference. Risk preference is the term we use to describe the fact that people are not expected-value decision makers; that is, they are not willing to choose among alternatives simply by comparing their expectations. (The expectation of an alternative is computed by multiplying each monetary prize by the probability of receiving it and then summing the products for all prizes.)

Suppose we said, "How many of us would flip a coin, double or nothing, for next year's income?" Whatever we would have gotten next year, we will either get twice as much, if we call the coin correctly, or nothing, if we call it incorrectly. Now that situation has an expected value equal to next year's salary, and anyone who is willing to make a decision on the basis of expected value should be marginally willing to engage in such a proposition. It is doubtful that we would have many takers, because there is nothing in it for a person. Suppose we pay each person a dollar to play; so now there is something in it for the taker. But most of us still would not do it. What if we say, "All right, what fraction of next year's salary would we have to pay to induce a person to engage in this gamble?" If we pay the fraction 1, then the taker has nothing to lose. He will get next year's salary anyway, and everyone will try it. If we pay the fraction 0, no one will try it. The real question is what fraction of next year's salary do we need to offer? Typically, numbers like 60-95 percent might be appropriate, 95 percent corresponding to the person who has substantial financial commitments and just does not see how he is going to make it, whereas the smaller fraction would correspond perhaps to the footloose bachelor who figures he can always go and live on the beach if he does not get any money next year. So it will be very specific to the person, to his own environment, his own tastes; and, in that sense, everything we are talking about is unique to the individual. It is appropriate to the decision maker and is not for the public at large.

To extend our medical discussion to this case, all we have to do is think about an imperfect magic drug. Unfortunately, the magic drug that might cure us will, now, also be able to kill us. So we will have to choose between going to the hospital with the day of pain or taking the magic drug, now costless, but which will kill us with probability p and cure us with probability $1-p$ with no pain—no side effects (see Fig. 5). The question is what is the probability p such that we are indifferent toward the day in the hospital with pain and taking the magic drug. Think about it; imagine being placed in this situation. It is not a very unrealistic situation; there are cases just like this that occur in medical practice.

What if $p = 1/1000$? One chance in a thousand we are going to die from the drug versus a day in the hospital with pain. The answer would probably be: "Dying's pretty bad, I don't like that." What about one chance in ten thousand? A typical reaction: "I'm feeling pretty lucky today—one in ten thousand, I might just do that." The point is, once we establish the value a person places on his life—which is another long story—and the value of a cure relative to a day in the

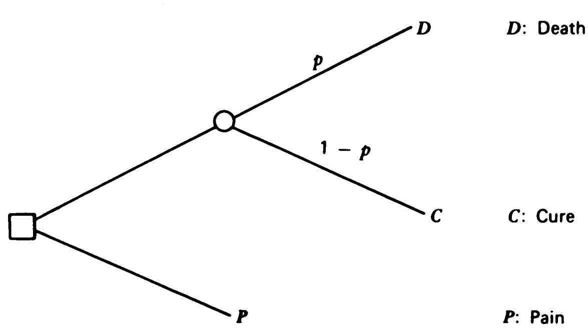

Figure 5 Risk preference.

hospital with pain, then the number p would certainly give us insight into their attitude toward risk and would allow us indeed to start building a description of their risk preference.

The Logical Decision. When this has been done, when we have carried out this procedure and have established preferences, the values placed on outcomes, the attitude toward time, the attitude toward risk (and there is a methodology for doing all of this), when we have established the models necessary for the decision one is making and have assessed probabilities as required on the uncertain variables, then we need nothing but logic to arrive at a decision. And a good decision is now very simply defined as the decision that is logically implied by the choices, information, and preferences that we have expressed. There is no ambiguity from that point on—there is only one logical decision.

This allows us to begin to assign praise or blame to the process of making the decision rather than to the ultimate outcome. We can do an analysis of the decision and make sure it is a high quality decision before we learn whether or not it produced a good outcome. This gives us many opportunities. It gives us the opportunity to revise the analysis—to look for weak spots in it—in other words, to tinker with it in the same way we can tinker with an engineering model of any other process.

The Value of Information

If this were all decision analysis did, it would be impressive enough, but from it we also get other benefits. We obtain sensitivities to the various features of the problem and we learn something that I think is unique to decision analysis called the "value of information." The value of information is what it would be worth to resolve uncertainty once and for all on one or more of the variables of the problem. In other words, suppose we are uncertain about something and do not know what to do. We postulate a person called a "clairvoyant." The clairvoyant is competent and truthful. He will tell us what is going to happen—for a price. The question is what should that price be. What can we afford to pay to eliminate uncertainty for the purpose of making this decision?

Of course we do not have real clairvoyants in the world—at least not very often—but the clairvoyant plays the same role in decision analysis as does the Carnot engine in thermodynamics. It is not the fact that we can or cannot make it, but that it serves as a bench mark for any other practical procedure against which it is compared. So the value of clairvoyance on any uncertainty represents an upper bound on what any information-gathering process that offers to shed light on the uncertainty might be worth.

For example, if we find in the medical problem that the value of clairvoyance on whether or not we are going to die from the drug is \$500, then that means that we should not pay more than \$500 for any literature search or anything else that would provide only imperfect information with respect to whether or not we are going to have this problem.

That is a revelation in itself to many people—the fact that one can establish a hard dollars and cents number on the value of information to us in making a decision, and hence can use that number to guide what information-gathering processes we might participate in.

The Medical Problem Evaluated. It is hard to demonstrate very simply how to do such a calculation, but let us try by taking the medical example and putting some numbers in it (see Fig. 6). The patient has the choice of taking the magic medicine or not. If he does not take it, then he is going to get the pain; we will consider that as a reference point of value \$0. If he does take the medicine, let us suppose he has one chance in a thousand of dying and 999 in a thousand of getting the instant cure. We have also put in numbers here saying that the cure is worth \$100 more than the pain. He is a relatively poor person, but he would pay \$100 more for the painless cure than he would for spending a painful day in the hospital. Now for death—what is the value of life to a person? This person has set the value of his life at \$100,000.

Notice that we "set" the value of his life. What is meant by this is that he wants the designers of public highway systems and airplanes to use the number \$100,000 in valuing his life. Why does he not make it a million dollars? If he does, he will have more expensive rides in airplanes, more expensive automobiles, and so forth. He does not get something for nothing. If he makes it too small, he had better be

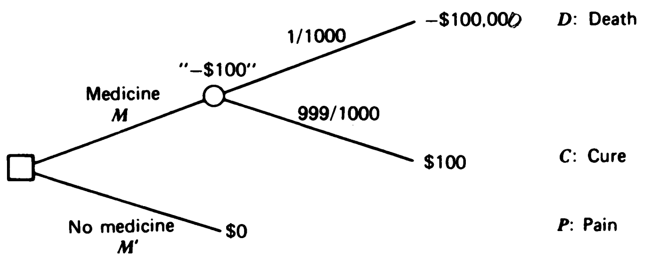

Figure 6 The medical decision.

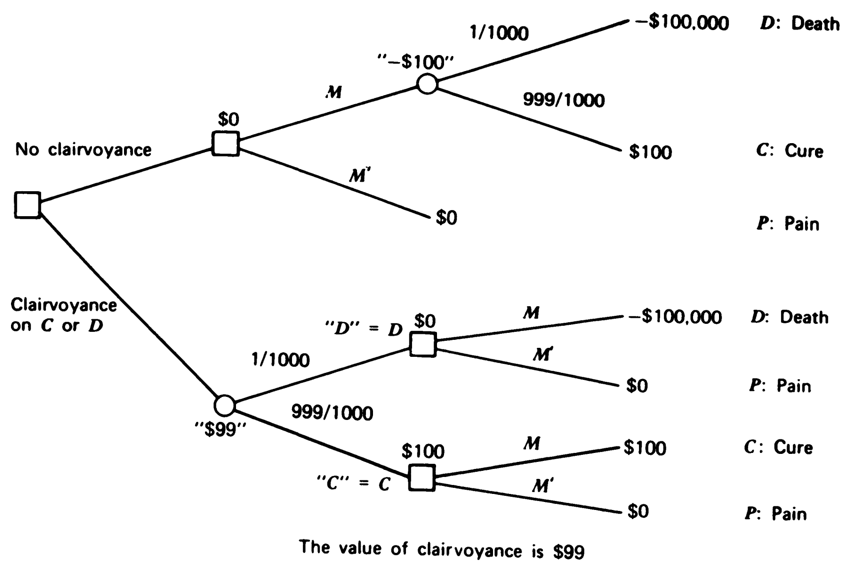

Figure 7 Value of clairvoyance computation.

wearing a helmet every time he enters his car. So it is a decision for him as to what number he wants the decision makers to use in this completely logical world that we are talking about.

The number -$100 in quotes (in Fig. 6) means that our patient has said that one chance in a thousand of losing $100,000 and 999 chances in a thousand of winning $100 has a value to him of −$100. In other words, we have to pay him $100 to get him to take on this uncertain proposition. It is clear that, comparing -$100 to $0, he is better off deciding not to take the medicine. So for him the probabilities, values, and attitude toward risk leading to the −$100 assessment of this whole uncertain proposition, the best decision is to forget about the medicine.

Clairvoyance. Now the clairvoyant arrives. If the individual we are talking about does not patronize the clairvoyant, then he does not take the medicine and makes nothing. If, on the other hand, he does buy the clairvoyance on the question of whether death will occur, what will happen? First, the clairvoyant will tell him whether he is going to die if he takes the medicine (see Fig. 7). We have "*D*" in quotes here, meaning that the clairvoyant says he is going to die, equivalent to his actually dying because the clairvoyant is truly prophetic. "*C*" means the clairvoyant says he is going to be cured. Since the probability the clairvoyant will say he is going to die has to be the same as the probability that he really will die, he has to assign one chance in a thousand to getting that report from the clairvoyant. Now suppose the clairvoyant says he is going to die. Obviously, he ought not to take the medicine in that case, and he will make nothing. If the clairvoyant says he is going to be cured without dying, then he is better off taking the medicine, and he will make $100. Since the payoff from the clairvoyant's saying that he is going to die is $0 and from not going to die is $100, and since there are 999 chances out of a 1000 that the clairvoyant will say he is not going to die, just by looking at that lottery we can see it will be worth almost $100 to him. He has 999 chances out of a 1000 of winning $100, and only one chance in 1000 in winning $0.

Let us suppose he evaluates the whole uncertain proposition at $99. If he does not buy the clairvoyance, he is looking at $0; if he does buy it, he is looking at a proposition that is worth about $99 to him. Thus, the value of the clairvoyance would be $99.

So here is an uncertain proposition with all kinds of big numbers running around in it, yet a very simple calculation based on his attitudes toward risk, life, death, and pain says he should not be willing to pay more than $99 to know for sure whether he would get the unfortunate event of death if he should take the drug.

Similarly, in any other decision problem—and there are some very, very complicated ones, involving many jointly-related variables—we can establish an upper bound on the value of information-gathering on any aspect of that problem. We can subsequently determine the best information-gathering strategy to precede the actual making of the decision.

The Decision Analysis Cycle

Let us begin with a word on methodology and then go on to an example. When doing a decision analysis it helps to organize your thoughts along the following lines. First, constructing a deterministic model of the problem and then measuring the sensitivity to each of the problem variables will reveal which uncertainties are important. Next, assessing probabilities on these uncertainties and establishing risk preference will determine the best decision. Finally, performing a value of clairvoyance analysis allows us to evaluate getting information on each of the uncertainties in the problem. The problem could be very complicated, involving many variables and months of modelling and analysis, but the basic logic is the same. The phases are: deterministic to evaluate sensitivities, probabilistic to find the best decision, and informational to determine in what direction new information would be most valuable. Of course you can repeat the process as many times as is economically valuable.

That is just to give an idea of how one does a professional decision analysis. Let us now turn to a case history to demonstrate the kind of problem that can be attacked in this way. Everything said so far has a naive ring to it. We can talk about betting on next year's salary, but we are really interested in not just the theory of decision analysis, but the practice of it.

A Power System Expansion Decision

Let us take an example from the public area. It concerns the planning of the electrical system of Mexico and is one of the largest decision analyses that has been done. It has been chosen because it comes closest to a problem in systems engineering. The specific question posed was: Should the Mexican electrical system install a nuclear plant and, if so, what should its policy toward nuclear plants in general be? Of course, we can not really answer that question without deciding how they are going to expand, operate, and price their system over time from here on out. So the real question is how to run the electrical system of Mexico for the rest of the century (see Fig. 8).

The Mexican electrical system is nationalized and very large—the size of several United States state-sized electrical systems. Because it is a complete national system, its planners have unique problems and also unique opportunities. The basic idea in working this problem was to look first of all at the various environmental factors that might influence the decision and then to look at the various measures of value that would result from particular methods of operation.

The Inputs

First, let us discuss the inputs. There are four input models: financial, energy, technology, and market. The financial models are concerned with the financial environment of the Mexican electrical system both in the world and the Mexican financial market. The inputs that these models provide are the amounts of money and the rates at which money can be borrowed from that source over time, with uncertainty if necessary. An input to this model is something called x which is picked up from the lower right. It is the book profit of the system. There is a feedback between the profitability of the system over time and the amount that it can borrow to support future

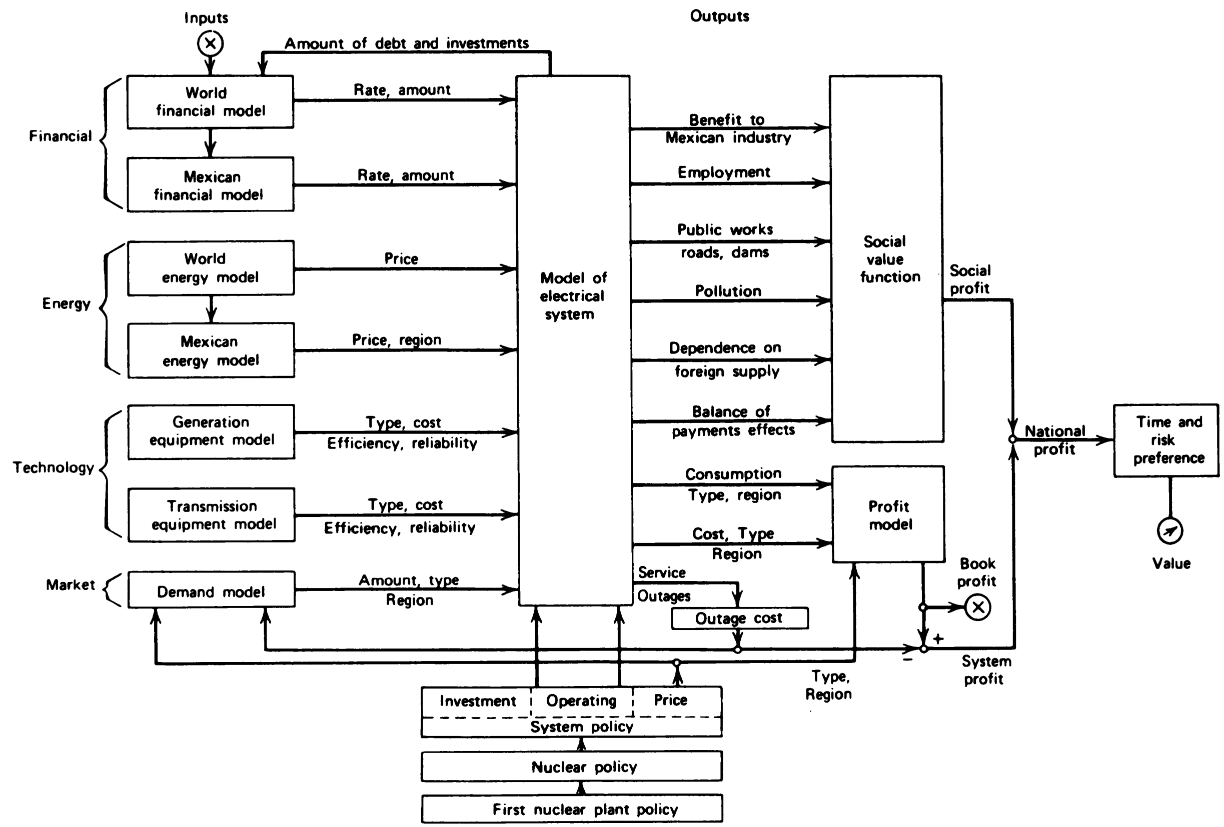

Figure 8 A decision analysis model of the Mexican electrical system.

expansion. The current amounts of debt and investment are also fed back.

The second type of input is energy costs, both in the world market and in the Mexican market. The interesting thing about Mexico is that it has just about every type of energy available: coal, oil, uranium, and thermal fields, and, of course, there are world markets in uranium and oil, at least, whose price movements over time would influence the economics of the Mexican system.

Next comes technology. This model describes generation and transmission equipment according to type, cost, efficiency, reliability. It includes such features as the advent of better reactors in the future and the possibilities of new and improved transmission systems which might make some of their remote hydro locations more desirable.

The last input model is the demand or market model, indicating by type and region the amount of electricity that would be consumed, given a pricing policy and given a quality of service. So these are the inputs to the model of the Mexican electrical system, which can then be run.

The Outputs

We will not go into the details of the rather sophisticated model which was prepared to describe operation and expansion of the Mexican electrical system. Of more interest in this discussion is the kind of outputs that were produced. There were the very logical ones of the consumption of electricity and the cost of producing the electricity by region to give a profit for the electrical system. This profit was what might be called the operating profit or book profit of the system, and is what the investor would see if he looked at the books of the Mexican electrical system. One modification to that profit which was considered was an economic penalty for system outages. A measure of the service provided by the system is added to the book profit to give something called system profit—which the investor does not see, but which the designer of the system does see. This penalty makes him unwilling to make a system that has outages for hours at a time, even though it might be more profitable if he looked only at the book profit.

The Social Value Function. But what is unusual about the outputs here is that many of them do not appear on the balance sheet of the corporation at all, but are what we might call social outputs; they enter into something called the social value function.

The decision maker in this case was the head of the Mexican electrical system. He felt many pressures on his position—not just the regular financial pressures of operating an electrical system, but social pressures coming about from the fact that his is a nationalized industry. For example, one of the things that was of concern to him was the benefit to Mexican industry. What would be the Mexican manufactured component of any system that might be installed? Another one was employment. How many Mexicans would be employed at what level if they went one route as opposed to another? Now we can see that the way we design the system is going to have major impacts on these kinds of outputs. If we have a nuclear system, then we might provide training for a few high-level technicians, but most of the components would be manufactured abroad; we do not have the army of Mexican laborers that we would if we built a hydro system in a remote location.

Another side effect is the public works that are produced by the generation choice. For example, with hydro you have roads and dams—that is access, flood control, and so on, that we would not have if we installed a large nuclear plant in the central valley of Mexico. Balance of payments is still another consideration. Mexico at that time had not devalued its currency; the currency was artificially pegged with respect to the free world rate. The question is, if we are going to have an import quota system to try to maintain this kind of disparity in the price of money, should we include that mechanism within the model or should we say other parts of the government are going to be responsible for making such adjustments. That is what the balance of payment effect is all about.

There are two outputs left that illustrate two different points. One is called dependence on foreign supply. At the time that this study started, there was a worry in the minds of the Mexicans that a nation supplying nuclear equipment might become hostile for some political reason and cut off the supply of repair parts, fuel, or maintenance facilities, much as the United States did with respect to Cuba. If that happened, of course Mexico would be in trouble. The question was, would this have a major effect on the decision, or would it not. They could buy insurance against it by stock-piling uranium until such time as they were able to establish alternate sources of supply. But it was a real worry, because they wanted to make sure they would be protected against any politically generated stoppage of equipment or supplies. By the end of the study, this whole area was of much less importance.

The other output was pollution. Originally the decision makers were not too interested in pollution. They said they could not afford to worry about it. And yet, if you have visited Mexico City, you know that atmospheric polution is very high. By the time this study was over, about one year later, they were very glad that they had provided a place in the model for pollution because they were now getting the same kind of citizen complaint that we get in the United States. Some of the things they were planning, like giant coal plants in the middle of Mexico City, were not acceptable any more.

The social outputs from the operating model entered the social value function to produce what we call "social profit." It represents social effects that do not appear on the balance sheet of the electrical system, per se. Social profit is combined with the system profit to produce national profit. Time and risk preference are expressed on national profit to give an evaluation of the system as a whole.

The problem that remained was to find a way to expand the Mexican electrical system that would produce the highest overall evaluation. Various optimization procedures were used to suggest installations of different types (gas turbines, nuclear, conventional, and hydro plants) to achieve this objective over the rest of the century.

The Nature of Policy

Let us briefly examine the question of what a policy for expansion of such a system means. A common policy in the past had been to establish a so-called plant list, which was a list of when each type of plant would be installed—in 1979 we are going to have an X-type plant in location Y. That is a little bit like asking a new father, "When is your son going to wear size-ten pants?" He could look at projected growth charts and say, "Well, I think it will be when he is nine years old." Another way to answer the question is to say, "Well, I will buy him size-ten pants when his measurements get into such and such a region." This is what we might call a closed-loop policy because we cannot say in advance when we are going to do it, but we have built a rule that will tell us the right time to do it.

So when we ask how is the system going to be expanded from here on out, no one can tell us: They can show us expected times for different things to happen, but indeed, only the program can determine what the effect on expansion of the future evolution of the system's environment will be. It has what we might call a self-healing property. If we foul it up by forcing it to put in a giant plant that it cannot immediately assimilate, then it is self-healing in the sense that it will delay and adjust the sizes and types of future plants until it gets back on the optimum track again. As a matter of fact, it is so much self-healing that it is hard to foul it up very much no matter what we do, because in the course of time it is a growing system that finds a way to get around any of our idiocies. In actuality, when they compared what this optimization system was doing with the designs produced by their conventional techniques using the same information, this system yielded superior results in every case.

The size of the Mexican study is interesting. It took approximately eight man-years, and was completed in one calendar year by a staff of decision analysts from the Stanford Research Institute Decision Analysis Group, plus four representatives of the Mexican Electricity Commission who were very competent in nuclear engineering and power system design. The programs and analyses are now being used in Mexico for continued planning of system expansion.

Other Applications

Other applications include industrial projects—should companies merge, should they bring out a new product, or should they bring a mine into production? All of these things are what we might call fairly conventional decision analyses by the criteria that we in the profession use.

Some interesting decision analyses have been done in the medical area, such as one recently performed on the treatment of pleural effusion, that is, water in the cavity between the lung and the chest wall. This was a one-year study done by a graduate student who, as far as the doctor (who was the lung expert) is concerned, completely encoded everything the doctor knew about pleural effusion. Later the doctor was asked if he developed this symptom would he prefer to be treated by this large decision model or by one of his colleagues. He said, without hesitation, he would rather use the model.

Another study that has just recently been completed is whether to seed a hurricane threatening the coast of the United States. It was based on a large experiment a few years ago on hurricane "Debbie" which indicated, but certainly not conclusively, that seeding a hurricane with silver iodide crystals would cause the wind to diminish about 15 percent. This in turn would lead to something like a 50 percent decrease in damage. The question now is—if you are the decision maker in the White House and here comes a big one, hurricane "Zazie," headed right for Miami—what do you do? Should you send the planes out to seed it knowing that, even so, there is a chance that it might get worse just because of natural causes and wipe out two cities instead of one? Or should you sit on your hands and possibly watch people get killed and property destroyed when they might have been saved? There is a tough problem. It has severe social impacts and is definitely a decision under uncertainty. Study of this problem was presented very recently to the President's Scientific Advisory Committee. They have formed a subpanel to see whether the conclusions should be put into effect.

Conclusion

We have tried to characterize what is a new profession—a profession that brings to the making of decisions the same kind of engineering concern and competence applied to other engineering questions. It seems fair to say that the profession has now come of age. We are able to work on virtually any decision where there is a decision maker who is worried about making that decision, regardless of the context in which it may arise. The only proviso is that the resources that he is allocating must be real world resources. We are not competent to allocate prayer because we can not get our hands on it—or love, which is infinite. But when it comes down to allocating money, or time, or anything else that a person or organization might have to allocate, this logic has a lot to be said for it. And indeed, as we have seen, the key is the idea of separating the good decision from the good outcome. Once we have done that then we have the same ability to analyze, to measure, to compare that gives strength to any other engineering discipline.

Question Period

QUESTION. Is the professional decision maker the man who is right out in the forefront making the decisions in his own name, or will there be a professional decision analyst who is like the ghost writer standing behind the man, the president, the corporate executive?

ANSWER. That is a good question. In the legal profession there is a maxim that the lawyer who defends himself in court has a fool for a client. And I think the same is true of decision analysis. I know that I would never want to be my own decision analyst because I am not detached. I want the answer to come out certain ways, subconsciously. For example, if I want to make a case for why I should buy a new stereo system, I will work like a dog to make sure that I have lots of variables in the analysis indicating that I am going to use it a lot, it is going to be very valuable to me, and it is not going to cost much. But when I bring in one of my friends who is a decision analyst, he will say, "Just wait a second." "How many days are there in the month?" "How many hours in a day?" "How often are you going to listen?" And pretty soon he has it down to size where I can say, "Yes, I am kidding myself, it just will not all fit together." So I think we will never get to the stage—nor should we—where the decision maker is the decision analyst. I think these are two very different roles and one can subvert the other.

What has turned out, however, is that some of the presidents of corporations who have been exposed to this kind of thing have begun to think the way that I am indicating here and to do very simple analyses on their own. And that is great. Everyone should know a little science, a little auto mechanics, and a little of everything, so they are able to do simple problems relatively well. But they should also realize that when they have a tough problem—one involving complexity, dynamics, modeling, and all the other things we have examined—then it is really time for a professional. We might take a medical analogy again. Most of the times when a person has a headache, aspirin is alright, but every once in a while it is a brain tumor and it is better not to take the aspirin. The important thing is to know the difference.

QUESTION. What about the systems analyst versus the decision analyst?

ANSWER. I see the decision analyst as the person who combines the complexity and the dynamic aspects of systems analysis with the ability to treat uncertainty and to measure preference—activities that are usually foreign to systems analysis. One of the problems with the systems field is that systems analysis is a much misunderstood term. Many groups and stakeholders in the systems professions have entirely different attitudes about what a systems analyst is. It can be everything from someone who riffles punched cards in a computing installation at one end of the spectrum to someone who know operations research, management science, and all of engineering rolled up into one. I do not know what a systems analyst is. He is somewhere in that spectrum, but I cannot say where.

QUESTION. Is there any information available on whether decision-making actually leads to statistically significant decisions?

ANSWER. I think you made a "no-no" there. Let us go back to the $5 and the $0–$100 coin toss. How can we measure in a one-shot decision what is statistically significant? This raises the issue of what view of probability we are taking. What we are saying is that the whole concept of statistical significance is pretty much irrelevant from the point of view of decision-making because we usually make decisions in one-shot situations. We cannot fire off a thousand Apollo rockets and see how many are going to succeed and how many are going to fail. We have to make a one-shot decision—do we go now, or do we not go now? And the question of statistical significance just does not come into it at all. I never find myself using those words. I find no use for them in making logical decisions.

QUESTION. Can you give some references for rating the qualitative effects of decisions?

ANSWER. You mean the so-called "intangible aspects"? There is a large amount of literature on the whole field. In general, what we say is that it is not a matter of tangibles and intangibles. If we take the Mexican example, people would say pollution is an intangible, or dependence on foreign supply is an intangible. But they are really tangibles. Why would you be willing to pay for things that are not tangibles? What we are saying is, let us take all the things that have value to you—positive and negative—and put values on them. In other words, if you would like to go out tonight and smell fresh air as opposed to smoggy air, let us talk about what that is worth to you. It is not worth $100 for you to do it for one night because you would go broke—that value would not be consistent with the other demands on you. But it is worth a penny, I will bet. Thus we begin to put dollars and cents values on what many people consider intangibles. Finally, we find ourselves making comparisons among values represented in dollars, not because dollars are in some sense the ultimate measure of everything, but because money is the Lagrange multiplier that our society has prepared for trading off one kind of thing against another kind of thing.

QUESTION. What are the axioms you must believe in order to reason this way?

ANSWER. Let me discuss just one of them. You must have transitive preferences, that is, you must reason such: if you like *A* better than *B* and *B* better than *C*, you must like *A* better than *C*. One of the points I was trying to make originally was that people often are not transitive. I might very well express to you intransitive preferences. But the question is, when you illustrate to me that I am being intransitive, do I like it or not? I do not like it, and the reason I do not like it is that someone can make a money pump out of me in that situation. I will switch *A* for *B* and *B* for *C*, and *C* back for *A*, all the while paying happily to make the transition—and he is just taking my money away, little by little. So the whole idea of intransitive preference is one of the things I do not want, and it is the cornerstone of what we do here, because the opposite of it is to be drained of your resources.

QUESTION. Could you comment further on the value of a life? Could you, for example, infer the value placed on life by observing the way corporations make decisions involving life?

ANSWER. Well, I have never done that. Of course there are studies all over the world on the value of a life. And unfortunately it varies greatly from one society to the next. But there was a comparative study I saw a few years ago indicating that at that time, for example, the value of a life was $100,000 in the United States, about $10,000 in Japan, and about $700 in South Vietnam.

There are many ways you can go about establishing a value for life. For example, you can examine the cash amounts awarded by juries for people killed in automobile accidents. The real issue is not what is the exact value of life, but rather are you being consistent in setting the value from one situation to another. The point is, life is precious, it is infinitely valuable. We are not talking about what you are willing to sell your life for—that is not the issue at all. The question is, what are you willing to buy it for. It is inconsistent to say a life is precious and, then, go out and not put the seat belt on when you get in your car. You are not being the same kind of person you would like to be at other times.

So what I like to do is pose a number for myself. (I cannot say what it should be for anyone else.) I want people to use this number as the value of my life when they make decisions that affect me. If I place it too high I will be running out of money very soon, because my car will weigh ten tons and will look like a tank. If I place it too low, I will not be able to venture outdoors.

You asked more specifically, could you determine from previous corporate decisions what value must have been placed on life. Well, first of all, I doubt that any corporation has ever established a number in the sense I am suggesting now. Perhaps they did it intuitively, but not explicitly. I would guess that if they did set such a value it would be $100,000 to $500,000 in the United States today.

QUESTION. Are the probabilities ever so hard to assign that you have to tell a client you cannot do an analysis for him?

ANSWER. No, I have never had that happen. Let me give you an example that arose in determining whether a new power nuclear reactor design should be introduced. The critical variable was the lifetime of the fuel cladding. The cladding was to be made from a material that had never been formulated before and, yet, the decision to go ahead on this new design would depend upon how this material performed.

The three people who were most knowledgeable in this company on the question of how long the material would

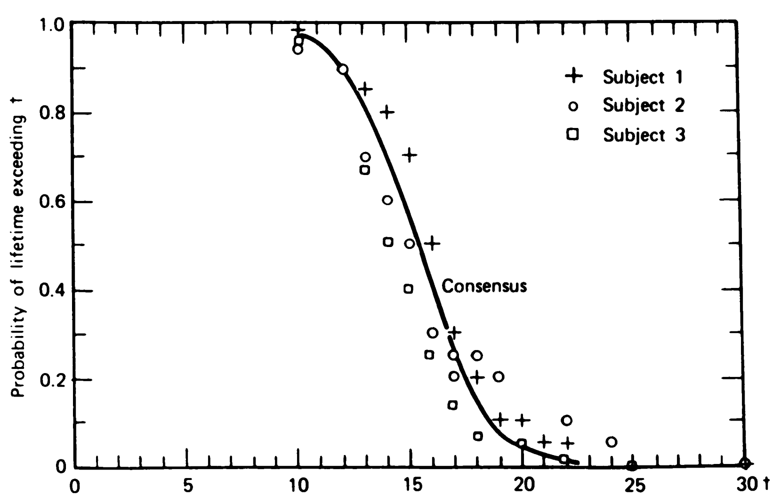

Figure 9 Priors on material lifetime.

last were assembled and told, "Look, we are going to have to come up with probability distributions on the life of this material, even though no one has ever built it before." They were interviewed separately, so they could not hear each other's answers. They gave the responses shown in Fig. 9. The first subject assessed about an 80 percent chance that the life would exceed 13. By an interview procedure, we developed the series of crosses. Then the next person (circles) was interviewed. He came up with the answers you see in Fig. 9, independently. Subject three's knowledge is represented by squares. Now when you look at these results, you find that they are remarkably similar. There are inconsistencies, however. For example, at one point in the questioning the subject represented by circles said that there was a 0.25 chance of exceeding a life of 17, and at another point, a 0.2 chance. There is a difference of 0.05 in probability. But when you think of how different these answers might have come out, I find their agreement remarkable.

Then the three men got back together again and started exchanging information. The last two subjects were relatively pessimistic about a short life compared to the first subject. He pointed out to them that certain things they were worried about were not really of concern because of experiments he had done recently. In other words, they exchanged information.

The same thing happened down at the far end. At the conclusion of the meeting they were willing to sketch the solid curve to represent the consensus of their opinion about how long this material would last. When this curve was shown to their boss, the manager of the whole operation, he for the first time felt that he was understanding what they were trying to communicate to him. He realized that they were not trying to be evasive, but were genuinely uncertain about what the life of this material would be, uncertain to the extent indicated by the consensus curve.

So, to answer your question, first I have never had the problem because you can always do it by means of a reference process. The basic idea is to say: would you rather win a million dollars, if the life of the material comes out more than 14; or win a million dollars, if a coin comes up "heads"? Then you can adjust the reference process, whether it is a coin coming up "heads" or a die coming up some number, until you are indifferent between the two processes.

But the point is that when a person is dealing with something every day of his life, when he is as expert at it as these subjects were, their answers are in much greater agreement than you would expect, if you asked them questions about the length of the Danube River, or other things they have not seen or thought about very much. What in the abstract seems as though it might be a problem, when you speculate about assigning probabilities, turns out to be not much of a problem, when you actually face a decision.

QUESTION. Can decision analysis tell whether it is worthwhile to do decision analysis?

ANSWER. You can do a "back-of-the-envelope" analysis, and then the question is how much would it be worth to do a more refined analysis. The decision to employ a decision analyst is, itself, a decision that can be analyzed like any other decision.

I have a rule of thumb that I have found helpful in my own life. I would like to spend at least one percent of the resources I am allocating on making sure that I am getting a good allocation of those resources. If I am going to buy a $2000 car, I want to spend at least $20 in making sure that I find the right car for me. Not making sure—in an absolute sense—but making sure in the sense that given the limits of my time and interest in the subject, I am doing the best I can.

In thinking about professional analyses, you should realize that decision analysts are in high demand. They command a premium professional salary. Consequently, decision analysis does not come cheap. However, we have not yet had a decision analysis that I am aware of where the decision maker did not feel very good about the insights he received relative to the money he spent.

QUESTION. Is it true that rather than removing intuition entirely from the process, you have simply removed it from the decision-making process and pushed it back to within the decision analysis?

ANSWER. Right. We are not eliminating judgment, or feelings, or opinions, or anything like that. Rather, we are quantifying them and putting them into a form where logic can operate upon them, rather than be buried in a man's mind where we cannot get access to them. This is very much a subjective and judgmental process in the sense that the probability assignments and all evaluations and preferences have to come from the decision maker. This is really just a matter of rendering unto Caesar the things that are Caesar's. What we say is, let the manager who makes the probability assignments and who has the preferences devote his time to making sure they represent his true feelings and his true attitudes. Let the logic be handled by the computer, which is eminently qualified to do such a job. It is really "divide and conquer". It sounds like a small thing, but the power of it is very great.

QUESTION. I do not see how a decision can be good unless the decision maker has good preferences.

ANSWER. Well, a decision is good if it is consistent with the decision maker's choices, information, and preferences. Looking on, I might think he is an idiot to make such a decision, but it is his prerogative. Our theory is amoral, in the sense that a person can go to Las Vegas, gamble, and lose money, but he says, "I have a ball there." "I value the experience very highly." Alright, given *his* values, he is making a statement with which I cannot disagree.

QUESTION. Would you comment on the balance between seeking new alternatives and analyzing the ones you have got?

ANSWER. Well, you cannot beat a new alternative. But how do you get a good new alternative? I find, for example, as an old engineer that there is nothing like doing an analysis of the existing design to see its weaknesses and to suggest improvements in that design. I think the same is true of a decision. We often find that we have two alternatives with the property that one is weak in one area, while the other is weak in a different area. Someone will suggest combining the two to create a new alternative with both good features. Often, this is feasible. These are new ideas that were not suggested originally by the individual who had the decision problem. There is no magical way of getting better alternatives by doing it this way, but it often turns out that creativity is a by-product of the process.

Question Period

QUESTION. What makes you believe that you obtain better outcomes by decision analysis than you would by following intuition?

ANSWER. It is really an act of faith. Let us take the case of the man with the \$5 payment for the \$0–\$100 coin toss. I say to him, "That is a good decision." "I have looked into your finances, and we agree it is a good decision." And he says, "Yes". He calls the coin and he misses, because after all he still has a probability of one-half of failing. Whereas some other person says, "I have looked into the Swami's eyes and I know I must call 'heads'." So he calls "heads" and wins. Which is the better decision-making procedure? It is really an act of faith that a logical procedure based on principles you believe in is better than another procedure. We can never prove that someone who appeals to astrology is acting in any way inferior to what we are proposing. It is up to you to decide whose advice you would seek.

QUESTION. Is it always possible to get better, more complete information and, hence, make a better decision?

ANSWER. Not always. For example, if a major hurricane bears down on Miami in the next hurricane season, where are you going to get more complete information? The decision will have to be made with the presently available information.

One of the persistent features of human nature is this quest for certainty. If anything too much money is spent on information-gathering, rather than too little. We keep pursuing this "Holy Grail" of certainty, instead of trying to find better alternatives or just making the decision and getting on to something else. I see the whole move toward data bases as symtomatic of this desire—this quest for certainty, hopeless as it is.

QUESTION. Is not a major function of an executive the ability to recognize when a decision has to be made?

ANSWER. Yes. For example, a president of a major company faces the decision of introducing a new product. He knows he has the decision, he is worried about it, and he does not know what to do. He has complex alternatives which are not easily evaluated, and he knows that intuition is not going to be much help to him. Therefore, he calls on a decision analyst to sort out the alternatives, get probabilities assigned, build models, present lotteries for his inspection, and so on.

QUESTION. The decision analyst will create the probabilities, will he not?

ANSWER. No, the probabilities do not come from the analyst but from the decision maker, his experts, and possibly external experts. I am not an expert on hurricanes; I am not an expert in medical problems; I am not an expert in the Mexican electrical system's rate of growth, or anything like that. God forbid, we should try to become experts in all the different things we work on. But what people in this profession are expert in is the modeling of the decision problem and the extraction of information from experts and preferences from decision makers in order to develop a better decision. It is a very careful separation of function.

QUESTION. If I were to adopt this approach and apply it to a variety of different decisions and if, after a while, I were to discover that this led to favorable outcomes less frequently than my old usual approach, then I would be forced to conclude that you were giving us a rather esoteric meaning of the word "good."

ANSWER. That is what you would be forced to conclude.

QUESTION. But does it not have to lead to more good outcomes, if it is to have practical value?

ANSWER. The question is: would you make the same decision if you faced the same situation again without knowing how it was going to turn out. I think it is a good decision to pay $5 for the $0–$100 coin toss. I would even be willing to purchase several of them. Suppose I keep losing. I would look at the coin, consider all kinds of hypotheses about people cheating me, two-headed coins, and sleight-of-hand, but suppose I am convinced there really is no "hanky-panky" going on. Well, then, I would not depart from this theory. I would say, "O.K., it is still a good decision, give me another one," even though I had lost five or six in a row—which is not an unlikely event. But I am going to stay with this theory until I find a better one, and I have not found one yet.

QUESTION. But should our goal not be to maximize the likelihood of good outcomes?

ANSWER. Of course. We all want joy. We all want good outcomes. Let that be stipulated right now. Everyone wants good rather than bad, more rather than less—the question is how do we get there. The only thing you can control is the decision and how you go about making that decision. That is the key. When you focus on that, I think you will want to do it the way we have discussed.

BIBLIOGRAPHY

Ginsberg, Allen S., *Decision Analysis in Clinical Patient Management with an Application to the Pleural Effusion Problem*, Ph.D. dissertation, Stanford University, July 1969.

Howard, Ronald A., "Decision Analysis: Applied Decision Theory," *Proceedings of the Fourth International Conference on Operational Research*, pp. 55–71, Wiley, 1966.

Howard, Ronald A., *Dynamic Probabilistic Systems*, Vol. 1, *Markov Models;* Vol. 2, *Semi-Markov and Decision Processes*, Wiley, 1971.

Howard, Ronald A., *Dynamic Programming and Markov Processes*, The M.I.T. Press, 1960.

Howard, Ronald A. (Ed.), "Special Issue on Decision Analysis," IEEE Transactions on *Systems Science and Cybernetics*, Vol. SSC-4, No. 3, September 1968.

Howard, R. A., J. E. Matheson, and D. W. North, "The Decision to Seed Hurricanes," in *Science*, (4040), **176**, 1191-1202, June 17, 1972.

Luce, R. Duncan and Howard Raiffa, *Games and Decisions*, Wiley, 1957.

Raiffa, Howard, *Decision Analysis: Introductory Lectures on Choices under Uncertainty*, Addison-Wesley, 1968.

Savage, L. J., *The Foundations of Statistics*, Wiley, 1954.

Schlaifer, Robert, *Analysis of Decisions Under Uncertainty*, McGraw-Hill, 1969.

Stanford Research Institute, Decision Analysis Group, *Decision Analysis of Nuclear Plants in Electrical System Expansion*, prepared for Comision Federal de Electricidad, Mexico City, SRI Project 6496, December 1968.

DECISION ANALYSIS: APPLIED DECISION THEORY

Ronald A. Howard

Department of Engineering-Economic Systems
Stanford University

Reprinted with permission from Proceedings of the Fourth International Conference on Operational Research (D. B. Hertz and J. Melese, eds.) New York: Wiley-Interscience, 1966. pp. 55-71

DECISION ANALYSIS: APPLIED DECISION THEORY

Analyse des Décisions: Théorie Appliquée des Décisions

RONALD A. HOWARD

Institute in Engineering-Economic Systems
Stanford University, California
United States of America

1. INTRODUCTION

Decision theory in the modern sense has existed for more than a decade. Most of the effort among the present developers of the theory has been devoted to Bayesian analysis of problems formerly treated by classical statistics. Many practical management decision problems, however, can be handled by formal structures that are far from novel theoretically. The world of top management decision making is not often structured by simple Bernoulli, Poisson, or normal models.

Indeed, Bayes's theorem itself may not be so important. A statistician for a major company wrote a report in which he commented that for all the talk about the Bayesian revolution he did not know of a single application in the company in which Bayes's theorem was actually used. The observation was probably quite correct—but what it shows by implication is that the most significant part of the revolution is not Bayes's theorem or conjugate distributions but rather the concept of probability as a state of mind, a 200-year-old concept. Thus the real promise of decision theory lies in its ability to provide a broad logical basis for decision making in the face of uncertainty rather than in any specific models.

The purpose of this article is to outline a formal procedure for the analysis of decision problems, a procedure that I call "decision analysis." We shall also discuss several of the practical problems that arise when we attempt to apply the decision analysis formalism.

2. DECISION ANALYSIS

To describe decision analysis it is first necessary to define a decision. A decision is an irrevocable allocation of resources, irrevocable in the sense that it is impossible or extremely costly to change back to the situation that existed before making the decision. Thus for our purposes a decision is not a mental commitment to follow a course of action but rather the actual pursuit of that course of action. This definition often serves to identify the real decision maker within a loosely structured organization. Finding the exact nature of the decision to be made, however, and who will make it, remains one of the fundamental problems of the decision analyst.

Having defined a decision, let us clarify the concept by drawing a necessary distinction between a good decision and a good outcome. A good decision is a logical decision—one based on the uncertainties, values, and preferences of the decision maker. A good outcome is one that is profitable or otherwise highly valued. In short, a good outcome is one that we wish would happen. Hopefully, by making good decisions in all the situations that face us we shall ensure as high a percentage as possible of good outcomes. We may be disappointed to find that a good decision has produced a bad outcome or dismayed to learn that someone who has made what we consider to be a bad decision has enjoyed a good outcome. Yet, pending the invention of the true clairvoyant, we find no better alternative in the pursuit of good outcomes than to make good decisions.

Decision analysis is a logical procedure for the balancing of the factors that influence a decision. The procedure incorporates uncertainties, values, and preferences in a basic structure that models the decision. Typically, it includes technical, marketing, competitive, and environmental factors. The essence of the procedure is the construction of a structural model of the decision in a form suitable for computation and manipulation; the realization of this model is often a set of computer programs.

2.1. The Decision Analysis Procedure

Table 1 lists the three phases of a decision analysis that are worth distinction: the deterministic, probabilistic, and post-mortem phases.

TABLE 1

The Decision Analysis Procedure

I. Deterministic phase
   1. Define the decision
   2. Identify the alternatives
   3. Assign values to outcomes
   4. Select state variables
   5. Establish relationship at state variables
   6. Specify time preference

   Analysis: (a) Determine dominance to eliminate alternatives
   (b) Measure sensitivity to identify crucial state variables

II. Probabilistic phase
   1. Encode uncertainty on crucial state variables
      Analysis: Develop profit lottery
   2. Encode risk preference
      Analysis: Select best alternative

III. Post-mortem phase
   Analysis: (a) Determine value of eliminating uncertainty in crucial state variables
   (b) Develop most economical information-gathering program

2.1.1. *The Deterministic Phase*

The first step in the deterministic phase is to answer the question, "What decision must be made?" Strange as it may seem, many people with what appear to be decision problems have never asked themselves that question. We must distinguish between situations in which there is a decision to be made and situations in which we are simply worried about a bad outcome. If we have resources to allocate, we have a decision problem, but if we are only hand wringing about circumstances beyond our control no formal analysis will help. The difference is that between selecting a surgeon to operate on a member of your family and waiting for the result of the operation. We may be in a state of anguish throughout, but decision analysis can help only with the first question.

The next step is to identify the alternatives that are available, to answer the question, "What courses of action are open to us?" Alternative generation is the most creative part of the decision analysis procedure. Often the introduction of a new alternative eliminates the need for further formal analysis. Although the synthesis of new alternatives necessarily does not fall within the province of the decision analysis procedure, the procedure does evaluate alternatives and thereby suggests the defects in present alternatives that new alternatives might remedy. Thus the existence of an analytic procedure is the first step toward synthesis.

We continue the deterministic phase by assigning values to the various outcomes that might be produced by each alternative. We thus answer the question, "How are you going to determine which outcomes are good and which are bad?" In business problems this will typically be a measure of profit. Military and governmental applications should also consider profit, measured perhaps with more difficulty, because these decision makers are also allocating the economic resources of the nation. Even when we agree on the measure of profit to be assigned to each outcome, it may be difficult to make the assignment until the values of a number of variables associated with each outcome are specified. We call these variables the state variables of the decision. Their selection is the next step in the deterministic phase.

A typical problem will have state variables of many kinds: costs of manufacture, prices charged by competitors, the failure rate of the product, etc. We select them by asking the question, "If you had a crystal ball, what numerical questions would you ask it about the outcome in order to specify your profit measure?" At the same time that we select these variables we should assign both nominal values for them and the range over which they might vary for future reference.

Next we establish how the state variables are related to each other and to the measure of performance. We construct, in essence, a profit function that shows how profit is related to the factors that underlie the decision. The construction of this profit function requires considerable judgment to avoid the twin difficulties of excessive complexity and unreal simplicity.

If the results of the decision extend over a long time period, it will be necessary to have the decision maker specify his time preference for profit. We must ask, "How does profit received in the future compare in value to profit received today?" or an equivalent question. In cases in which we can assume a perfect financial environment the present value of future profit at some rate of interest will be the answer. In many large decision problems, however, the nature of the undertaking has an effect on the basic financial structure of the enterprise. In these cases a much more realistic modeling of the time preference for profit is necessary.

Now that we have completed the steps in the deterministic phase we have a deterministic model of the decision problem. We next perform two closely related analyses. We perform them by setting the state variables to their nominal values and then sweeping each through its range of values, individually and jointly, as judgment dictates. Throughout this process we observe which alternative would be best and how much value would be associated with each alternative. We often observe that regardless of the values the state variables take on in their ranges one alternative is always superior to another, a condition we describe by saying that the first alternative dominates the second. The principle of dominance may often permit a major reduction in the number of alternatives that need be considered.

As a result of this procedure we have performed a sensitivity analysis on the state variables. We know how much a 10 percent change in one of the variables will affect profit, hence the optimum alternative. Similarly, we know how changes in state variables may interact to affect the decision. This sensitivity analysis shows us where uncertainty is important. We identify those state variables to which the outcome is sensitive as "crucial" state variables. Determining how uncertainties in the crucial state variable influence the decision is the concern of the probabilistic phase of the decision analysis.

2.1.2. Probabilistic Phase

The probabilistic phase begins by encoding uncertainties on each of the crucial state variables; that is, gathering priors on them. A subset of the crucial state variables will usually be independent—for these only a single probability distribution is necessary. The remainder will have to be treated by collecting conditional as well as marginal distributions. We have more to say on this process later.

The next step is to find the uncertainty in profit for each alternative implied by the functional relationship of profit to the crucial state variables and the probability distribution on those crucial state variables for the alternative. We call this derived probability distribution of profit the profit lottery of the alternative. In a few cases the profit lottery can be derived analytically and in many by numerical analysis procedures. In any case it may be approximated by a Monte Carlo simulation. Regardless of the procedure used, the result is a probability distribution on profit (or perhaps on discounted profit) for each of the alternatives that remain in the problem.

Now we must consider how to choose between two alternatives with different profit lotteries. In one case the choice is easy. Suppose that we plot the profit lottery for each alternative in complementary cumulative form; that is, plot the

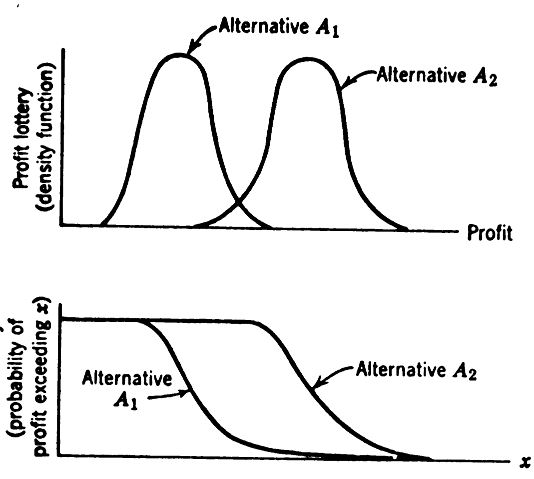

Figure 1. Stochastic dominance.

probability of profit exceeding x for any given x. Suppose further, as shown in Figure 1, that the complementary cumulative for alternative A_2 always lies above that for alternative A_1. This means that for any number x there is a higher probability of profit exceeding that number with alternative A_2 than with alternative A_1. In this case we would prefer alternative A_2 to alternative A_1, provided only that we liked more profit better than less profit. We describe this situation by saying that the profit from alternative A_2 is stochastically greater than the profit from alternative A_1 or equivalently by saying that alternative A_2 stochastically dominates alternative A_1. Stochastic dominance is a concept that appeals intuitively to management; it applies in a surprising number of cases.

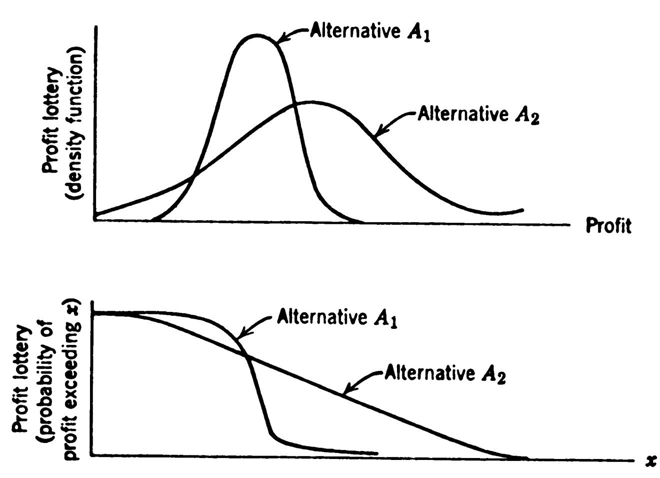

Figure 2. Lack of stochastic dominance.

Figure 2, however, illustrates a case in which stochastic dominance does not apply. When faced with a situation like this, we must either abandon formal methods and leave the selection of the best alternative to judgment or delve into the measurement of risk preference. If we choose to measure risk preference, we begin the second step of the probabilistic phase. We must construct a utility function for the decision maker that will tell us whether or not, for example, he would prefer a certain 4 million dollars profit to equal chances of earning zero or 10 million dollars. Although these questions are quite foreign to management, they are being asked increasingly often with promising results. Of course, when risk preference is established in the form of a utility function, the best alternative is the one whose profit lottery has the highest utility.

2.1.3. Post-Mortem Phase

The post-mortem phase of the procedure is composed entirely of analysis. This phase begins when the best alternative has been selected as the result of the probabilistic phase. Here we use the concepts of the clairvoyant lottery to establish a dollar value of eliminating uncertainty in each of the state variables individually and jointly. Being able to show the impact of uncertainties on profit is one of the most important features of decision analysis. It leads directly to the next step of the post-mortem, which is finding the most economical information-gathering program, if, in fact, it would be profitable to gather more information. The information-gathering program may be physical research, a marketing survey, or the hiring of a consultant. Perhaps in no other area of its operations is an enterprise in such need of substantiating analysis as it is in the justification of information-gathering programs.

Of course, once the information-gathering scheme, if any, is completed, its information modifies the probability distributions on the crucial state variables and consequently affects the decision. Indeed, if the information-gathering program were not expected to modify the probability distributions on the crucial state variables it would not be conducted. We then repeat the probabilistic phase by using the new probability distributions to find the profit lotteries and then enter the post-mortem phase once more to determine whether further information gathering is worthwhile. Thus the decision analysis is a vital structure that lets us compare at any time the values of such alternatives as acting, postponing action and buying information, or refusing to consider the problem further. We must remember that the analysis is always based on the current state of knowledge. Overnight there can arrive a piece of information that changes the nature of the conclusions entirely. Of course, having captured the basic structure of the problem, we are in an excellent position to incorporate any such information.

Finally, as the result of the analysis the decision maker embarks on a course of action. At this point he may be interested in the behavior of several of the state variables for planning purposes; for example, having decided to introduce a new product, he may want to examine the probability distributions for its sales in future years to make subsidiary decisions on distribution facilities or on the size of the sales force. The decision-analysis model readily provides such planning information.

2.2. The Advantages of Decision Analysis

Decision analysis has many advantages, of which we have described just a few, such as its comprehensiveness and vitality as a model of the decision and its ability to place a dollar value on uncertainty. We should point out further that the procedure is relevant to both one of a kind and repetitive decisions. Decision analysis offers the operations research profession the opportunity to extend its scope beyond its traditional primary concern with repetitively verifiable operations.

One of the most important advantages of decision analysis lies in the way it encourages meaningful communication among the members of the enterprise because it provides a common language in which to discuss decision problems. Thus engineers and marketing planners with quite different jargons can appreciated one another's contributions to a decision. Both can use the decision-analysis language to convey their feelings to management quickly and effectively.

A phenomenon that seems to be the result of the decision-analysis language is the successive structuring of staff groups to provide reports that are useful in decision-analysis terms. Thus, if the decision problem being analyzed starts in an engineering group, that group ultimately seeks inputs from marketing, product planning, the legal staff, and so on, that are compatible with the probabilistic analysis. Soon these groups begin to think in probabilistic terms and to emphasize probabilistic thinking in their reports. The process seems irreversible in that once the staff of an organization becomes comfortable in dealing with probabilistic phenomena they are never again satisfied with deterministic or expected value approaches to problems. Thus the existence of decision-analysis concepts as a language for communication may be its most important advantage.

2.3. The Hierarchy of Decision Analysis

It is informative to place decision analysis in the hierarchy of techniques that have been developed to treat decision problems. We see that a decision analysis requires two supporting activities. One is a lower order activity that we call alternative evaluation; the second, a higher order activity that we call goal setting. Performing a decision analysis requires evaluating alternatives according to the goals that have been set for the decision. The practitioners of operations research are quite experienced in alternative evaluation in both industrial and military contexts. In fact, in spite of the lip service paid to objective functions, only rare operations researchers have had the scope necessary to consider the goal-setting problems.

All mankind seems inexpert at goal setting, although it is the most important problem we face. Perhaps the role of decision analysis is to allow the discussion of decisions to be carried on at a level that shows the explicit need for goals or criteria for selection of the best alternative. We need to make goals explicit only if the decision maker is going to delegate the making of the decision or if he is unsure of his ability to be consistent in selecting the best alternative. We shall not comment on whether there is a trend toward more or less delegation of decision making. However, it is becoming clear to those with decision-making responsibilities that the increasing complexity of the operations under their control requires correspondingly more formal approaches to the problem of organizing the information that bears on a decision if inconsistent decisions are to be avoided.

The history of the analysis of the procurement of military weapons systems points this out. Recent years have shown the progression of procurement thinking from effectiveness to cost effectiveness. In this respect the military authorities have been able to catch up in their decision-making apparatus to what industry had been doing in its simpler problems for years. Other agencies of government are now in the process of making the same transition. Now all must move on to the inclusion of uncertainty, to the establishment of goals that are reflected in risk and time preferences.

These developments are now on the horizon and in some cases in sight; for example, although we have tended to think of the utility theory as an academic pursuit, one of our major companies was recently faced with the question, "Is 10 million dollars of profit sufficient to incur one chance in 1 million of losing 1 billion dollars?" Although the loss is staggering, it is realistic for the company concerned. Should such a large company be risk-indifferent and make decisions on an expected value basis? Are stockholders responsible for diversifying their risk externally to the company or should the company be risk-averting on their behalf? For the first time the company faced these questions in a formal way rather than deciding the particular question on its own merits and this we must regard as a step forward.

Decision analysis has had its critics, of course. One said, "In the final analysis, aren't decisions politically based?" The best answer to that came from a high official in the executive branch of our government who said, "The better the logical basis for a decision, the more difficult it is for extraneous political factors to hold sway." It may be discouraging in the short run to see logic overridden by the tactical situation, but one must expect to lose battles to win the war.

Another criticism is, "If this is such a good idea, why haven't I heard of it before?" One very practical reason is that the operations we conduct in the course of a decision analysis would be expensive to carry out without using computers. To this extent decision analysis is a product of our technology. There are other answers, however. One is that the idea of probability as a state of mind and not of things is only now regaining its proper place in the world of thought. The opposing heresy lay heavy on the race for the better part of a century. We should note that most of the operations research performed in World War II required mathematical and probabilistic concepts that were readily available to Napoleon. One wonders about how the introduction of formal methods for decision making at that time might have affected the course of history.

3. THE PRINCIPLES OF THE DECISION ANALYST

Next we turn to the principles of the decision analyst, the professional who embarks on preparing a decision analysis. His first principle is to identify and isolate the components of the decision—the uncertainty, risk aversion, time preference, and problem structure. Often arguments over which is the best decision arise because the participants do not realize that they are arguing on different grounds. Thus it is possible for A to think that a certain alternative is riskier than it is in B's opinion, either because A assigns different probabilities to the outcomes than B but both are equally risk-averting, or because A and B assign the same probabilities to the outcomes but differ in their risk aversion. If we are to make progress in resolving the argument, we must identify the nature of the difficulty and bring it into the open. Similar clarifications may be made in the areas of time preference or in the measurement of the value of outcomes.

One aid in reducing the problem to its fundamental components is restricting the vocabulary that can be used in discussing the problem. Thus we carry on the discussion in terms of events, random variables, probabilities, density functions, expectations, outcomes, and alternatives. We do not allow fuzzy thinking about the nature of these terms. Thus "The density function of the probability" and "The confidence in the probability estimate" must be nipped in the bud. We speak of "assigning," not "estimating," the probabilities of events and think of this assignment as based on our "state of information." These conventions eliminate statements like the one recently made on a TV panel of doctors who were discussing the right of a patient to participate in decision making on his treatment. One doctor asserted that the patient should be told of "some kind of a chance of a likelihood of a bad result." I am sure that the doctor was a victim of the pressures of the program and would agree with us that telling the patient the probability the doctor would assign to a bad result would be preferable.

One principle that is vital to the decision analyst is professional detachment in selecting alternatives. The analyst must not become involved in the heated political controversies that often surround decisions except to reduce them to a common basis. He must demonstrate his willingness to change the recommended alternative in the face of new information if he is to earn the respect of all concerned. This professional detachment may, in fact, be the analyst's single most valuable characteristic. Logic is often severely strained when we are personally involved.

The detachment of the analyst has another positive benefit. As an observer he may be able to suggest alternatives that may have escaped those who are intimately involved with the problem. He may suggest delaying action, buying insurance, or performing a test, depending on the nature of the decision. Of course, the comprehensive knowledge of the properties of the existing alternatives that the decision analyst must gain is a major aid in formulating new alternatives.

Since it is a rare decision that does not imply other present and future decisions, the decision analyst must establish a scope for the analysis that is broad enough to provide meaningful answers but not broad enough to impose impractical computational requirements. Perhaps the fundamental question in establishing scope is how much to spend on decision analysis. Because the approach could be applied both to selecting a meal from a restaurant menu and to allocating the federal budget, the analyst needs some guidelines to determine when the analysis is worthwhile.

The question of how much decision analysis is an economic problem susceptible to a simpler decision analysis, but rather than pursue that road let us pose an arbitrary and reasonable but indefensible rule of thumb: spend at least 1 percent of the resources to be allocated on the question of how they should be allocated. Thus, if we were going to buy a 2000-dollar automobile, the rule indicates a 20-dollar analysis, whereas for a 20,000-dollar house it would specify a 200-dollar analysis. A 1-million-dollar decision would justify 10,000 dollars' worth of analysis or, let us say, about three man-months. The initial reaction to this guideline has been that it is conservative in the sense of not spending much on analysis; yet, when we apply it to many decisions now made by business and government, the reaction is that the actual expenditures on analysis are only one-tenth or one-hundredth as large as the rule would prescribe. Of course, we can all construct situations in which a much smaller or larger expenditure than given by the rule would be appropriate, and each organization can set its own rule, perhaps making the amount spent on analysis nonlinear in the resources to be allocated. Nevertheless, the 1 percent figure has served well to illustrate where decision analysis can be expected to have the highest payoff.

The professional nature of the decision analyst becomes apparent when he balances realism in the various parts of the decision-analysis model. Here he can be guided only by what used to be called engineering judgment. One principle he should follow is to avoid sophistication in any part of the problem when that sophistication would not affect the result. We can describe this informally by saying that he should strive for a constant "wince" level as he surveys all parts of the analysis. One indication that he has achieved this state is that he would be torn among many possibilities for improvement if we allowed him to devote more time and resources to the decision model.

4. THE ENCODING OF SUBJECTIVE INFORMATION

One unique feature of decision analysis is the encoding of subjective information, both in the form of risk aversion and in the assignment of probabilities.

4.1. Risk Aversion and Time Preference

Since we are dealing in most cases with enterprises rather than individuals, the appropriate risk aversion and time preference should be that of the enterprise. The problem of establishing such norms is beyond our present scope. It is easy, however, to demonstrate to managers, or to anyone else for that matter, that the phenomenon of risk aversion exists and that it varies widely from individual to individual. One question useful in doing this is, "How much would you have to be paid to call a coin, double or nothing, for next year's salary?" Regardless of the salary level of the individuals involved, this is a provocative question. We point out that only a rare individual would play such a game for a payment of zero and that virtually everyone would play for a payment equal to next year's salary, since then there would be nothing to lose. Thereafter we are merely haggling over the price. Payments in the range of 60 percent to 99 percent of next year's salary seem to satisfy the vast majority of professional individuals.

The steps required to go from a realization of personal risk aversion and time preference to corporate counterparts and finally to a reward system for managers that will encourage them to make decisions consistent with corporate risk aversion and time preference remain a fascinating area of research.

4.2. Encoding of Uncertainty

When we begin the probabilistic phase of the decision analysis, we face the problem of encoding the uncertainty in each of the crucial state variables. We shall want to have the prior probability distributions assigned by the people within the enterprise who are most knowledgeable about each state variable. Thus the priors on engineering variables will typically be assigned by the engineering department; on marketing variables, by the marketing department, and so on. However, since we are in each case attempting to encode a probability distribution that reflects a state of mind and since most individuals have real difficulty in thinking about uncertainty, the method we use to extract the priors is extremely important. As people participate in the prior-gathering process, their attitudes are indicated successively by, "This is ridiculous," "It can't be done," "I have told you what you want to know but it doesn't mean anything," "Yes, it seems to reflect the way I feel," and "Why doesn't everybody do this?" In gathering the information we must be careful to overcome the defenses the individual develops as a result of being asked for estimates that are often a combination of targets, wishful thinking, and expectations. The biggest difficulty is in conveying to the man that you are interested in his state of knowledge and not in measuring him or setting a goal for him.

If the subject has some experience with probability, he often attempts to make all his priors look like normal distributions, a characteristic we may designate as "bellshaped" thinking. Although normal distributions are appropriate priors in some circumstances, we must avoid making them a foregone conclusion.

Experience has shown certain procedures to be effective in this almost psychoanalytic process of prior measurement. The first procedure is to make the measurement in a private interview to eliminate group pressure and to overcome the vague notions that most people exhibit about matters probabilistic. Sending around forms on which the subjects are supposed to draw their priors has been worse than useless, unless the subjects were already experienced in decision analysis.

Next we ask questions of the form, "What are the chances that x will exceed 10," because people seem much more comfortable in assigning probabilities to events than they are in sketching a density function. As these questions are asked, we skip around, asking the probability that x will be "greater than 50, less than 10, greater than 30," often asking the same question again later in the interview. The replies are recorded out of the view of the subject in order to frustrate any attempt at forced consistency on his part. As the interview proceeds, the subject often considers the questions with greater and greater care, so that his answers toward the end of the interview may represent his feelings much better than his initial answers. We can change the form of the questions by asking the subject to divide the domain of the random variable into n mutually exclusive regions with equal probability. (Of course, we would never put the question to him that way.) We can use the answers to all these questions to draw the complementary cumulative distribution for the variable, a form of representation that seems easiest to convey to people without formal probabilistic training.

The result of this interview is a prior that the subject is willing to live with, regardless of whether we are going to use it to govern a lottery on who buys coffee or on the disposal of his life savings. We can test it by comparing the prior with known probabilistic mechanisms; for example, if he says that a is the median of the distribution of x, then he should be indifferent about whether we pay him one hundred dollars if x exceeds a or if he can call the toss of a coin correctly. If he is not indifferent, then we must require him to change a until he is. The end result of such questions is to produce a prior that the subject is not tempted to change in any way, and we have thus achieved our final goal. The prior-gathering process is not cheap, but we perform it only on the crucial state variables.

In cases in which the interview procedure is not appropriate, the analyst can often obtain a satisfactory prior by drawing one himself and then letting the subject change it until the subject is satisfied. This technique may also be useful as an educational device in preparation for the interview.

If two or more variables are dependent, we must gather priors on conditional as well as marginal distributions. The procedure is generally the same but somewhat more involved. However, we have the benefit of being able to apply some checks on our results. Thus, if we have two dependent variables x and y, we can obtain the joint distribution by measuring the prior on x and the conditional on y, given x, or, alternatively, by measuring the prior on y and the conditional on x, given y. If we follow both routes, we have a consistency check on the joint distribution. Since the treating of joint variables is a source of expense, we should formulate the problem to avoid them whenever possible.

To illustrate the nature of prior gathering we present the example shown in Figure 3. The decision in a major problem was thought to depend primarily on the average lifetime of a new material. Since the material had never been made and test results would not be available until three years after the decision was required, it was necessary to encode the knowledge the company now had concerning the life of the material. This knowledge resided in three professional metallurgists who were experts in that field of technology. These men were interviewed separately according to the principles we have described. They produced the points labeled "Subjects 1, 2, and 3" in Figure 3. These results have several interesting features. We note, for example, that for $t = 17$ Subject

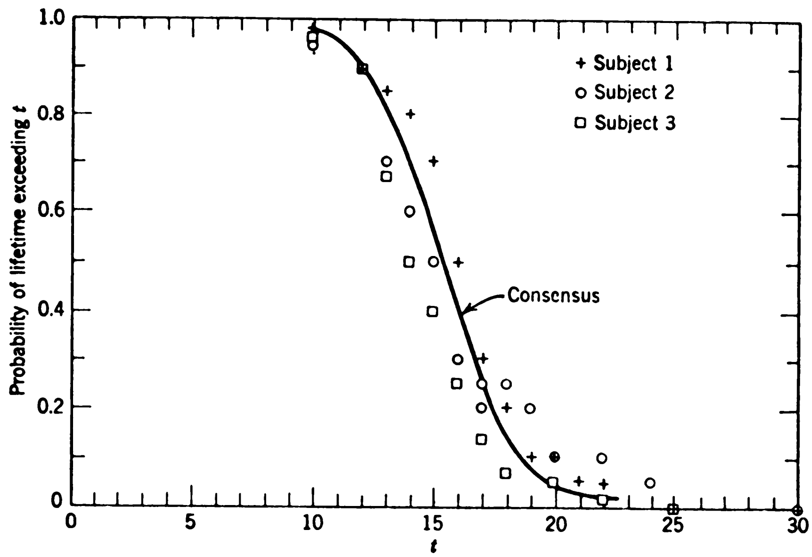

Figure 3. Priors on lifetime of material.

2 assigned probability 0.2 and 0.25 at various points in the interview. On the whole, however, the subjects were remarkably consistent in their assignments. We observe that Subject 3 was more pessimistic than Subject 1.

At the conclusion of the three interviews the three subjects were brought together and shown the results. At this point a vigorous discussion took place. Subjects 1 and 3, in particular, brought forth information of which the other two members of the group were unaware. As the result of this information exchange, the three group members drew the consensus curve—each subject said that this curve represented the state of information about the material life at the end of the meeting.

It has been suggested that the proper way to reconcile divergent priors is to assign weights to each, multiply, and add, but this experiment is convincing evidence that any such mechanistic procedure misses the point. Divergent priors are an excellent indicator of divergent states of information. The experience just described not only produced the company's present encoding of uncertainty about the lifetime of the material but at the same time encouraged the exchange of information within the group.

5. A DECISION-ANALYSIS EXAMPLE

To illustrate the flavor of application let us consider a recent decision analysis in the area of product introduction. Although the problem was really from another industry, let us suppose that it was concerned with the development and production of a new type of aircraft. There were two major alternatives: to develop and sell a new aircraft (A_2) or to continue manufacturing and selling

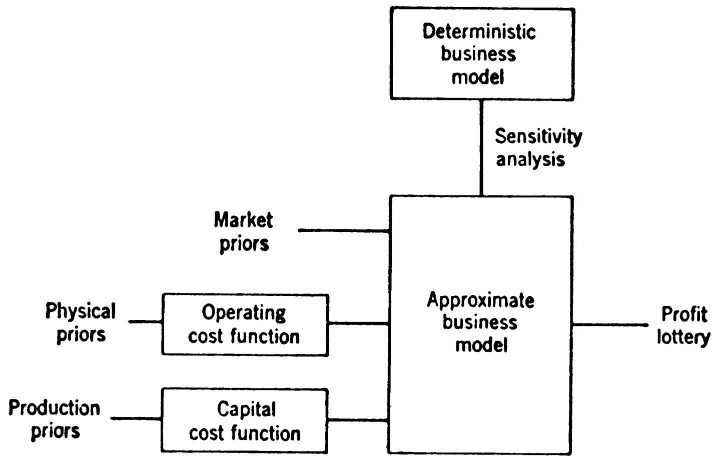

Figure 4. Decision analysis for new product introduction.

the present product (A_1). The decision was to be based on the present value of future expected profits at a discounting rate of 10 percent per year. Initially, the decision was supposed to rest on the lifetime of the material for which we obtained the priors in Figure 3; however, a complete decision analysis was desired. Since several hundred million dollars in present value of profit were at stake, the decision analysis was well justified.

The general scheme of the analysis appears in Figure 4. The first step was to construct a model of the business, a model that was primarily a model of the market. The profit associated with each alternative was described in terms of the price of the product, its operating capital costs, the behavior of its competitors, and the national characteristics of customers. The actual profit and discounted profit were computed over a 22-year time period. A suspicion grew that this model did not adequately capture the regional nature of demand. Consequently a new model was constructed that included the market characteristics, region by region and customer by customer. Moving to the more detailed basis affected the predictions so much that the additional refinement was clearly justified. Other attempts at refinement, however, did not affect the results sufficiently to justify a still more refined model. Now, the sensitivity analysis was performed to determine the crucial state variables, which turned out to be the operating cost, capital cost, and a few market parameters. Because of the complexity of the original business model, an approximate business model essentially quadratic in form was constructed to show how profit depended on these crucial state variables in the domain of interest. The coefficients of the approximate business model were established by runs on the complete business model.

The market priors were directly assigned with little trouble. However, because the operating and capital costs were the two most important variables in the problem, these priors were assigned according to a more detailed procedure. First, the operating cost was related to various physical features of the design by the engineering department. This relationship was called the operating-cost function. One of the many input physical variables was the average lifetime of the material whose priors appear in Figure 3. All but two of the 12 physical input variables were independent. The priors on the whole set of input variables were gathered and used with the operating-cost function in a Monte Carlo simulation that produced a prior for the operating cost of the product.

The capital-cost function was again developed by engineering but was much simpler in form. The input certainties were the production costs for various parts of the product. Again, a Monte Carlo analysis produced a prior on capital cost. .

Once we had established priors on all inputs to the approximate business model, we could determine the profit lottery for each alternative, in this case by using numerical analysis.

The present-value profit lotteries for the two alternatives looked very much like those shown in Figure 1. The new product alternative A_2 stochastically dominated the alternative A_1 of continuing to manufacture the present product. The result showed two interesting facets of the problem. First, it had been expected that the profit lottery for the new product alternative would be considerably broader than it was for the old product. The image was that of a profitable and risky new venture compared with a less profitable but less risky standard venture. In fact, the results showed that the uncertainties in profit were about the same for both alternatives, thus showing how initial concepts may be misleading.

The second interesting facet was that the average lifetime of the material whose priors appear in Figure 3 was actually of little consequence in the decision. It was true enough that profits were critically dependent on this lifetime if the design were fixed, but if the design were left flexible to accommodate to different average material lifetimes profits would be little affected. Furthermore, leaving the design flexible was not an expensive alternative; therefore another initial conception had to be modified.

However, the problem did not yield so easily. Figure 5 shows the present value of profits through each number of years t for each alternative. Note that if we ignore returns beyond year 7 the new product has a higher present value but that if we consider returns over the entire 22-year period the relationship reverses, as we have already noted. When management saw these results, they were considerably disturbed. The division in question had been under heavy pressure to show a profit in the near future—alternative A_2 would not meet that requirement. Thus the question of time preference that had been quickly passed off as one of present value at 10 percent per year became the central issue in the decision. The question was whether the division was interested in the quick kill or the long pull. At last report the division was still trying to convince the company to extend its profit horizon.

This problem clearly illustrates the use of decision analysis in clarifying the

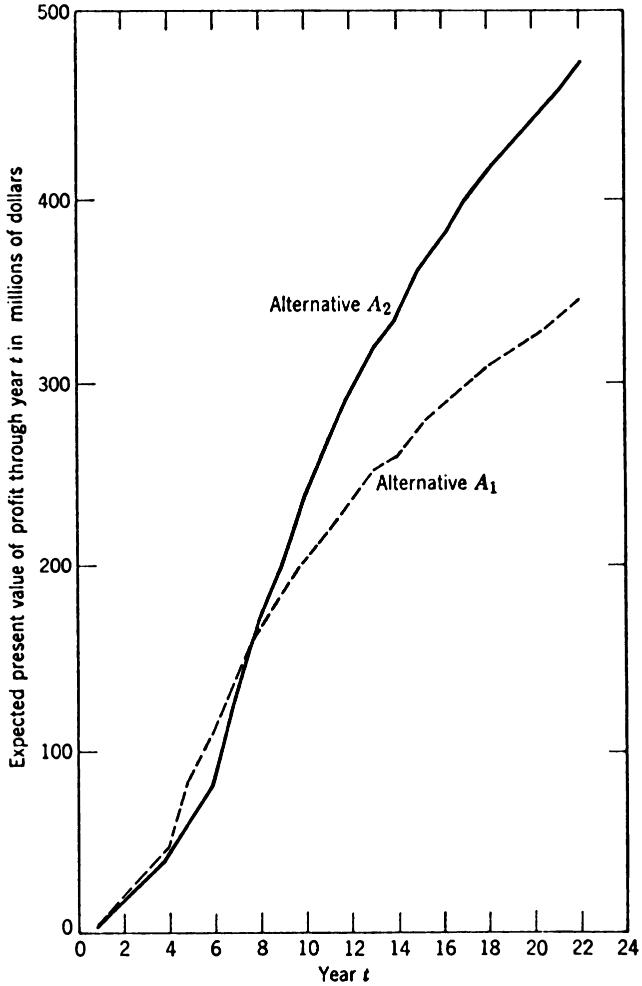

Figure 5. Expected present value of profit.

issues surrounding a decision. A decision that might have been made on the basis of a material lifetime was shown to depend more fundamentally on the question of time preference for profit. The nine man-months of effort devoted to this analysis were considered well spent by the company. The review committee for the decision commented, "We have never had such a realistic analysis of a new business venture before." The company is now interested in instituting decision-analysis procedures at several organizational levels.

6. CONCLUSION

Decision analysis offers operations research a second chance at top management. By foregoing statistical reproducibility we can begin to analyze the one-of-a-kind problems that managers have previously had to handle without assistance. Experience indicates that the higher up the chain of management we progress the more readily the concepts we have outlined are accepted. A typical reaction is, "I have been doing this all along, but now I see how to reduce my ideas to numbers."

Decision analysis is no more than a procedure for applying logic. The ultimate limitation to its applicability lies not in its ability to cope with problems but in man's desire to be logical.

ANALYSE DES DECISIONS: THEORIE APPLIQUEE DES DECISIONS

RÉSUMÉ

Au cours de ces dernières années, la théorie de décision a été de plus en plus acceptée en tant que cadre conceptuel pour la prise de décision. Cependant, cette théorie a surtout affecté les statisticiens plutôt que les personnes qui en ont le plus besoin: les responsables de décisions. Cette étude décrit un procédé qui permet de replacer des problèmes de décision réels dans la structure de la théorie de décision. Le procédé d'analyse de décision englobe chaque étape, du mesurage des choix de risques et des jugements portant sur des facteurs critiques par l'établissement de structures des facteurs relatifs à la technique, au marché, à la rivalité commerciale et à l'environnement, jusqu'au mesurage des préférences subjectives et de la valeur de la prédiction. L'analyse de décision met en perspective les nombreux instruments de simulation, d'analyse numérique, et de transformations de probabilités qui deviennent de plus en plus commodes depuis le développement des systèmes d'ordinateurs électroniques dont les différentes "stations" dépendent d'une "centrale" unique.

Le procédé est appliqué à un problème de décision réelle qui s'étend sur des dizaines d'annés et dont la valeur actuelle est de plusieurs centaines de millions de dollars. Cette étude analyse le problème de la détermination des dépenses consacrées à l'analyse de décisions. L'une des plus importantes propriétés de ce procédé tient au nombre des bénéfices auxiliaires créés au cours de l'élaboration de ce genre d'étude. L'expérience montre que ces bénéfices peuvant excéder en valeur le coût des dépenses consacrées a l'élaboration de la décision.

5

A TUTORIAL INTRODUCTION TO DECISION THEORY

D. Warner North
Decision Focus Incorporated

Reprinted with permission from IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, September 1968, pp. 200-210.

A Tutorial Introduction to Decision Theory

D. WARNER NORTH

Abstract—Decision theory provides a rational framework for choosing between alternative courses of action when the consequences resulting from this choice are imperfectly known. Two streams of thought serve as the foundations: utility theory and the inductive use of probability theory.

The intent of this paper is to provide a tutorial introduction to this increasingly important area of systems science. The foundations are developed on an axiomatic basis, and a simple example, the "anniversary problem," is used to illustrate decision theory. The concept of the value of information is developed and demonstrated. At times mathematical rigor has been subordinated to provide a clear and readily accessible exposition of the fundamental assumptions and concepts of decision theory. A sampling of the many elegant and rigorous treatments of decision theory is provided among the references.

INTRODUCTION

THE NECESSITY of making decisions in the face of uncertainty is an integral part of our lives. We must act without knowing the consequences that will result from the action. This uncomfortable situation is particularly acute for the systems engineer or manager who must make far-reaching decisions on complex issues in a rapidly changing technological environment. Uncertainty appears as the dominant consideration in many systems problems as well as in decisions that we face in our personal lives. To deal with these problems on a rational basis, we must develop a theoretical structure for decision making that includes uncertainty.

Confronting uncertainty is not easy. We naturally try to avoid it; sometimes we even pretend it does not exist. Our primitive ancestors sought to avoid it by consulting soothsayers and oracles who would "reveal" the uncertain future. The methods have changed: astrology and the reading of sheep entrails are somewhat out of fashion today, but predictions of the future still abound. Much current scientific effort goes into forecasting future economic and technological developments. If these predictions are assumed to be completely accurate, the uncertainty in many systems decisions is eliminated. The outcome resulting from a possible course of action may then be presumed to be known. Decision making becomes an optimization problem, and techniques such as mathematical programming may be used to obtain a solution. Such problems may be quite difficult to solve, but this difficulty should not obscure the fact that they represent the limiting case of perfect predictions. It is often tempting to assume perfect predictions, but in so doing we may be eliminating the most important features of the problem.[1] We should like to include in the analysis not just the predictions themselves, but also a measure of the confidence we have in these predictions. A formal theory of decision making must take uncertainty as its departure point and regard precise knowledge of outcomes as a limiting special case.

Before we begin our exposition, we will clarify our point of view. We shall take the engineering rather than the purely scientific viewpoint. We are not observing the way people make decisions; rather we are participants in the decision-making process. Our concern is in actually making a decision, i.e., making a choice between alternative ways of allocating resources. We must assume that at least two distinct alternatives exist (or else there is no element of choice and, consequently, no problem). Alternatives are distinct only if they result in different (uncertain) rewards or penalties for the decision maker; once the decision has been made and the uncertainty resolved, the resource allocation can be changed only by incurring some penalty.

What can we expect of a general theory for decision making under uncertainty? It should provide a framework in which all available information is used to deduce which of the decision alternatives is "best" according to the decision maker's preferences. But choosing an alternative that is consistent with these preferences and present knowledge does not guarantee that we will choose the alternative that by hindsight turns out to be most profitable.

We might distinguish between a good decision and a good *outcome*. We are all familiar with situations in which careful management and extensive planning produced poor results, while a disorganized and badly managed competitor achieved spectacular success. As an extreme example, place yourself in the position of the company president who has discovered that a valuable and trusted subordinate whose past judgment had proved unfailingly accurate actually based his decisions upon the advice of a gypsy fortune teller. Would you promote this man or fire him? The answer, of course, is to fire him and hire the gypsy as a consultant. The availability of such a clairvoyant to provide perfect information would make decision theory unnecessary. But we should not confuse the two. Decision theory is not a substitute for the fortune teller. It is rather a procedure that takes account of all available information to give us the best possible logical

Manuscript received May 8, 1968. An earlier version of this paper was presented at the IEEE Systems Science and Cybernetics Conference, Washington, D.C., October 17, 1966. This research was supported in part by the Graduate Cooperative Fellowship Program of the National Science Foundation at Stanford University, Stanford, Calif.

The author is with the Systems Sciences Area, Stanford Research Institute, Menlo Park, Calif. 94025.

[1] For further discussion of this point, see Howard [10] and Klein and Meckling [14].

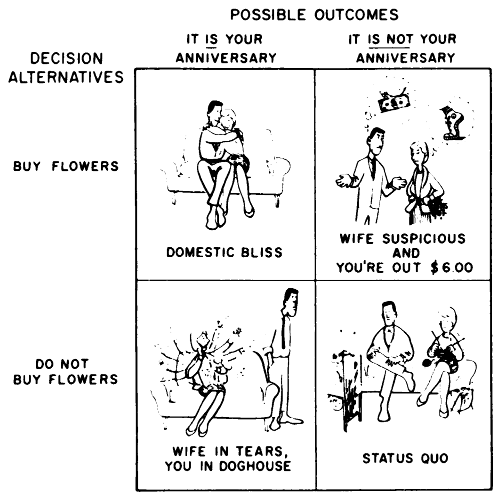

Fig. 1. Anniversary problem payoff matrix.

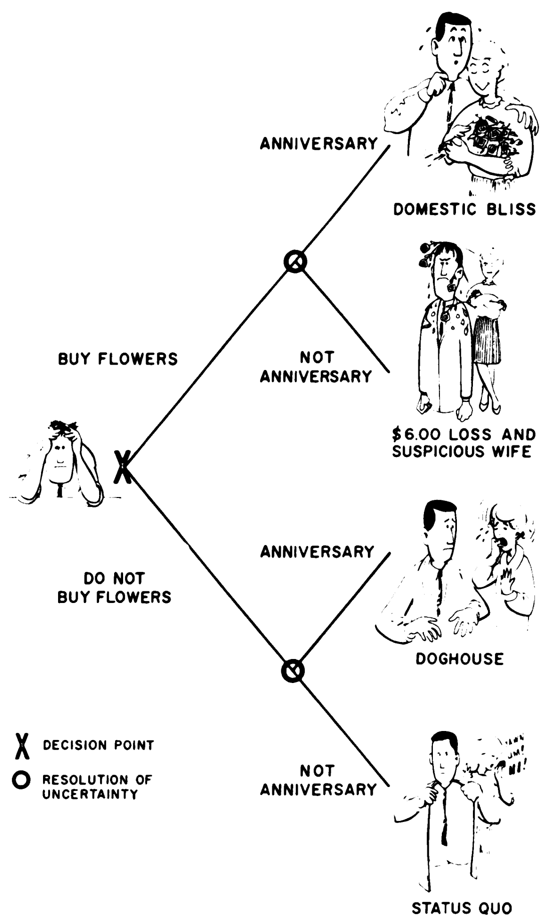

Fig. 2. Diagram of anniversary decision.

decision. It will minimize the consequences of getting an unfavorable outcome, but we cannot expect our theory to shield us from all "bad luck." The best protection we have against a bad outcome is a good decision.

Decision theory may be regarded as a formalization of common sense. Mathematics provides an unambiguous language in which a decision problem may be represented. There are two dimensions to this representation that will presently be described: value, by means of utility theory, and information, by means of probability theory. In this representation, the large and complex problems of systems analysis become conceptually equivalent to simple problems in our daily life that we solve by "common sense." We will use such a problem as an example.

You are driving home from work in the evening when you suddenly recall that your wedding anniversary comes about this time of year. In fact, it seems quite probable (but not certain) that it is today. You can still stop by the florist shop and buy a dozen roses for your wife, or you may go home empty-handed and hope the anniversary date lies somewhere in the future (Fig. 1). If you buy the roses and it is your anniversary, your wife is pleased at what a thoughtful husband you are and your household is the very epitome of domestic bliss. But if it is not your anniversary, you are poorer by the price of the roses and your wife may wonder whether you are trying to make amends for some transgression she does not know about. If you do not buy the roses, you will be in the clear if it is not your anniversary; but if it is, you may expect a temper tantrum from your wife and a two-week sentence to thedoghouse. What do you do?

We shall develop the general tools for solving decision problems and then return to this simple example. The reader might consider how he would solve this problem by "common sense" and then compare his reasoning with the formal solution which we shall develop later (Fig. 2).

The Machinery of Decision Making

Utility Theory

The first stage in setting up a structure for decision making is to assign numerical values to the possible outcomes. This task falls within the area covered by the modern theory of utility. There are a number of ways of developing the subject; the path we shall follow is that of Luce and Raiffa [16].[2]

The first and perhaps the biggest assumption to be made is that any two possible outcomes resulting from a decision can be compared. Given any two possible outcomes or prizes, you can say which you prefer. In some cases you might say that they were equally desirable or undesirable, and therefore you are indifferent. For example, you might prefer a week's vacation in Florida to a season ticket to the symphony. The point is not that the vacation costs more than the symphony tickets, but rather

[2] The classical reference on modern utility theory is von Neumann and Morgenstern [22]. A recent survey of the literature on utility theory has been made by Fishburn [5].

that you prefer the vacation. If you were offered the vacation or the symphony tickets on a nonnegotiable basis, you would choose the vacation.

A reasonable extension of the existence of your preference among outcomes is that the preference be transitive; if you prefer A to B and B to C, then it follows that you prefer A to C.[3]

The second assumption, originated by von Neumann and Morgenstern [22], forms the core of modern utility theory: you can assign preferences in the same manner to lotteries involving prizes as you can to the prizes themselves. Let us define what we mean by a lottery. Imagine a pointer that spins in the center of a circle divided into two regions, as shown in Fig. 3. If you spin the pointer and it lands in region I, you get prize A; if it lands in region II, you get prize B. We shall assume that the pointer is spun in such a way that when it stops, it is equally likely to be pointing in any given direction. The fraction of the circumference of the circle in region I will be denoted P, and that in region II as $1 - P$. Then from the assumption that all directions are equally likely, the probability that the lottery gives you prize A is P, and the probability that you get prize B is $1 - P$. We shall denote such a lottery as $(P,A;1 - P,B)$ and represent it by Fig. 4.

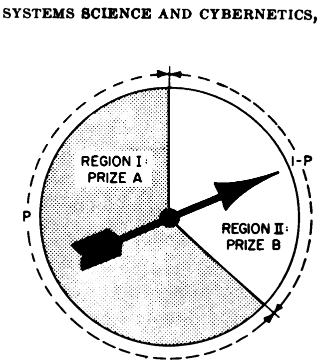

Fig. 3. A lottery.

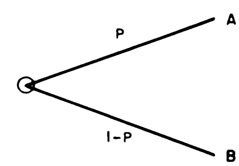

Fig. 4. Lottery diagram.

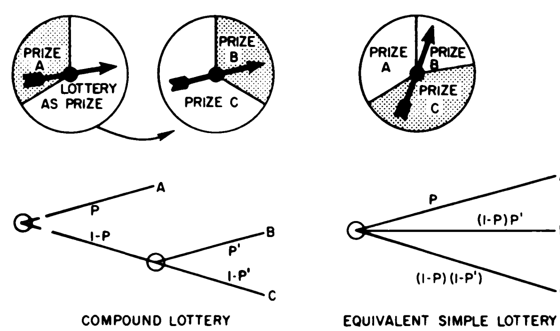

Fig. 5. "No fun in gambling."

Now suppose you are asked to state your preferences for prize A, prize B, and a lottery of the above type. Let us assume that you prefer prize A to prize B. Then it would seem natural for you to prefer prize A to the lottery, $(P,A;1 - P,B)$, between prize A and prize B, and to prefer this lottery between prize A and prize B to prize B for all probabilities P between 0 and 1. You would rather have the preferred prize A than the lottery, and you would rather have the lottery than the inferior prize B. Furthermore, it seems natural that, given a choice between two lotteries involving prizes A and B, you would choose the lottery with the higher probability of getting the preferred prize A, i.e., you prefer lottery $(P,A;1 - P,B)$ to $(P',A;1 - P',B)$ if and only if P is greater than P'.

The final assumptions for a theory of utility are not quite so natural and have been the subject of much discussion. Nonetheless, they seem to be the most reasonable basis for logical decision making. The third assumption is that there is no intrinsic reward in lotteries, that is, "no fun in gambling." Let us consider a compound lottery, a lottery in which at least one of the prizes is not an outcome but another lottery among outcomes. For example, consider the lottery $(P,A;1 - P,(P',B;1 - P',C))$. If the pointer of Fig. 3 lands in region I, you get prize A; if it lands in region II, you receive another lottery that has different prizes and perhaps a different division of the circle (Fig. 5). If you spin the second pointer you will receive prize B or prize C, depending on where this pointer lands. The assumption is that subdividing region II into two parts whose proportions correspond to the probabilities P' and $1 - P'$ of the second lottery creates an equivalent simple lottery in which all of the prizes are outcomes. According to this third assumption, you can decompose a compound lottery by multiplying the probability of the lottery prize in the first lottery by the probabilities of the individual prizes in the second lottery; you should be indifferent between $(P,A;1 - P,(P',B;1 - P',C))$ and $(P,A;P' - PP',B;1 - P - P' + PP',C)$. In other words, your preferences are not affected by the way in which the uncertainty is resolved—bit by bit, or all at once. There is no value in the lottery itself; it does not matter whether you spin the pointer once or twice.

Fourth, we make a continuity assumption. Consider three prizes, A, B, and C. You prefer A to C, and C to B (and, as we have pointed out, you will therefore prefer A to B). We shall assert that there must exist some probability P so that you are indifferent to receiving prize C or

[3] Suppose not: you would be at least as happy with C as with A. Then if a little man in a shabby overcoat came up and offered you C instead of A, you would presumably accept. Now you have C; and since you prefer B to C, you would presumably pay a sum of money to get B instead. Once you had B, you prefer A; so you would pay the man in the shabby overcoat some more money to get A. But now you are back where you started, with A, and the little man in the shabby overcoat walks away counting your money. Given that you accept a standard of value such as money, transitivity prevents you from becoming a "money pump."

the lottery $(P,A;1 - P,B)$ between A and B. C is called the certain equivalent of the lottery $(P,A;1 - P,B)$, and on the strength of our "no fun in gambling" assumption, we assume that interchanging C and the lottery $(P,A;1 - P,B)$ as prizes in some compound lottery does not change your evaluation of the latter lottery. We have *not* assumed that, given a lottery $(P,A;1 - P,B)$, there exists a Prize C intermediate in value between A and B so that you are indifferent between C and $(P,A;1 - P,B)$. Instead we have assumed the existence of the probability P. Given prize A preferred to prize C preferred to prize B, for some P between 0 and 1, there exists a lottery $(P,A;1 - P,B)$ such that you are indifferent between this lottery and Prize C. Let us regard the circle in Fig. 3 as a "pie" to be cut into two pieces, region I (obtain prize A) and region II (obtain prize B). The assumption is that the "pie" can be divided so that you are indifferent as to whether you receive the lottery or intermediate prize C.

Is this continuity assumption reasonable? Take the following extreme case:

A = receive \$1;
B = death;
C = receive nothing (status quo).

It seems obvious that most of us would agree A is preferred to C, and C is preferred to B; but is there a probability P such that we would risk death for the possibility of gaining \$1? Recall that the probability P can be arbitrarily close to 0 or 1. Obviously, we would not engage in such a lottery with, say, $P = 0.9$, i.e., a 1-in-10 chance of death. But suppose $P = 1 - 1 \times 10^{-50}$, i.e., the probability of death as opposed to \$1 is not 0.1 but 10^{-50}. The latter is considerably less than the probability of being struck on the head by a meteor in the course of going out to pick up a \$1 bill that someone has dropped on your doorstep. Most of us would not hesitate to pick up the bill. Even in this extreme case where death is a prize, we conclude the assumption is reasonable.

We can summarize the assumptions we have made into the following axioms.

A, B, C are prizes or outcomes resulting from a decision.

Notation:

$>$ means "is preferred to;"
$A > B$ means A is preferred to B;
$\sim$ means "is indifferent to;"
$A \sim B$ means the decision maker is indifferent between A and B.

Utility Axioms:

1) Preferences can be established between prizes and lotteries in an unambiguous fashion. These preferences are transitive, i.e.,

$$A > B,\quad B > C \quad \text{implies } A > C$$
$$A \sim B,\quad B \sim C \quad \text{implies } A \sim C.$$

2) If $A > B$, then $(P,A;1 - P,B) > (P',A;1 - P',B)$ if and only if $P > P'$.

3) $(P,A;1 - P,(P',B;1 - P',C)) \sim (P,A;P' - PP',B;1 - P - P' + PP',C)$, i.e., there is "no fun in gambling."

4) If $A > C > B$, there exists a P with $0 < P < 1$ so that

$$C \sim (P,A;1 - P,B)$$

i.e., it makes no difference to the decision maker whether C or the lottery $(P,A;1 - P,B)$ is offered to him as a prize.

Under these assumptions, there is a concise mathematical representation possible for preferences: a utility function $u(\ )$ that assigns a number to each lottery or prize. This utility function has the following properties:

$$u(A) > u(B) \text{ if and only if } A > B \tag{1}$$

$$\text{if } C \sim (P,A;1 - P,B), \quad \text{then } u(C) = P \cdot u(A) + (1 - P) \cdot u(B) \tag{2}$$

i.e., the utility of a lottery is the mathematical expectation of the utility of the prizes. It is this "expected value" property that makes a utility function useful because it allows complicated lotteries to be evaluated quite easily.

It is important to realize that all the utility function does is provide a means of consistently describing the decision maker's preferences through a scale of real numbers, providing these preferences are consistent with the previously mentioned assumptions 1) through 4). The utility function is no more than a means to logical deduction based on given preferences. The preferences come first and the utility function is only a convenient means of describing them. We can apply the utility concept to almost any sort of prizes or outcomes, from battlefield casualties or achievements in space to preferences for Wheaties or Post Toasties. All that is necessary is that the decision maker have unambiguous preferences and be willing to accept the basic assumptions.

In many practical situations, however, outcomes are in terms of dollars and cents. What does the utility concept mean here? For an example, let us suppose you were offered the following lottery: a coin will be flipped, and if you guess the outcome correctly, you gain \$100. If you guess incorrectly, you get nothing. We shall assume you feel that the coin has an equal probability of coming up heads or tails; it corresponds to the "lottery" which we have defined in terms of a pointer with $P = 1/2$. How much would you pay for such a lottery? A common answer to this academic question is "up to \$50," the average or expected value of the outcomes. When real money is involved, however, the same people tend to bid considerably lower; the average bid is about \$20.[4] A group of Stanford University graduate students was actually confronted with a \$100 pile of bills and a 1964 silver quarter to flip. The average of the sealed bids for this game was slightly under \$20, and only 4 out of 46 ventured to bid as high as \$40. (The high bidder, at \$45.61, lost and the proceeds were used for a class party.) These results are quite typical; in fact, professional engineers and managers are, if any-

[4] Based on unpublished data obtained by Prof. R. A. Howard of Stanford University, Stanford, Calif.

thing, more conservative in their bids than the less affluent students.

The lesson to be learned here is that, by and large, most people seem to be averse to risk in gambles involving what is to them substantial loss. They are willing to equate the value of a lottery to a sure payoff or certain equivalent substantially less than the expected value of the outcomes. Similarly, most of us are willing to pay more than the expected loss to get out of an unfavorable lottery. This fact forms the basis of the insurance industry.

If you are very wealthy and you are confronted with a small lottery, you might well be indifferent to the risk. An unfavorable outcome would not deplete your resources, and you might reason that you will make up your losses in future lotteries; the "law of averages" will come to your rescue. You then evaluate the lottery at the expected value of the prizes. For example, the (1/2, \$0; 1/2, \$100) lottery would be worth 1/2(\$0) + 1/2(\$100) = \$50 to you. Your utility function is then a straight line, and we say you are an "expected value" decision maker. For lotteries involving small prizes, most individuals and corporations are expected value decision makers. We might regard this as a consequence to the fact that any arbitrary utility curve for money looks like a straight line if we look at a small enough section of it. Only when the prizes are substantial in relation to our resources does the curvature become evident. Then an unfavorable outcome really hurts. For these lotteries most of us become quite risk averse, and expected value decision making does not accurately reflect our true preferences.

Let us now describe one way you might construct your own utility curve for money, say, in the amounts of \$0 to \$100, in addition to your present assets. The utility function is arbitrary as to choice of zero point and of scale factor; changing these factors does not lead to a change in the evaluation of lotteries using properties (1) and (2). Therefore, we can take the utility of \$0 as 0 and the utility of \$100 as 1. Now determine the minimum amount you would accept in place of the lottery of flipping a coin to determine whether you receive \$0 or \$100. Let us say your answer is \$27. Now determine the certain equivalent of the lotteries (1/2, \$0; 1/2, \$27), and (1/2, \$27; 1/2, \$100), and so forth. We might arrive at a curve like that shown in Fig. 6.

We have simply used the expected value property (2) to construct a utility curve. This same curve, however, allows us to use the same expected utility theorem to evaluate new lotteries; for example, (1/2, \$30; 1/2 \$80). From Fig. 6, $u(\$30) = 0.54$, $u(\$80) = 0.91$, and therefore $1/2\ u(\$30) + 1/2\ u(\$80) = u(x) \rightarrow x = \$49$. If you are going to be consistent with the preferences you expressed in developing the utility curve, you will be indifferent between \$49 and this lottery. Moreover, this amount could have been determined from your utility curve by a subordinate or perhaps a computer program. You could send your agent to make decisions on lotteries by using your utility curve, and he would make them to reflect *your* preference for amounts in the range \$0 to \$100.

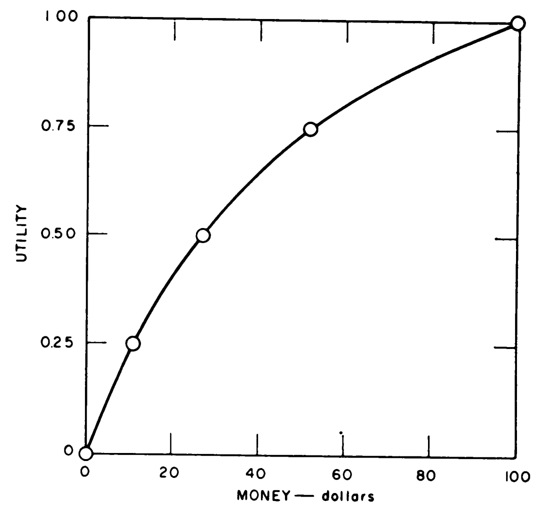

Fig. 6. Utility curve for money: \$0 to \$100.

Even without such a monetary representation, we can always construct a utility function on a finite set of outcomes by using the expected value property (2). Let us choose two outcomes, one of which is preferred to the other. If we set the utilities arbitrarily at 1 for the preferred outcome and 0 for the other, we can use the expected value property (2) of the utility function to determine the utility of the other prizes. This procedure will always work so long as our preferences obey the axioms, but it may be unwieldy in practice because we are asking the decision maker to assess simultaneously his values in the absence of uncertainty and his preference among risks. The value of some outcome is accessible only by reference to a lottery involving the two "reference" outcomes. For example, the reference outcomes in the anniversary problem might be "domestic bliss" = 1 and "doghouse" = 0. We could then determine the utility of "status quo" as 0.91 since the husband is indifferent between the outcome "status quo" and a lottery in which the chances are 10 to 1 of "domestic bliss" as opposed to the "doghouse." Similarly, we might discover that a utility of 0.667 should be assigned to "suspicious wife and \$6 wasted on roses," since our friend is indifferent between this eventuality and a lottery in which the probabilities are 0.333 of "doghouse" and 0.667 of "domestic bliss." Of course, to be consistent with the axioms, our friend must be indifferent between "suspicious wife, etc.," and a 0.73 probability of "status quo" and a 0.27 probability of "doghouse." If the example included additional outcomes as well, he might find it quite difficult to express his preferences among the lotteries in a manner consistent with the axioms. It may be advisable to proceed in two stages; first, a numerical determination of value in a risk-free situation, and then an adjustment to this scale to include preference toward risk.

Equivalent to our first assumption, the existence of transitive preferences, is the existence of some scale of value by which outcomes may be ranked; A is preferred to B if and only if A is higher in value than B. The numerical structure we give to this value is not important since a monotonic transformation to a new scale preserves the ranking of outcomes that corresponds to the original preferences. No matter what scale of value we use, we can construct a utility function on it by using the expected value property (2), so long as our four assumptions hold. We may as well use a standard of value that is reasonably intuitive, and in most situations money is a convenient standard of economic value. We can then find a monetary equivalent for each outcome by determining the point at which the decision maker is indifferent between receiving the outcome and receiving (or paying out) this amount of money. In addition to conceptual simplicity, this procedure makes it easy to evaluate new outcomes by providing an intuitive scale of values. Such a scale will become necessary later on if we are to consider the value of resolving uncertainty.

We will return to the anniversary decision and demonstrate how this two-step value determination procedure may be applied. But first let us describe how we shall quantify uncertainty.

The Inductive Use of Probability Theory

We now wish to leave the problem of the evaluation of outcomes resulting from a decision and turn our attention to a means of encoding the information we have as to which outcome is likely to occur. Let us look at the limiting case where a decision results in a certain outcome. We might represent an outcome, or an event, which is certain to occur by 1, and an event which cannot occur by 0. A certain event, together with another certain event, is certain to occur; but a certain event, together with an impossible event, is certain not to occur. Most engineers would recognize the aforementioned as simple Boolean equations: $1 \cdot 1 = 1$, $1 \cdot 0 = 0$. Boolean algebra allows us to make complex calculations with statements that may take on only the logical values "true" and "false." The whole field of digital computers is, of course, based on this branch of mathematics.

But how do we handle the logical "maybe?" Take the statement, "It will rain this afternoon." We cannot now assign this statement a logical value of true or false, but we certainly have some feelings on the matter, and we may even have to make a decision based on the truth of the statement, such as whether to go to the beach. Ideally, we would like to generalize the inductive logic of Boolean algebra to include uncertainty. We would like to be able to assign to a statement or an event a value that is a measure of its uncertainty. This value would lie in the range from 0 to 1. A value of 1 indicates that the statement is true or that the event is certain to occur; a value of 0 indicates that the statement is false or that the event cannot occur. We might add two obvious assumptions. We want the value assignments to be unambiguous, and we want the value assignments to be independent of any assumptions that have not been explicitly introduced. In particular, the value of the statement should depend on its content, not on the way it is presented. For example, "It will rain this morning or it will rain this afternoon," should have the same value as "It will rain today."

These assumptions are equivalent to the assertion that there is a function P that gives values between 0 and 1 to events ("the statement is true" is an event) and that obeys the following probability axioms.[5]

Let E and F be events or outcomes that could result from a decision:

1) $P(E) \geq 0$ for any event E;
2) $P(E) = 1$, if E is certain to occur;
3) $P(E \text{ or } F) = P(E) + P(F)$ if E and F are mutually exclusive events (i.e., only one of them can occur).

E *or* F means the event that either E or F occurs. We are in luck. Our axioms are identical to the axioms that form the modern basis of the theory of probability. Thus we may use the whole machinery of probability theory for inductive reasoning.

Where do we obtain the values $P(E)$ that we will assign to the uncertainty of the event E? We get them from our own minds. They reflect our best judgment on the basis of all the information that is presently available to us. The use of probability theory as a tool of inductive reasoning goes back to the beginnings of probability theory. In Napoleon's time, Laplace wrote the following as a part of his introduction to *A Philosophical Essay on Probabilities* ([15], p. 1):

> Strictly speaking it may even be said that nearly all our knowledge is problematical; and in the small numbers of things which we are able to know with certainty, even in the mathematical sciences themselves, the principal means for ascertaining truth—induction and analogy—are themselves based on probabilities . . ..

Unfortunately, in the years following Laplace, his writings were misinterpreted and fell into disfavor. A definition of probability based on frequency came into vogue, and the pendulum is only now beginning to swing back. A great many modern probabilists look on the probability assigned to an event as the limiting fraction of the number of times an event occurred in a large number of independent repeated trials. We shall not enter into a discussion of the general merits of this viewpoint on probability theory. Suffice it to say that the situation is a rare one in which you can observe a great many independent identical trials in order to assign a probability. In fact, in decision theory we are often interested in events that will occur just once. For us, a probability assessment is made on the basis of a state of mind; it is not a property of physical objects to be measured like length, weight, or temperature. When we assign the probability of 0.5 to a coin coming up heads, or equal probabilities to all possible orientations of a pointer, we may be reasoning on the basis of the symmetry of the

[5] Axioms 1) and 2) are obvious, and 3) results from the assumption of invariance to the form of data presentation (the last sentence in the preceding paragraph). Formal developments may be found in Cox [3], Jaynes [12], or Jeffreys [13]. A joint axiomatization of both probability and utility theory has been developed by Savage [20].

physical object. There is no reason to suppose that one side of the coin will be favored over the other. But the physical symmetry of the coin does not lead immediately to a probability assignment of 0.5 for heads. For example, consider a coin that is placed on a drum head. The drum head is struck, and the coin bounces into the air. Will it land heads up half of the time? We might expect that the probability of heads would depend on which side of the coin was up initially, how hard the drum was hit, and so forth. The probability of heads is not a physical parameter of the coin; we have to specify the flipping system as well. But if we knew exactly how the coin were to be flipped, we could calculate from the laws of mechanics whether it would land heads or tails. Probability enters as a means of describing our feelings about the likelihood of heads when our knowledge of the flipping system is not exact. We must conclude that the probability assignment depends on our present state of knowledge.

The most important consequence of this assertion is that probabilities are subject to change as our information improves. In fact, it even makes sense to talk about probabilities of probabilities. A few years ago we might have assigned the value 0.5 to the probability that the surface of the moon is covered by a thick layer of dust. At the time, we might have said, "We are 90 percent certain that our probability assignment after the first successful Surveyor probe will be less than 0.01 or greater than 0.99. We expect that our uncertainty about the composition of the moon's surface will be largely resolved."

Let us conclude our discussion of probability theory with an example that will introduce the means by which probability distributions are modified to include new information: Bayes' rule. We shall also introduce a useful notation. We have stressed that all of our probability assignments are going to reflect a state of information in the mind of the decision maker, and our notation shall indicate this state of information explicitly.

Let A be an event, and let x be a quantity about which we are uncertain; e.g., x is a random variable. The values that x may assume may be discrete (i.e., heads or tails) or continuous (i.e., the time an electronic component will run before it fails). We shall denote by $\{A|S\}$ the probability assigned to the event A on the basis of a state of information S, and by $\{x|S\}$ the probability that the random variable assumes the value x, i.e., the probability mass function for a discrete random variable or the probability density function for a continuous random variable, given a state of information S. If there is confusion between the random variable and its value, we shall write $\{x = x_0|S\}$, where x denotes the random variable and x_0 the value. We shall assume the random variable takes on some value, so the probabilities must sum to 1:

$$\int_x \{x|S\} = 1. \tag{3}$$

$\int$ is a generalized summation operator representing summation over all discrete values or integration over all continuous values of the random variable. The expected value, or the average of the random variable over its probability distribution, is

$$\langle x|S\rangle = \int_x x\{x|S\}. \tag{4}$$

One special state of information will be used over and over again, so we shall need a special name for it. This is the information that we now possess on the basis of our prior knowledge and experience, before we have done any special experimenting or sampling to reduce our uncertainty. The probability distribution that we assign to values of an uncertain quantity on the basis of this prior state of information (denoted $\mathcal{E}$) will be referred to as the "prior distribution" or simply the "prior."

Now let us consider a problem. Most of us take as axiomatic the assignment of 0.5 to the probability of heads on the flip of a coin. Suppose we flip thumbtacks. If the thumbtack lands with the head up and point down, we shall denote the outcome of the flip as "heads." If it lands with the head down and the point up, we shall denote the outcome as "tails." The question which we must answer is, "What is p, the probability of heads in flipping a thumbtack?" We will assume that both thumbtack and means of flipping are sufficiently standardized so that we may expect that all flips are independent and have the same probability for coming up heads. (Formally, the flips are Bernoulli trials.) Then the long-run fraction of heads may be expected to approach p, a well-defined number that at the moment we do not know.

Let us assign a probability distribution to this uncertain parameter p. We are all familiar with thumbtacks; we have no doubt dropped a few on the floor. Perhaps we have some experience with spilled carpet tacks, or coin flipping, or the physics of falling bodies that we believe is relevant. We want to encode all of this prior information into the form of a probability distribution on p.

This task is accomplished by using the cumulative distribution function, $\{p \le p_0|\mathcal{E}\}$, the probability that the parameter p will be less than or equal to some specific value of the parameter p_0. It may be convenient to use the complementary cumulative

$$\{p > p_0|\mathcal{E}\} = 1 - \{p \le p_0|\mathcal{E}\}$$

and ask questions such as, "What is the probability that p is greater than $p_0 = 0.5$?"

To make the situation easier to visualize, let us introduce Sam, the neighborhood bookie. We shall suppose that we are forced to do business with Sam. For some value p_0 between 0 and 1, Sam offers us two packages:

Package 1: If measurement of the long run fraction of heads p shows that the quantity is less than or equal to p_0, then Sam pays us **$1**. If $p > p_0$, then we pay Sam **$1**.

Package 2: We divide a circle into two regions (as shown in Fig. 3). Region I is defined by a fraction P of the circumference of the circle, and the remainder of the circle constitutes region II. Now a pointer is spun in such a way that when it stops, it is equally likely to be pointing in any given direction. If the pointer stops in region I, Sam pays us $1; if it lands in region II, we pay Sam $1.

Sam lets us choose the fraction P in Package 2, but then he chooses which package we are to receive. Depending on the value of p_0, these packages may be more or less attractive to us, but it is the relative rather than the absolute value of the two packages that is of interest. If we set P to be large, we might expect that Sam will choose package 1, whereas if P is small enough, Sam will certainly choose package 2. Sam wishes (just as we do) to have the package with the higher probability of winning $1. (Recall this is our second utility axiom.) We shall assume Sam has the same information about thumbtacks that we do, so his probability assignments will be the same as ours. The assumption [utility axiom 4)] is that given p_0, we can find a P such that Packages 1 and 2 represent equivalent lotteries, so $P = \{p \leq p_0|\mathcal{E}\}$.[6] The approach is similar to the well-known method of dividing an extra dessert between two small boys: let one divide and the other choose. The first is motivated to make the division as even as possible so that he will be indifferent as to which half he receives.

Suppose Sam starts at a value $p_0 = 0.5$. We might reason that since nails always fall on the side (heads), and a thumbtack is intermediate between a coin and a nail heads is the more likely orientation; but we are not too sure; we have seen a lot of thumbtacks come up tails. After some thought, we decide that we are indifferent about which package we get if the fraction P is 0.3, so $\{p \leq 0.5|\mathcal{E}\} = 0.30$.

Sam takes other values besides 0.5, skipping around in a random fashion, i.e., 0.3, 0.9, 0.1, 0.45, 0.8, 0.6, etc. The curve that results from the interrogation might look like that shown in Fig. 7. By his method of randomly skipping around, Sam has eliminated any bias in our true feelings that resulted from an unconscious desire to give answers consistent with previous points. In this fashion, Sam has helped us to establish our prior distribution on the parameter p. We may derive a probability density function by taking the derivative of the cumulative distribution function (Fig. 8): $\{p|\mathcal{E}\} = (d/dp_0)\{p \leq p_0|\mathcal{E}\}$.

Now supposing we are allowed to flip the thumbtack 20 times and we obtain 5 heads and 15 tails. How do we take account of this new data in assigning a probability distribution based on the new state of information, which we denote as $\mathcal{E}, E$: our prior experience $\mathcal{E}$ plus E, the 20-flip experiment? We will use one of the oldest (1763) results of probability theory, Bayes' rule. Consider the prior probability that p will take on a specific value and the 20-flip experiment E will have a certain specific outcome (for example, $p = 0.43$; E = 5 heads, 15 tails). Now we can write this joint probability in two ways:

$$\{p,E|\mathcal{E}\} = \{p|E,\mathcal{E}\}\{E|\mathcal{E}\} \tag{5}$$

[6] We have equated the subjective probability that summarized our information about thumbtacks to the more intuitive notion of probability based on symmetry (in Package 2). Such a two-step approach to probability theory has been discussed theoretically by Anscombe and Aumann [1].

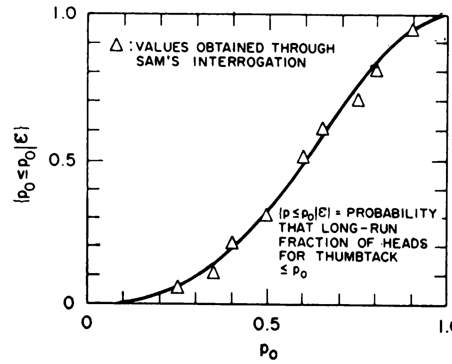

Fig. 7. Cumulative distribution function for thumbtack flipping.

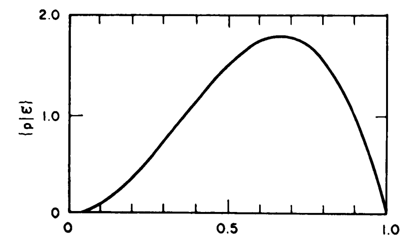

Fig. 8. Prior probability density function.

i.e., as the product of the probability we assign to the experimental outcome E times the probability we would assign to the value of p after we knew the experimental outcome E in addition to our prior information; or

$$\{p,E|\mathcal{E}\} = \{E|p,\mathcal{E}\}\{p|\mathcal{E}\} \tag{6}$$

i.e., the product of the probability of that experimental outcome if we knew that p were the probability of getting heads times our prior probability assessment that p actually takes on that value.

We assumed that probabilities were unambiguous, so we equate these two expressions. Providing $\{E|\mathcal{E}\} \neq 0$, i.e., the experimental outcome is not impossible, we obtain the posterior (after the experiment) probability distribution on p

$$\{p|E,\mathcal{E}\} = \frac{\{E|p,\mathcal{E}\}\{p|\mathcal{E}\}}{\{E|\mathcal{E}\}}. \tag{7}$$

This expression is the well-known Bayes' rule.

$\{E|\mathcal{E}\}$ is the "pre-posterior" probability of the outcome E. It does not depend on p, so it becomes a normalizing factor for the posterior probability distribution. $\{E|p,\mathcal{E}\}$ is the probability of the outcome E if we knew the value p for the probability of heads. This probability is a function of p, usually referred to as the "likelihood function." We notice since p must take on some value, the expectation of the likelihood function over the values of p gives the pre-posterior probability of the experimental outcome:

$$\{E|\mathcal{E}\} = \int_p \{E|p,\mathcal{E}\}\{p|\mathcal{E}\}. \tag{8}$$

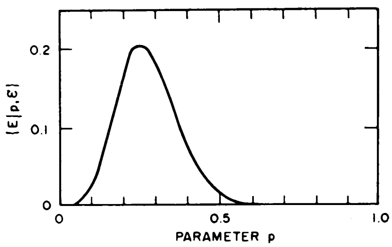

Fig. 9. Likelihood function for 5 heads in 20 trials.

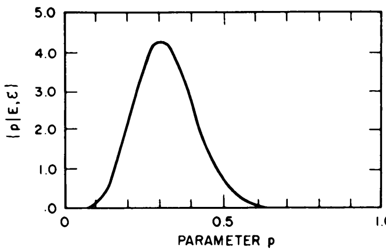

Fig. 10. Posterior probability density function.

For the specific case we are treating, the likelihood function is the familiar result from elementary probability theory for r successes in n Bernoulli trials when the probability of a success is p:

$$\{E|p,\mathcal{E}\} = \frac{n!}{r!(n-r)!} p^r(1-p)^{n-r}. \qquad (9)$$

This function is graphed for $r = 5$ heads in $n = 20$ trials in Fig. 9. Multiplying it by the prior $\{p|\mathcal{E}\}$ (Fig. 8) and normalizing by dividing by $\{E|\mathcal{E}\}$ gives us the posterior distribution $\{p|E,\mathcal{E}\}$ (Fig. 10). In this way, Bayes' rule gives us a general means of revising our probability assessments to take account of new information.[7]

Solution of Decision Problems

Now that we have the proper tools, utility theory and probability theory, we return to the anniversary decision problem. We ask the husband, our decision maker, to assign monetary values to the four possible outcomes. He does so as follows:

Domestic bliss	(flowers + anniversary):	\$100
Doghouse	(no flowers, anniversary):	\$ 0
Status quo	(no flowers, no anniversary):	\$ 80
Suspicious wife	(flowers, no anniversary):	\$ 42.

(For example, he is indifferent between "status quo" and "doghouse" provided in the latter case he receives \$80.) His preference for risk is reflected by the utility function of Fig. 6, and he decides that a probability assessment of 0.2 sums up his uncertainty about the possibility of today being his anniversary: the odds are 4 to 1 that it is not his anniversary. Now let us look at the two lotteries that represent his decision alternatives. If he buys the flowers, he has a 0.2 probability of "domestic bliss" and an 0.8 probability of "suspicious wife." The expected utility of the lottery is $0.2(1.0) + 0.8(0.667) = 0.734 = u(\$50)$. On the other hand, if he does not buy the flowers, he has an 0.8 chance of "status quo" and a 0.2 chance of "doghouse." The expected utility of this alternative is $0.8(0.91) + 0.2(0) = 0.728 = u(\$49)$. The first alternative has a slightly higher value to him so he should buy the flowers. On the basis of his values, his risk preference, and his judgment about the uncertainty, buying the flowers is his best alternative. If he were an expected value decision maker, the first lottery would be worth $0.2(\$100) + 0.8(\$42) = \$53.60$ and the second $0.2(0) + 0.8(\$80) = \$64$. In this case he should not buy the flowers.

The foregoing example is, of course, very trivial, but conceptually any decision problem is exactly the same. There is only one additional feature that we may typically expect: in general, decision problems may involve a sequence of decisions. First, a decision is made and then an uncertain outcome is observed; after which another decision is made, and an outcome observed, etc. For example, the decision to develop a new product might go as follows. A decision is made as to whether or not a product should be developed. If the decision is affirmative, an uncertain research and development cost will be incurred. At this point, a decision is made as to whether to go into production. The production cost is uncertain. After the production cost is known, a sale price is set. Finally, the uncertain sales volume determines the profit or loss on the product.

We can handle this problem in the same way as the anniversary problem: assign values to the final outcomes, and probabilities to the various uncertain outcomes that will result from the adoption of a decision alternative. We can represent the problem as a decision tree (Fig. 11), and the solution is conceptually easy. Start at the final outcome, sales volume (the ends of the tree). Go in to the first decision, the sales price (the last to be made chronologically). Compute the utility of the decision alternatives, and choose the one with the highest value. This value becomes the utility of the chance outcome leading to that decision (e.g., production cost). The corresponding certain equivalent in dollars reflects the expected utility of reaching that point in the tree. In this fashion, we work backwards to the start of the tree, finding the best decision alternatives and their values at each step.

Many decision problems encountered in actual practice are extremely complex, and a decision tree approach may not always be appropriate. If all quantities concerned in the problem were considered uncertain (with prior distributions), the problem might be computationally intractable. It is often advisable to solve the model deterministically as a first approximation. We approximate all uncertain quantities with a single best estimate and then examine the decision; i.e., if research and development costs, production costs, and sales volume took the

[7] For certain sampling processes having special statistical properties, assumption of a prior probability distribution from a particular family of functions leads to a simple form for Bayes' rule. An extensive development of this idea of "conjugate distributions" has been accomplished by Raiffa and Schlaifer [19].

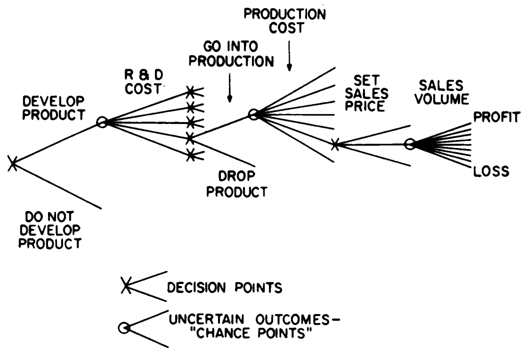

Fig. 11. Product development decision tree.

values we consider most likely, would it then be advisable to develop the product? This deterministic phase will usually give us some insight into the decision. Moreover, we can perform a sensitivity analysis by varying quantities that we believe are uncertain to determine how they affect the decision. The decision may be quite insensitive to some quantities, and these quantities may be treated as certain (uncertainty is neglected if it appears not to affect the decision). On the other hand, if a variation that lies within the range of uncertainty of a factor causes a major shift in the decision (i.e., from "develop the product" to "do not develop the product"), we shall certainly wish to encode our feelings about the uncertainty of that quantity by a prior distribution.[8]

The Value of Resolving Uncertainties

There is a class of alternatives usually available to the decision maker that we have not yet mentioned: activities that allow him to gather more information to diminish the uncertainties before he makes the decision. We have already seen how new information may be incorporated into probability assessments through Bayes' rule, and we noted that we can assign a probability distribution to the results of the information gathering by means of the pre-posterior probability distribution. Typical information-gathering activities might include market surveys, pilot studies, prototype construction, test marketing, or consulting with experts. These activities invariably cost the decision maker time and resources; he must pay a price for resolving uncertainty.

Let us return to the husband with the anniversary problem. Suppose he has the option of calling his secretary. If it is his anniversary, his secretary will certainly tell him. But if it is not, she may decide to play a trick and tell him that today is his anniversary. He assigns probability 0.5 to such practical joking. In any event, the secretary will spread the word around the office and our friend will get some good natured heckling, which he views as having a value of minus **$10**.

[8] The decision analysis procedure has been described in detail by Howard [8].

How will the secretary's information change his assessment of the probability that today is his anniversary? If she says, "No, it is not your anniversary," he may be sure that it is not; but if she says "Yes, it is," she could be joking. We can compute the new assessment of the probability from Bayes' rule. This new probability is equal to the probability 0.2 that she says yes and it really is his anniversary, divided by his prior estimate, $0.2 + 0.5 \times 0.8 = 0.6$, that she will say yes regardless of the date of his anniversary. Hence the probability assignment revised to include the secretary's yes answer is 0.333.

What is the value of this new alternative to our friend? If his secretary says no (probability 0.4), he may return home empty-handed and be assured of "status quo." On the other hand, if she says yes (probability 0.6), he will buy the flowers. In either case, he has incurred a cost of **$10** which must be subtracted from the values of the outcomes. Calling the secretary then has a utility of

$$0.4\,u(\$70) + 0.6\,[0.333\,u(\$90) + 0.667\,u(\$32)] = 0.344 + 0.416 = 0.760 = u(\$53.50).$$

Since this value of **$53.50** exceeds the value of **$50** for his previous best alternative (buy flowers), our friend should call his secretary. If the husband were an expected value decision maker, the alternative of calling the secretary would have a value of

$$0.4\,(\$70) + 0.6\,[0.333\,(\$90) + 0.667\,(\$32)] = \$58.80$$

which is less than the value of **$64** for the "do not buy flowers" alternative; in this case our friend should not call his secretary. It is evident that in this example preference toward risk is very important in determining the decision maker's best course of action.

In the complex decision problems normally encountered in practice, there are usually several alternative options available for diminishing the uncertainty associated with the unknown factors. In theory, the expected gain for each type of sampling could be computed and compared with the cost of sampling as we have just done in the simple anniversary example. But these calculations can be quite involved as a rule, and there may be a great many alternative ways of gathering information. Often the relevant questions are, first, "Should we sample at all?" and then, "What kind of sampling is best for us?"

It is often useful to look at the limiting case of complete resolution of uncertainty, which we call perfect information. We can imagine that a gypsy fortune teller who always makes correct predictions is, in fact, available to us. The value of perfect information is the amount that we are willing to pay her to tell us exactly what the uncertain quantity will turn out to be. Note that her answer may be of little value to us—we may be planning to take the best decision alternative already. On the other hand, her perfect information may be quite valuable; it may allow us to avoid an unfavorable outcome. We are going to have to pay her before we hear her information; our payment will reflect what we expect the information to be on the basis of our prior probability assessment.

In the husband's anniversary problem, perfect information might correspond to a secretary who is certain to tell him if today is his anniversary. If he could act on this information, he would buy flowers if it were his anniversary and would not buy flowers otherwise. Since he feels that there is a 0.2 chance the secretary will tell him that it is his anniversary, the expected utility of the outcomes if he bases his decision on perfect information is $0.2\,u(\$100 - b) + 0.8\,u(\$80 - b)$ where b is the amount he must pay to get the information. By setting this expression equal to 0.734, the expected utility of his best alternative based on prior information, we can solve for $b = \$33.50$. The husband should consider for more detailed analysis only those opportunities for resolving his uncertainty that "cost" him $33.50 or less. If he were an expected value decision maker, perfect information would be of less value to him; he would be willing to pay a maximum of only $20 for it.[9]

Summary

Decision theory is a way of formalizing common sense. The decision maker analyzes the possible outcomes resulting from his available alternatives in two dimensions: value (by means of utility theory) and probability of occurrence. He then chooses the alternative that he expects to have the highest value. He cannot guarantee that the outcome will be as good as he might hope for, but he has made the best decision he can, based on his preferences and available knowledge. Inference using Bayes' rule allows the decision maker to evaluate information gathering activities that will reduce his uncertainty.

Decision theory gives no magical formulas for correct decisions. In fact, it forces the decision maker to rely more strongly than ever on his own preferences and judgments. But it does give him a logical framework in which to work, a framework that is adaptable in principle to all decision problems, from the simplest to the most complex. As modern society continues to grow in size and complexity, such a framework for decision making will become more and more necessary.

[9] Additional discussion regarding the value of information in decision theory is available from many sources, most notably Howard [8b], [9], [11] and Raiffa and Schlaifer [19].

Acknowledgment

The author is deeply indebted to Prof. R. A. Howard for valuable guidance and advice. Much of the viewpoint presented in this paper was acquired as a result of two years of assisting Prof. Howard in his course in Decision Analysis at Stanford University, Stanford, Calif.

References

[1] F. J. Anscombe and R. J. Aumann, "A definition of subjective probability," *Ann. Math. Statist.*, vol. 34, pp. 199–205, 1963.
[2] T. Bayes, "An essay toward solving a problem in the doctrine of chances," *Phil. Trans. Roy. Soc.* (London), vol. 53, pp. 370–418, 1763.
[3] R. T. Cox, *The Algebra of Probable Inference.* Baltimore, Md.: The John Hopkins Press, 1961.
[4] P. C. Fishburn, *Decision and Value Theory.* New York: Wiley, 1964.
[5] ——, "Utility theory," *Management Sci.*, vol. 14, pp. 335–378, January 1968.
[6] ——, "On the prospects of a useful unified theory of value for engineering," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-2, pp. 27–35, August 1966.
[7] R. A. Howard, "Bayesian decision models for system engineering," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-1, pp. 36–40, November 1965.
[8] a) ——, "Decision analysis: applied decision theory," *Proc. 4th Internat'l Conf. Operational Res.* (Boston, Mass., 1966).
b) ——," The foundations of decision analysis" this issue, pp. 211–219.
[9] ——, "Information value theory," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-2, pp. 22–26, August 1966.
[10] ——, "The science of decision-making" (unpublished).
[11] ——, "Value of information lotteries," *IEEE Trans. Systems Science and Cybernetics*, vol. SSC-3, pp. 54–60, June 1967.
[12] E. T. Jaynes, "Probability theory in science and engineering," Field Research Laboratory, Socony Mobil Oil Co., Inc., Dallas, Tex., 1959.
[13] H. Jeffreys, *Theory of Probability*, 3rd ed. Oxford: Clarendon Press, 1961.
[14] B. Klein and W. Meckling, "Applications of operations research to development decisions," *Operations Res.*, vol. 6, pp. 352–363, May–June 1958.
[15] P. S. Laplace, *Essai Philosophique sur les Probabilités* (Paris, 1814), translation. New York: Dover, 1951.
[16] R. D. Luce and H. Raiffa, *Games and Decisions: Introduction and Critical Survey.* New York: Wiley, 1957.
[17] J. W. Pratt, H. Raiffa, and R. Schlaifer, "The foundations of decision under uncertainty: an elementary exposition," *J. Am. Statist. Assoc.*, vol. 59, pp. 353–375, June 1964.
[18] ——, *Introduction to Statistical Decision Theory.* New York: McGraw-Hill, 1965.
[19] H. Raiffa and R. Schlaifer, *Applied Statistical Decision Theory.* Boston, Mass.: Graduate School of Business, Harvard University, 1961.
[20] L. J. Savage, *The Foundations of Statistics.* New York: Wiley, 1954.
[21] R. Schlaifer, *Probability and Statistics for Business Decisions.* New York: McGraw-Hill, 1959.
[22] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, 2nd ed. Princeton, N.J.: Princeton University Press, 1947.

Reprinted from IEEE TRANSACTIONS ON SYSTEMS SCIENCE AND CYBERNETICS
Vol. SSC-4, No. 3, September 1968
pp. 200–210

6

A TUTORIAL IN DECISION ANALYSIS

Carl-Axel S. Staël von Holstein

Scandinavian Airlines System

Presented in a slightly different form at the Third Research Conference on Subjective Probability, Utility, and Decision Making, London, 7-9 September 1971. Revised July 1973.

A TUTORIAL IN DECISION ANALYSIS[1]

1. Introduction

The papers presented at this conference have primarily dealt with theory and experiments relating to probability encoding and/or preference theory. It is often relevant to ask "What is the purpose of this research?" and "Where would its results become useful?" Hopefully, the answer is that this research will improve on the methodology of decision-making.

This paper is intended as an introduction to decision analysis, which in my opinion represents the present state of the art with respect to the methodology of decision-making. This paper provides a frame of reference for the rest of the research presented at this conference and may even provide some ideas for future research.

This tutorial will relate closely to the ideas on the methodology of decision analysis first described by Howard (1966). To a large extent this paper will also draw on the experience of the Decision Analysis Group at Stanford Research Institute with respect to practical implementation of decision analysis.

Decision analysis can be briefly described as a merger of the two fields of decision theory and systems analysis. Decision theory provides the philosophy for logical behavior in simple decision situations under uncertainty. Systems analysis here represents systems and modeling methodology, which captures the interactions and dynamics of complex problems. The result is a theory and methodology that allow the analysis of complex, dynamic, and uncertain decision situations.

Most textbooks on quantitative methods for business students include at least one chapter on decision theory. The presentations generally concern very small well-structured decision situations and thereby introduce some of the basic concepts. It may then be easy for the new M.B.A. to believe that he can go out and tackle decision problems in the real world. Most likely he (and perhaps even more so his superiors) will be discouraged when he has to face the complexities of even "simple" practical problems. It is here we find a need for the modeling methodology, which is never taught in decision theory courses.

This paper does not deal with decision theory as such, although its elements, of course, are part of the presentation. There exist a number of excellent introductions to decision theory and references to them are given in Section 7. Instead an attempt is made to describe the engineering of decision analysis, i.e., to give some guidelines on how to attack decision problems. The decision analysis cycle represents a procedure that might serve as a frame of reference when working on decision problems. It must be stressed, however, that not every problem can be or should be treated in the same way, and the elements of the cycle should only be viewed as a convenient check list to ensure that no important element of the problem has been omitted.

1.1 The Decision Analysis Approach

The decision analysis cycle can be summarized as follows (see Figure 1). It is made up of three phases--the deterministic, probabilistic, and informational phases. The deterministic phase is concerned with the basic structuring of the problem. The structuring entails defining relevant variables, characterizing their relationships in formal models, and assigning values to possible outcomes. The importance of the different variables is measured through sensitivity analysis.

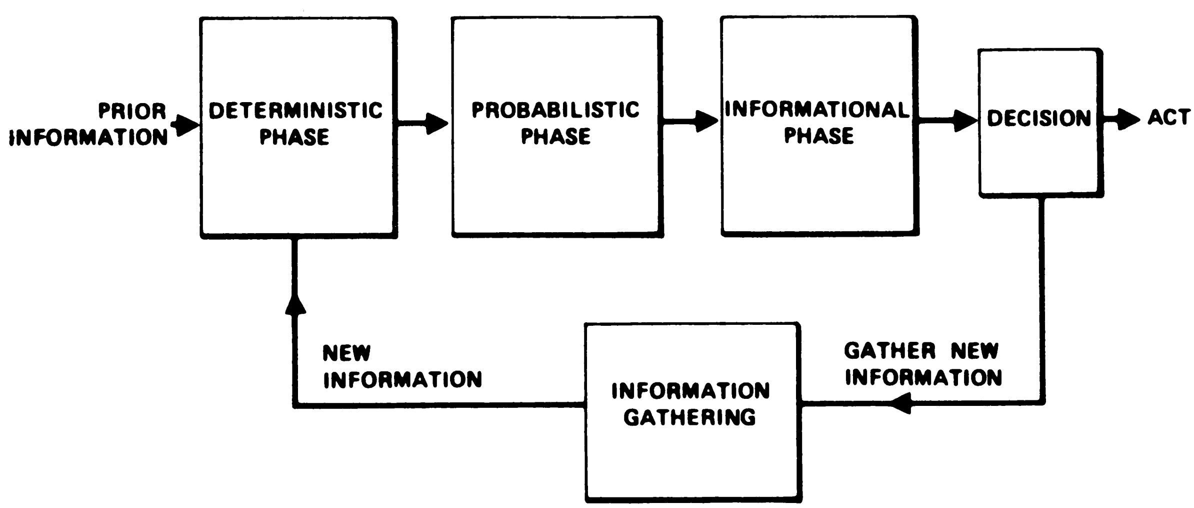

FIGURE 1 THE DECISION ANALYSIS CYCLE

Uncertainty is explicitly incorporated in the probabilistic phase by assigning probability distributions to the important variables. These distributions are transformed in the model to exhibit the uncertainty in the final outcome, which again is represented by a probability distribution. After the decision maker's attitude toward risk has been evaluated and taken into account, the best alternative in the face of uncertainty is then established.

The informational phase determines the economic value of information by calculating the worth of reducing uncertainty in each of the important variables in the problem. The value of additional information can then be compared with the cost of obtaining it. If the gathering of information is profitable, the three phases are repeated again. The analysis is completed when further analysis or information gathering is no longer profitable.

Throughout, the analysis is focused on the decision and the decision maker. That is, expanding the analysis is considered of value only if it helps the decision maker choose between the available alternatives.

1.2 A Case Study

The next three sections of the paper describe the three phases in more detail. The presentation is illuminated by a case study.[2] By request of the client company, the identity of the client and the specific decision problem are not identified. However, the major decision is similar to one that an agricultural subsidiary of a major diversified corporation might face in determining whether to market a newly developed biodegradable pesticide.

Some of the characteristics that made this decision problem a classical corporate application of decision analysis are listed below:

- The decision was one of a kind in that it represented a major change in the company's major product line.
- The decision concerned an investment of $150 million, which was a significant portion of the organization's resources--in fact, the investment was more than its normal annual capital expenditure budget.
- The problem structure was complex.
- The problem included many uncertain factors.
- The project would have long-term effects (10 to 20 years).

This investment opportunity had been evaluated five times over the preceding six years. Another study was under way when the decision analysis effort was initiated.

The analysis was performed by Dr. Carl S. Spetzler before he took up his present position with the Decision Analysis Group at Stanford Research Institute. He also formulated the case in its present disguised form, and I am grateful to him for permission to use the case in this paper.

2. The Deterministic Phase

2.1 The Deterministic Model

The basic steps of the deterministic phase are as follows:

- Define and bound the decision problem
- Identify the alternatives
- Establish the outcomes
- Select decision variables and state variables
- Build a structural model
- Build a value model
- Specify time preference
- Eliminate dominated alternatives
- Measure sensitivity to identify crucial state variables.

The first step is to define and bound the decision problem. This entails determining the resources to be allocated and this in turn is related to the organizational level at which the decision is to be made. Next, the available alternatives are identified. The introduction of a new alternative sometimes eliminates the need for further analysis. In our analysis the basic decision problems were determining whether the new product should be introduced and determining the best method of production.

The next step is to identify outcomes that would be sufficient to describe the results of the different alternatives. These might include sales volume, production process, government action, and so on. In relating these outcomes to the alternatives, we try to define the factors that are relevant to the decision. These factors can be separated into decision variables (factors that the decision maker can control) and state variables (factors outside the decision maker's control). In this study these variables were identified by interviewing various experts with respect to this problem and asking what factors must be considered. A very impressive list of factors was developed within a couple of days' time; these factors were then classified as decision and state variables.

The next step is to build a structural model that relates the outcomes to the decision and state variables. This is generally the most important step in the deterministic phase and often the most time-consuming part of the whole analysis. A logical diagram such as that shown in Figure 2 was developed as a start toward such a structural model.

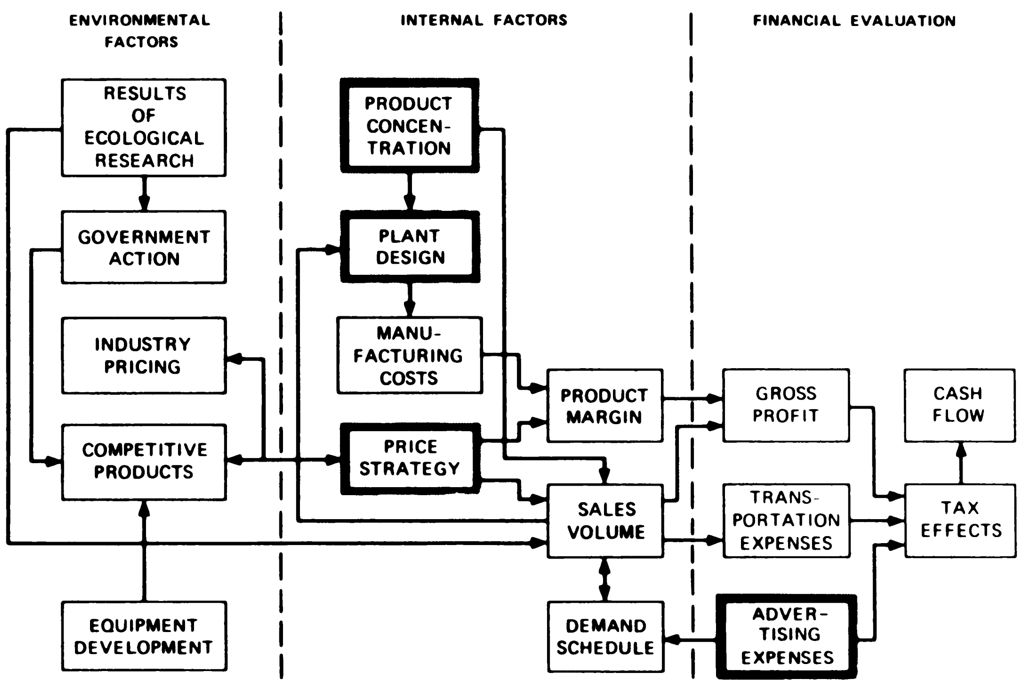

FIGURE 2 DETERMINISTIC MODEL STRUCTURE

Again, this diagram was the result of discussion with many individuals in the corporation. It began with very simple interrelationships and evolved into the form shown in Figure 2. The diagram was considered complete when the following analysis indicated no additional requirements. First, around each of the boxes an imaginary boundary was drawn and the arrows going in were listed. Then the question was asked: "Do you need any further inputs to determine the value of the variable in the box?" In many cases, simple calculations or algorithms were developed to prove that the inputs were sufficient. The process of developing such a diagram is useful for two reasons: (1) the analysts quickly become aware of the fundamental relationships that are of importance in the problem, and (2) the process allows the systematic gathering of information from various parts of the corporation.

After identifying the relationships among the variables the next task was to capture these relationships in mathematical forms. The basic alternative was assumed to be to continue without the product, and the model measured only the difference from this alternative of introducing the new product with one out of two production methods. The dynamic relationships of the variables were described in detail for the first 12 years, after which a terminal value based on the capitalized long-run profit of a stable business was assumed. The model outputs were measured in such terms as market share and net profit over time.

It has already been said that the outcome of a decision will generally be described by a set of outcomes such as sales volume and degree of government interference. The next step of the deterministic phase is to determine a single measure of value for this set of outcomes. In business problems the measure typically will be some measure of profitability. Determining this value means that trade-offs have to be considered between different outcomes. For this analysis it was assumed that the decision maker's values were completely reflected in the net contribution to profit after taxes. This net effect on profit was determined over time.

It often happens, as in this case, that the results of a decision extend over a long period of time. It is then necessary to have the decision maker specify his time preference for profit. This means that we must find a single measure for each time pattern. The present equivalent is such a measure, for which the decision maker is indifferent between receiving the present equivalent right away or waiting for the cash flow to be realized over a future time period. In many cases this present equivalent can be approximated by the present value of the cash flow discounted at an appropriate discount rate. In other cases the cash flow resulting from one project may have a major effect on the organization's financial structure and a more detailed analysis is then necessary. In the case under discussion it was assumed that the company's minimum acceptable rate of return completely reflected the decision maker's feeling of value over time. Therefore, the present value calculated at that discount rate was considered a reasonable measure of project worth as viewed by the decision maker.

2.2 Sensitivity Analysis

The analysis in the deterministic phase takes the form of measuring sensitivities to changes in state variables. The state variables are assigned nominal values (which might be, for instance, estimates of their mean values) and are then swept one by one through their ranges of possible values. We observe which alternative would be best and how much value is associated with this alternative. Sometimes we may observe that an alternative is dominated, which means that there is a better alternative for all values of the state variables. Dominance can often lead to a substantial reduction in the number of alternatives. The sensitivity analysis also tells us to what extent variations in the different state variables will affect profit (in terms of the present equivalent). The analysis indicates the variables for which uncertainty is important. These variables are said to be "crucial" and will have their uncertainties encoded in the next phase.

The preceding discussion assumes that the state variables can be considered one at a time. However, it may be necessary to study joint sensitivities when the variables are interrelated. Two or more variables are then varied simultaneously over their respective ranges; at the same time the other variables are kept at their nominal values. In the case study described in this paper, however, the interrelationships were handled in the following way. First a case was analyzed which combined all of the most likely forecasts for the input variables. Then the sensitivity to changes in each variable was analyzed by setting it to its high and low extreme values; at the same time all other variables were reset to new conditional most likely values. The results of the sensitivity analysis for some of the variables are given in Table 1. In the case study, seven crucial state variables were identified and three major decision alternatives with some minor variations remained. It is often the case in a decision analysis that only a few of the many variables under initial consideration are crucial state variables. This is of importance for the modeling process in the probabilistic phase. Probabilistic models with many variables are difficult to handle. Furthermore, the information required for such models is often difficult to derive. It is therefore a very important task to eliminate unnecessary decision alternatives and to limit the number of state variables to those crucial to the decision. Sensitivity analysis can also provide insight which is valuable to the building of the model, since variables to which the model shows a high sensitivity can often be further broken down to improve the model.

Table 1

RESULTS OF SENSITIVITY ANALYSIS

Variable	Present Value $ Million	
	Max	Min
Market Size	+149	-40
Manufacturing Costs	+32	-20
Government Action	+10	-48
Competition	+8	-10
Price of Substitutes	+24	-13
Results of Ecological Research	+11	-30

(Most likely case: $8 million)

3. The Probabilistic Phase

The deterministic analysis leads to the selection of a set of aleatory variables--variables to be formally treated by a probability assignment. The first step in the probabilistic phase is to encode the uncertainties of these variables. Next a probabilistic model is constructed which relates the uncertainty in profit for each decision alternative to the uncertainty in the aleatory variables. The resulting probability distribution for profit is termed the profit lottery.

The choice between a number of alternatives has now been reduced to a choice between profit lotteries. In some cases the choice is clear because one alternative stochastically dominates the others. This means that there will be one alternative, which for each level of profit has a higher probability of exceeding that level than all other alternatives. Otherwise, it will be necessary to encode the decision maker's risk attitude. The result of that procedure is substitution of each lottery by a single number, called the certain equivalent, which has the property that the decision maker is indifferent between having the certain equivalent for certain or having the lottery. The different alternatives can now be ranked in order of their certain equivalents, which indicate the decision maker's preferences.

The probabilistic phase is then concluded by performing further sensitivity analyses. Here the effect of a variable is measured when all other variables are taken as uncertain (rather than kept at nominal values).

This is a brief overview of the probabilistic phase. It introduces a number of concepts which may not be familiar to many readers. The different steps of the probabilistic phase are therefore discussed in greater detail in the remainder of this section.

3.1 The Probabilistic Model

The probabilistic model expands the deterministic model to include the uncertainties encoded for the aleatory variables. The purpose of the probabilistic model is to develop profit lotteries for the different alternatives. It should be designed to include dependencies between variables, if such exist.

In the case study described in this paper, the probabilistic model was made up of a decision tree and a financial model attached to the end nodes of the decision tree. The tree structure represented the relationships between the different variables whereas the financial model included the value and time preference models from the deterministic phase. The first tree, which was structured after the sensitivity analysis of the deterministic phase, had approximately 2,000 terminal nodes. This tree was simplified by reducing the number of decision nodes (alternatives) on the basis of back-of-the-envelope calculations. The tree was further developed after extensive interaction with the organization's staff.

The order of the variables in the decision tree was based on convenience in terms of contemplating the information required for the tree. The sequence was not a time sequence, although that would be the case if a set of variables were dependent through time. Each terminal node on the tree represented a sequence of values of the various factors that were included in the decision tree structure. A deterministic financial model was then developed which derived the effect on net profit of the corporation for each sequence of factors. Both the logical structure of the tree and the financial model were programmed for a timesharing computer, which allowed a rather fast analysis of the decision tree and made it possible to easily revise the program.

3.2 Encoding Uncertainty

Probability is the language of communication about uncertainty. The personal interpretation of probability represents a cornerstone in the decision analysis philosophy. Probability represents a state of information and it is only natural that two persons can make different probability assignments to the same event, since they are likely to have different information bases. Furthermore, a person is likely to revise a probability assignment if he receives new and relevant information. This will be made more explicit in the next section.

The decision maker is the person (or group of persons) who has the responsibility for the decision under consideration. It follows that a decision analysis must be based on the decision maker's beliefs and preferences. He may be willing to designate some other person or persons as his expert(s) for encoding the uncertainty in a particular variable if he feels that the expert has a more relevant information base. The decision maker can then either accept the expert's information as his input to the analysis or modify it to incorporate his own judgment.

In a practical application experts will be drawn from different fields. Market variables, such as sales volume, are likely to come from the marketing department; production variables, such as manufacturing cost, will be provided by engineers. Some variables may even require experts from outside the organization. However, it must be made clear that the fact that a person is an expert in a particular area of the problem does not mean that it will be easy to elicit his judgment in probabilistic terms. Most people have difficulty in thinking about uncertainty. This means that they cannot directly express their knowledge about a variable in terms of a probability distribution. Rather, encoding techniques that make use of simple concepts for which they may have some understanding are used.

There has been very little written on the subject of how one should go about encoding the opinions of experts in practical situations. Spetzler and Staël von Holstein (1972) give an extensive presentation of probability encoding methodology in decision analysis. Most of the remaining is either literature, theoretical or related to laboratory experiments. An overview is given by Staël von Holstein (1970).

Practical experience has led us to conclude that most people have difficulty in expressing their judgment except for choices between simple events. The use of reference processes has proved useful here. The probability wheel is one example of such a process. This is a disk with two sectors, one blue and the other red with a fixed pointer in the center of the disk. The disk is spun finally stopping with the pointer either in the blue or the red sector (see Figure 3). A simple adjustment changes the relative size of the two sectors and thereby also the probabilities of the pointer indicating either sector when the disk stops spinning.

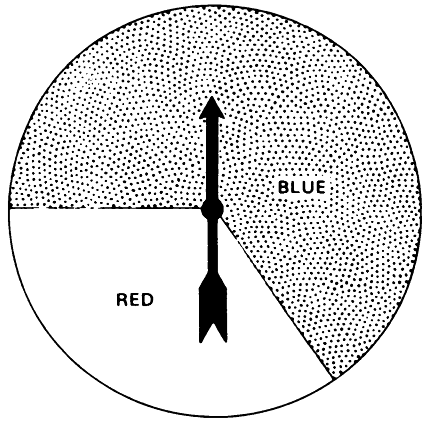

FIGURE 3 A PROBABILITY WHEEL

The probability wheel can be used in two ways. The expert can be asked whether he would prefer to bet either on a fixed event, e.g., that "next year's production will not exceed x units" or on the pointer ending up in the blue sector. The amount of blue in the wheel is then varied until the expert becomes indifferent. The amount of blue can also be kept fixed and the level of production is then varied until the indifference point is reached. When indifference has been obtained, the relative amount of blue in the wheel is assigned as the probability of the event.

A second approach is to use successive subdivisions. An interval is split into two (or more) parts and the expert is asked to choose which part he would prefer to bet on. The dividing point(s) is(are) changed until indifference is reached, and the subintervals are then assigned equal probabilities. Starting from an interval covering all possible outcomes, splitting into two parts will first give the median, then the quartiles, and so on. An illustrative example is given by Raiffa (1968, Section 7.3).

In concluding this rather long discussion on practical probability encoding, the following points should be stressed:

- It is better to conduct the encoding in private to eliminate group pressures and to make the process responsive to the expert whose judgment is being encoded.
- The quantity to be encoded should be important, otherwise credibility will be lost both with the decision maker and with the expert. The expert should be given an incentive to allocate time and effort to the process.

- The problem should be clearly structured. The production quantity above might have depended on whether or not a new process would work. It might then be easier to make probability assignments conditional on the process working and not working.

- In addition, the quantity should be described on a scale that is meaningful to the subject. For example, in the oil industry, the expert--depending on his occupation--may think in terms of gallons, barrels, or tank cars.

- Finally, the subject should not be worried about coherence. On the other hand, inconsistencies will be used as feedback in the encoding process to ensure that the final distribution is consistent with the subject's judgments.

It must also be stressed that the interviewing technique is by far superior to using questionnaires of various forms, unless the expert is very familiar with probability encoding. With questionnaires there is no way of finding out whether or not the assessor has understood the questions. An interactive computer program might provide a reasonable compromise. The Probability Encoding Program (PEP) developed by the Decision Analysis Group is an example of such a program. PEP makes use of successive subdivisions with two or three dividing points that are adjusted until the interviewed expert is indifferent. The fractiles encoded are those corresponding to probabilities 1/6, 1/3, 1/2, 2/3, and 5/6, and each fractile is encoded in two different ways to provide a coherence check.

Let us now return to the pesticide study. Seven variables had been selected as aleatory and their uncertainties were encoded with the help of the techniques mentioned above. More than one expert was used for each quantity and, as could be expected, the individual distributions often differed greatly. This was especially true of the distributions for "market size." It was generally found that agreement improved when individuals discussed their differing viewpoints and exchanged information. However, before spending much time and effort trying to get a consensus, the differences of opinion were tested in the probabilistic model to see whether they changed the choice of alternatives. If they did not, then it would suffice to fair one distribution to the set of distributions. However, it was found that the decision was indeed highly sensitive to judgment. Information had been encoded from the vice president of marketing, regional managers, market research staff, and various sales and marketing personnel and was presented to the decision maker for his reaction. After considering not only the information, but also the background and arguments presented by the individual assessors, he chose a distribution that he felt represented his best judgment regarding market size.

It is not uncommon that the decision maker has access to more than one expert. He is then confronted with the problem of how to reconcile the possibly different opinions. I would here like to mention a recent work by Morris (1971), who has given the most complete discussion of the problem of expert resolution within the decision analysis philosophy.

The cumulative probability distribution for market size that was used for the final analysis is shown in Figure 4. A step function which

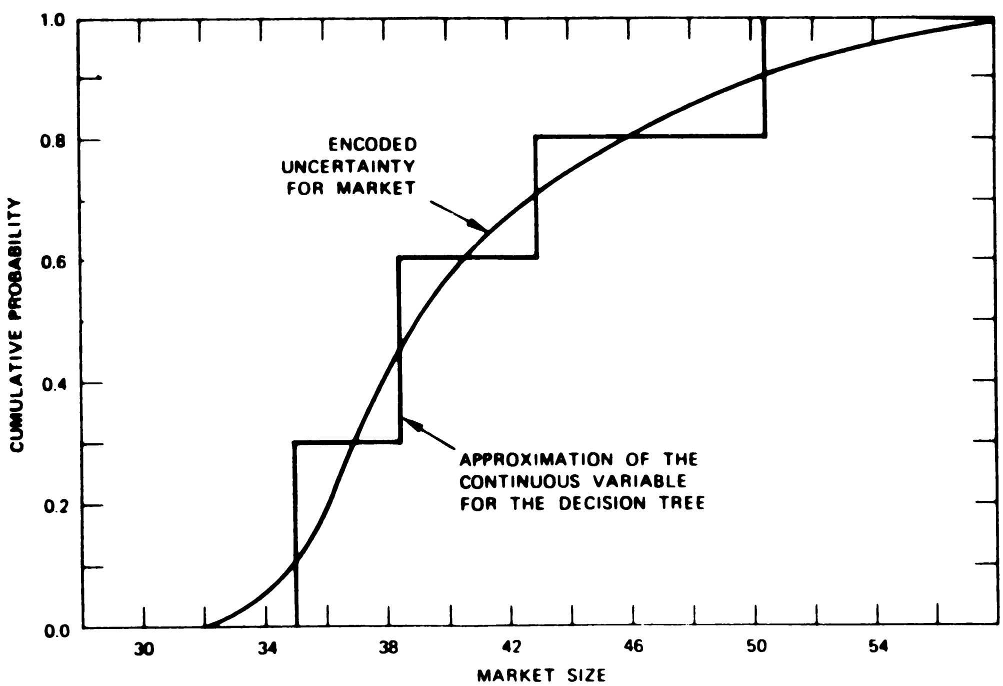

FIGURE 4 PROBABILITY DISTRIBUTION ON MARKET SIZE

was used as an approximation in the probabilistic model is also shown in the figure. Similar approximations were used for the other aleatory variables.

3.3 Developing a Profit Lottery

The profit lottery is simply the probability distribution of the present equivalent. The profit lottery is used to compare different alternatives and is discussed in Subsection 3.5. However, the calculation is not feasible, or at least it is not of any practical use, if the number of alternatives is too large. In such cases, it might be better to use other means to determine the optimal alternative and then only derive the profit lottery for the optimal alternative to display the uncertainty in profit.

The way the profit lottery is calculated depends on the structure of the probabilistic model. It might be possible to derive it analytically in a few cases, but it is more likely that some numerical approximation must be used. This approximation can be accomplished by formulating the probabilistic model as a decision tree or by using Monte Carlo simulation. In this paper, only the decision tree method is discussed, primarily because it seems to be the most useful approach. It provides further insight into the problem and also facilitates computations. The discussion will be in the context of the corporate decision.

In this decision problem it was possible to eliminate almost all decision nodes and it was therefore feasible to determine the complete profit lotteries for each decision alternative. In fact the tree that remained is better described as a probability tree than a decision tree. Figure 5 shows the tree structure that was used to analyze the various alternatives.

Each probability distribution has been approximated in the tree by a discrete distribution. The approximation to the encoded distribution improves as more branches are used, but at the same time the size of the tree grows rapidly, as does the cost of computation. A sensitivity analysis should decide the degree of approximation. Two to four branches for each distribution were used in the analysis of the case study. The approximated distribution for market size is given as an example in Figure 4.

Each end node of the tree represents an outcome and can be described by the values of the variables along the path leading to that node. There is a present value assigned to each node through the financial model. The probability of obtaining this present value is given by the product of the probabilities along the path. The present values are then sorted in increasing order and the cumulative probability distribution can then be plotted to summarize the profit lottery. The profit lottery for one alternative is shown in Figure 6.

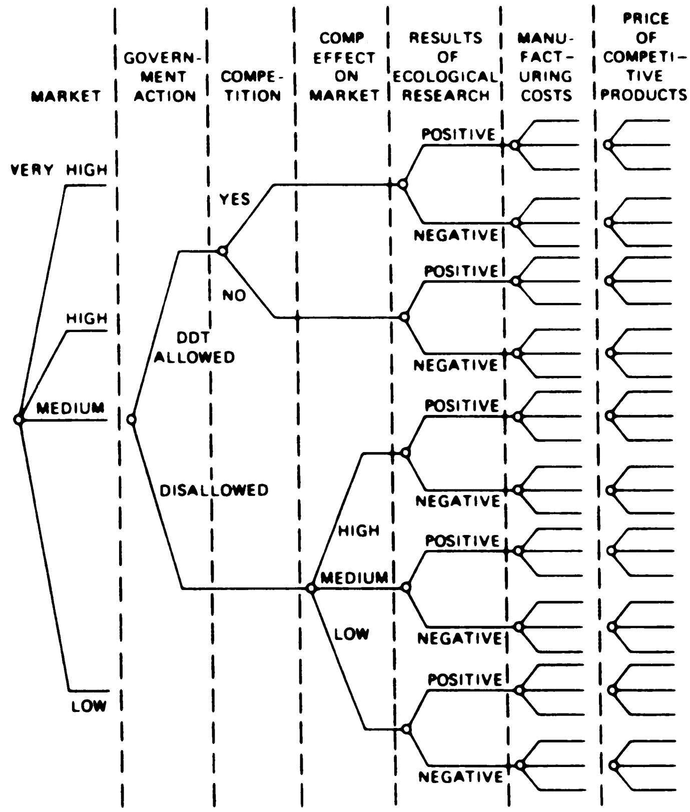

FIGURE 5 SIMPLIFIED PROBABILITY TREE STRUCTURE

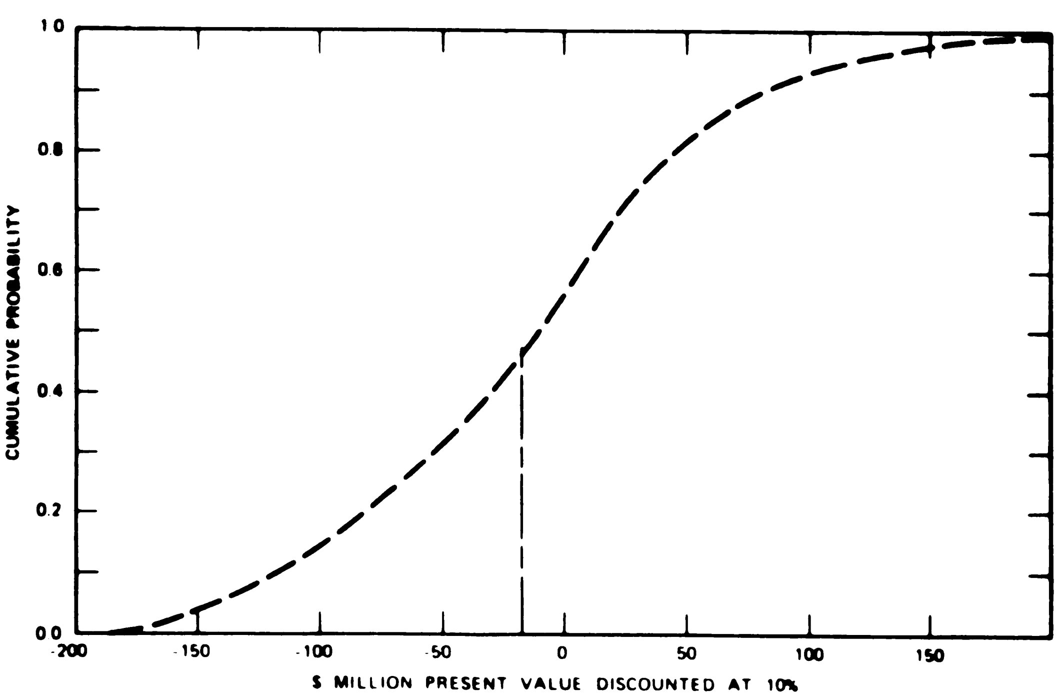

FIGURE 6 PROFIT LOTTERY

3.4 Encoding Risk Attitude

The question of risk attitude enters whenever a decision has to be made. A typical attitude might be that the decision maker prefers to have \$2 million for certain rather than a lottery with equal chances of winning \$20 million and losing \$5 million, i.e., he accepts a certain value smaller than the expected value of the lottery.

The risk attitude of the decision maker should be measured independently of any specific project. It is best done by questions similar to the one above. The decision maker is faced with two choices, one of which is riskless. For any lottery there is some riskless value which would make the decision maker indifferent. This value is called the certain equivalent. Generally, the certain equivalent is less than the expected value of the lottery and we then say that the decision maker is risk averse. The measurement procedure is continued with different lotteries until a good idea can be formed concerning the decision maker's risk attitude. The risk attitude is sometimes encoded in the form of a utility function and the best alternative is then the one that has the highest expected utility.

However, the encoding of risk attitude should be preceded by an analysis of sensitivity to risk attitude. It has been the experience of most practical studies at SRI that the decision is not very sensitive for reasonable risk attitudes, thus eliminating the need for further encoding. In some cases, because of stochastic dominance, there is not even a need for sensitivity analysis.

The decision maker can be expected to be risk neutral, i.e., willing to act on the basis of expected value, when the value of the project is not too large in relation to the organization's total worth. Otherwise his risk attitude is likely to be well approximated by an exponential utility function. This function has the property that if all outcomes of a lottery are augmented by the same amount, then the certain equivalent will increase by the same amount. This is an appealing property for a utility function, but it holds only for linear and exponential functions.

The exponential utility function facilitates the handling of risk attitude in two ways. The analysis of sensitivity to risk attitude is made easier since it is reduced to varying the one parameter of the utility function. If an exponential utility function seems to be a reasonable approximation to risk attitude, the encoding can be reduced to a few questions to check the consistency in the obtained values of the parameter. An example is discussed in Spetzler and Zamora (1971). The second advantage becomes more apparent in the informational phase.

Returning to the pesticide analysis, it was found that one alternative stochastically dominated all other alternatives except the non-action alternative. The profit lottery for this alternative is shown in Figure 6. As can be noticed, this profit lottery has an expected value of less than zero. In fact the probability of having less than zero present value is 55 percent. The conclusion from this profit lottery is that the decision maker would not be willing to introduce this new product, regardless of any risk aversion. Consequently, there is no need to measure the risk attitude in this example.

3.5 Determining Best Action

The basic question in the probabilistic phase is which alternative is best in the light of the available information. The answer is simply the alternative which has the highest certain equivalent. If risk indifference is assumed, then the certain equivalent is equal to the expected value. Otherwise, the risk attitude will have to be encoded in the form of a utility function and the best alternative is the one with the highest expected utility. The certain equivalent is then easily found since its utility is equal to this highest expected utility.

It is not, of course, necessary to describe each alternative by its profit lottery. The determination of the best alternative can easily be done within the decision tree structure. The alternative with the highest expected utility is found by performing a rollback analysis of the tree. The analysis works backwards from the end nodes through substituting certain equivalents for lotteries at probability nodes, and selecting the alternative with the highest certain equivalent at decision nodes.

3.6 Probabilistic Sensitivity Analysis

A decision analysis successively refines the decision model guided by sensitivity analyses. In the case study, the sensitivity analysis in the deterministic phase selected the aleatory variables and the model was improved as the analysis went into the probabilistic phase. Further sensitivity analyses helped determine the level of encoding of uncertainty and risk attitude. It is also probable that a study of sensitivity to time preference was made somewhere in the analysis.

It is now only natural that an analysis be performed to study the effect on the decision of the different variables within the probabilistic model. The probabilistic sensitivity indicates how the certain equivalent depends on a particular state variable when the other state variables are taken with their assigned probability distributions. It may show that variables that were thought of as important in the deterministic phase are relatively unimportant in the probabilistic environment. It also gives a measure of the robustness of the model.

Let us consider the effect of a "very high" market size in the pesticide example. This can easily be analyzed with the decision tree structure by restricting the tree in Figure 5 to only the very high path after the market node. This is equivalent to cutting off all other market branches and substituting probability one for the original probability of 0.2. It is then simple to reevaluate the tree. The resulting profit lottery is shown in Figure 7 together with the profit lotteries resulting from the other market sizes. The analysis indicates that the decision would not change until the volume is either "high" or "very high." In fact, even the profit lottery resulting from the "high" branch of the tree did not look particularly good if the possible risk aversion of the decision maker was considered. The analysis of market size thus showed that the decision was very sensitive to that variable.

A substantial effort in earlier studies of this decision had concerned the production process, and then primarily the manufacturing cost. The effect of manufacturing cost on the profit lottery is shown in Figure 8 and it is clear that it is unimportant.

The conclusion at the end of the probabilistic phase was that given the present level of information the best alternative was to do nothing.

4. The Information Phase

The analysis of the pesticide case in the preceding subsection gives some indication of what additional information would be most useful. For example, it is clear that better information on manufacturing cost would have little value since the decision would hardly change whatever the information might be. On the other hand, it is equally clear that if it were revealed that the market size would be "very high," the best decision would be to introduce the new product, which then would have a substantial expected value and little risk.

Information can be gathered in many ways--through discussion with experts, market surveys, pilot plants, and so on, depending on the context. The information is likely to have two characteristics: it will have some cost attached to it and it will not be perfect. The purpose of the informational phase is to evaluate different information gathering schemes and to then compare the values with the costs of using them. This means that we want to be able to answer questions such as: "Would it be worth $5 million to obtain perfect information on the market size?"

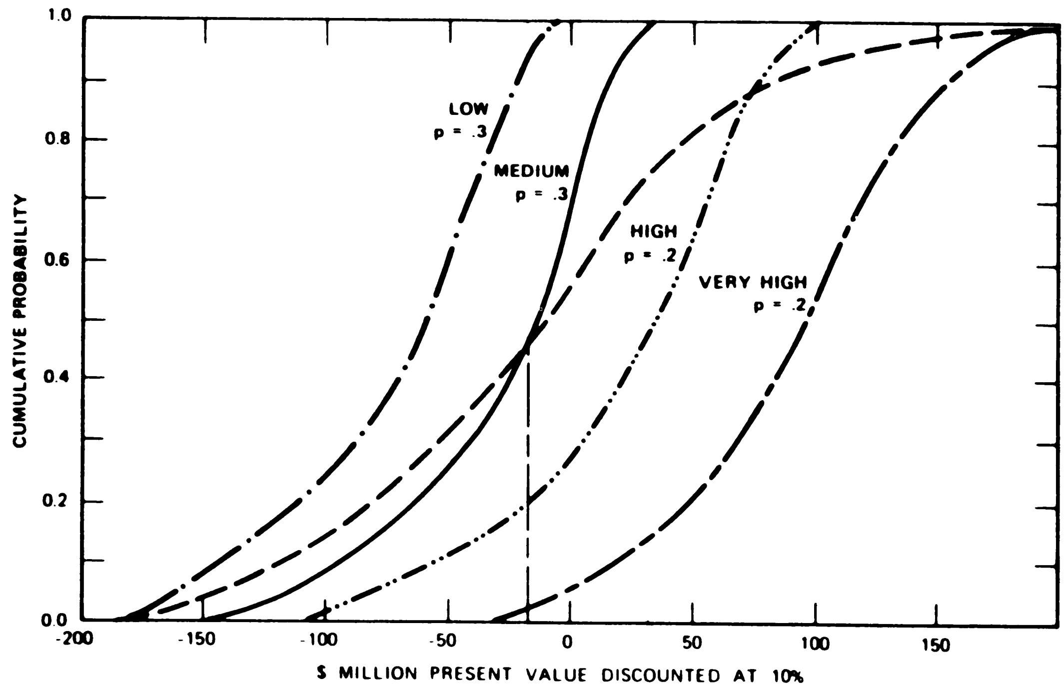

FIGURE 7 EFFECT OF MARKET SIZE ON PROFIT LOTTERY

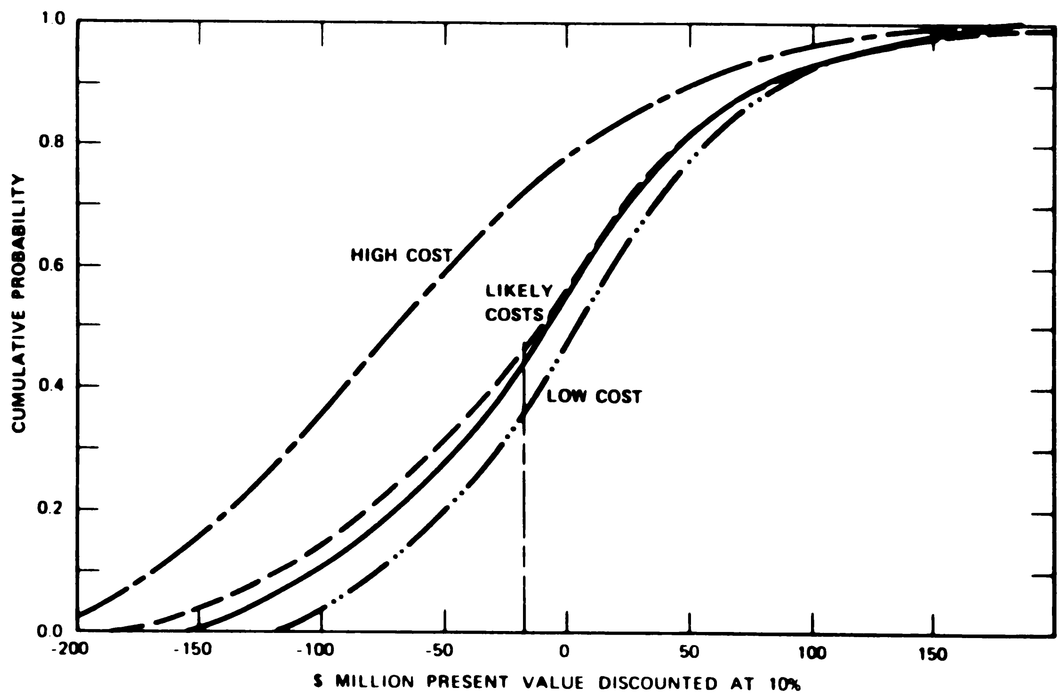

FIGURE 8 EFFECT OF MANUFACTURING COSTS ON PROFIT LOTTERY

The informational phase is thus intended to guide the decision maker in his search for further information that can help improve his basic decision. This seems to be one of the unique features of decision analysis. The reason is simply that the value of information cannot be calculated without a sound probabilistic model that is based on the decision maker's judgment and preferences. At the same time, this evaluation will be easily performed within the decision tree structure.

The first step in the informational phase is usually the calculation of the value of perfect information. Perfect information is a special case of any information gathering scheme and is seldom attainable in practice. However, there are two reasons why its value is determined first. One is that the value of perfect information represents the upper limit of the value of any imperfect information. It may be found, for instance, that the cost of any feasible information program is higher than the value of perfect information. In such cases consideration would not be given to getting more information. Another use of the value of perfect information is to suggest where it might be most valuable to look for feasible plans. If the value is low, then it will not be worth searching, but if it is high it might be worthwhile to expend some effort in looking for programs to improve on the information.

The value of perfect information is also easy to calculate as soon as the tree structure has been established. The procedure is simply to change the order of the nodes in the tree, placing the probability node representing the resolution of uncertainty before any decision nodes. The rest of the tree remains the same. This redesigning of the tree is very easily performed if the tree structure has already been programmed on a computer.

The value of perfect information for different state variables will suggest when imperfect information might be useful. It is then easy to incorporate the imperfect information scheme into the tree structure as a new probability node, representing the outcome of the information gathering, before the first decision node. The node representing the state variable in question remains in its old place in the tree, but the branches leaving the node now have new probabilities assigned to them based on the information received.

The value of information--perfect or imperfect--is easily calculated for the case of risk neutrality. It is then equal to the difference in expected values for the best alternatives with and without the information. The calculation is generally more complicated when the decision maker's risk attitude is to be considered. The value of information can then be found only through "trial and error." It is determined as the number that makes the expected utilities (or certain equivalents) equal for the best alternatives with or without information. However, the utility function will very likely be well approximated by an exponential utility function as was discussed in Subsection 3.4 and the value of information can then be calculated in the same way as when risk neutrality can be assumed. The reason is, of course, that the subtraction of the information cost from all terminal values leads to a reduction in the certain equivalent of the same amount.

Let us now look at the value of perfect information on market size in the pesticide example. Figure 7 shows that the best decision would be to continue to full-scale marketing if the market size was very large and to do nothing if the market size was less than very large. The profit lottery conditional on a very high market size extends from \$35 million to \$200 million with an expected value of more than \$90 million. However, this outcome of the information has only a 20 percent probability, and with 80 percent probability there would be no change in the decision. The value of perfect information is therefore about \$18 million.

A two-year market test was considered and its cost was roughly estimated at \$4 million. A first analysis of the decision on whether or not to test is presented in Figure 9. It assumes that the test would yield perfect information and that further analysis would have to assess the quality of the test. The expected value of the market test is around \$14 million. The decision, however, was not as clear as it may seem. The decision maker was essentially faced with a lottery with a rather high probability (80 percent) of losing about \$4 million on a market

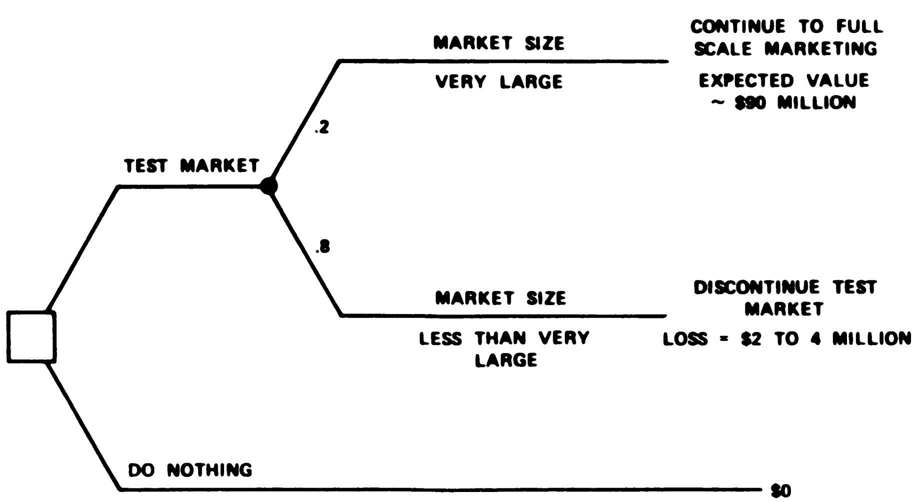

FIGURE 9 A FIRST CUT AT TEST MARKET DECISION

test for a product that seemed unprofitable in the first place and a low probability of gaining $90 million. To him the first outcome had a very negative value that was hardly compensated for by the large positive outcome, which would make him a "good guy." This is where the decision analysis ended; the first decision to be made concerned the market test, the subsequent decisions would be based on the outcome of the test.

5. Post Mortem

The decision analysis cycle may be iterated a number of times. If a decision is made to gather new information, it will lead to revised probability assignments and bring the analysis back to the informational phase. It is also likely that the new information will lead to new insight into the basic decision and will perhaps also suggest new alternatives, in which case we would be back to the deterministic phase.

It is often interesting to ask what decisions were actually made, and curiosity might also prompt a question with respect to the outcome of the decision. The second question is not very important. Knowledge of the outcome seems to create very good hindsight! This in turn leads to an evaluation of the decision maker based on the outcome rather than on the decision, which, of course, is completely wrong. Let me make the distinction very clear between good decisions and good outcomes. A good decision is based in a logical way on the decision maker's judgment and preferences, whereas a good outcome is loosely speaking a desirable outcome. There is no way to ensure good outcomes (unless there is no uncertainty entailed), but by making good decisions we are more likely to enjoy good outcomes.

It should therefore be clear that decision makers should be evaluated by the decisions rather than the outcomes. Far too often the evaluations of decisions are made after the outcomes have become known and it is then almost impossible to avoid having the actual outcome influence the opinion of what the best judgment should have been prior to this information. The only way to reach an honest evaluation of the decision is to have it well documented with respect to models and what has been included in them and what has been left out, and with respect to inputs in the form of assigned probability distributions and encoded preferences, and so on. This will be very difficult to accomplish if the original analysis was not made in the form of a decision analysis.

6. Further Aspects of Decision Analysis

This section considers a number of different topics related to decision analysis that did not find a place in the more formal presentation of the decision analysis cycle.

6.1 Decision Analysis as a Language

The preceding presentation has primarily been aimed at showing how decision analysis provides a logical and quantitative procedure for decision-making in uncertain situations. Another aspect is that it is also a language for communicating about decision problems. The models show what variables have been included and what their relationships have been assumed to be. Judgment about uncertainties is encoded in the form of probability distributions, which should facilitate communication between experts trying to reach a consensus. The formulation of a corporate risk policy should make it easier to communicate about risk attitudes and should eliminate the common observation that individual decisions are often based on very different risk attitudes (Spetzler, 1968).

Communication between different parts of the organization is often difficult to maintain because the different parts use different languages or jargons. Decision analysis has the great advantage of encouraging communication. For example, it provides a common language that makes it possible for engineers and sales people to understand their different contributions to the problem.

6.2 Interactions with the Decision Maker

How should the decision analyst interact with the decision maker? How are decision analysis recommendations implemented? These are two different questions that frequently are posed to members of the Decision Analysis Group, and the answers are closely related. The only way to get an analysis accepted and acted on by a decision maker is to interact with him throughout the analysis. He is very likely to be included at the beginning of the analysis when the problem is formulated. Later he may not be working actively on the project, but he will instead designate experts within (and sometimes outside of) the organization to work with the analyst. It is important, however, that contact is maintained, especially with respect to "educating" him in the decision analysis language. This is preferably done in the context of his own decision and can sometimes be a long process. When the decision maker understands that decision analysis provides him with the alternative that is consistent with his preferences and his judgment (or that of his chosen experts) then there is seldom a great problem in implementing the decision. Exceptions might occur when there are political problems, irrespective of whether or not the problem was related to government or industry, and even though the severity of such problems might be reduced if the problems are incorporated in the analysis.

It should be stressed that the decision analyst provides only expertise within his own field and that he does not pretend to be an expert in the problem area. He must be very careful not to make his own judgments or values influence a decision.

It is also important for the benefit of the project that the organization's staff get a good understanding of decision analysis. It will make it easier for them to see why certain factors are important or why judgment has to be encoded. It is often an important part of a project to train members of the organization in decision analysis, thus providing an in-house capability.

6.3 Engineering a Decision Analysis

It has been indicated in the discussion of the decision analysis cycle that decision analysis is very much an engineering approach to decision-making. That is, the aim is to construct a good enough model with the given resources. It is clear that very elaborate models can be constructed and that their inputs can be extensively refined, although the cost of doing so is likely to more than offset the value. A good engineering design should have the property that additional modeling will have equal value in all parts of the model.

The construction of the model is an iterative process in which the model becomes more refined as more is learned about the decision. This is often accomplished by careful selection of sensitivity analyses and information gathering schemes.

The decision tree represents a model and the choice of the number of branches to represent a probability distribution provides a good example of an engineering problem. Going from five to ten branches will in most cases have little effect on the decision, although the computational cost will increase significantly. Other examples are found in the modeling of time and risk preference. It is often a good approximation to use a discount rate to describe time preference and exponential (and sometimes even a linear) utility function to describe risk attitude.

7. Selected Reading

It has been said previously in this paper that decision theory provides the logical basis for decision-making in simple uncertain situations. Even though the literature on decision theory is extensive, there are surprisingly few sources of information that present a good discussion of the fundamental concepts. The article by North (1968) is the shortest exposition available. Lindley (1971) gives a penetrating nonmathematical discussion of the logical foundations; this study probably makes the easiest reading among the works cited here, and it is at the same time the most illuminating. Raiffa (1968) and Schlaifer (1969) both present very thorough expositions of decision theory. The former is the more technical of the two and contains a few advanced topics. Pratt *et al.* (1965) are more strongly directed toward statistical problems in decision theory.

The basic references to the methodology of decision analysis are provided by Howard (1966, 1968). These references differ in that the former is a completely verbal discussion, whereas the latter includes some mathematical formalism. A special issue on decision analysis published in 1968 by the *IEEE Transactions on Systems Science and Cybernetics* covers a variety of topics in addition to the two fundamental articles by North (1968) and Howard (1968). The contributions are drawn from a wide range of disciplines such as economics, statistics, psychology, and engineering and provide a great deal of insight into the general area of decision analysis.

There have not been many practical applications described in the literature, primarily because most of the studies that have been performed have contained proprietary material. Howard (1966) presents an example of a new-product introduction. Matheson (1969) gives short summaries of three applications; one application concerns a new-product introduction, another concerns space project planning (unmanned exploration of Mars), and the third concerns a decision as to when (and whether) to install a nuclear generating plant in Mexico. Spetzler and Zamora (1971) give a fairly detailed discussion of a case of a facilities investment and expansion problem. Howard *et al.* (1972) present the essential parts of the analysis of the strategic decision as to whether experimental seeding of hurricanes should be permitted.[3]

All these examples are drawn from work done by the Decision Analysis Group and give a representative picture of its projects over the past several years. The limitation to applications from this group is easily explained by the fact that there are no other published examples. A report by Brown (1971) on marketing applications--which he describes as "personalist decision analysis"--includes some case studies that seem to be applications of decision theory rather than decision analysis. The report is an outgrowth of an article (Brown, 1970) in which he discusses whether or not managers find "decision theory analysis" useful. The experience seems to based on studies that were performed either on a rather limited time scale or by people without proper qualifications. Decision theory is today being taught at most business schools in the United States, which means that every year thousands of M.B.A.s graduate and some of these will very likely try to apply decision theory. It seems very clear that their training is inadequate for the problems they sometimes attack (this is no criticism of the M.B.A.s; the criticism should be directed to their superiors). Therefore, it is not surprising that many managers are left with disillusioned views of decision theory. Some of these cases would make good examples of how decision analysis should not be performed. Hopefully, an increasing number of well-performed decision analyses, in the sense used in this paper, will change the evaluation.

Footnotes

[1] This paper was written while the author was with the Economic Research Institute at the Stockholm School of Economics. It was supported by a grant from the Bank of Sweden Tercentenary Fund. The author is now with the Decision Analysis Group, Stanford Research Institute. He would like to thank Michael Menke, Warner North, and Carl Spetzler for comments on an earlier draft of the paper.

[2] I did not include this case study in my talk at the conference, but gave instead a presentation of a decision analysis of hurricane modification. That study has now been written up by Howard et al. (1972).

[3] This study was presented in my talk at the conference.

References

Alpert, M., and Raiffa, H. A Progress Report on the Training of Probability Assessors. Unpublished manuscript, Harvard University, 1969.

Brown, R. V. Do Managers Find Decision Theory Useful?, Harvard Business Review, 48 (1970), 78-89.

__________, Marketing Applications of Personalist Decision Analysis. Working paper, Marketing Science Institute, 1971.

Howard, R. A. Decision Analysis: Applied Decision Theory. In D. B. Hertz and J. Melese (Eds.) Proceedings of the Fourth International Conference on Operational Research. New York: Wiley, 1966 (1968), pp. 55-71.

__________, The Foundations of Decision Analysis, IEEE Transactions on Systems Science and Cybernetics, SSC-4 (1968), 211-219.

__________, Matheson, J. E., and North, D. W., The Decision to Seed Hurricanes, Science, 176, (1972), 1191-1202.

Lindley, D. V. Making Decisions. New York: Wiley-Interscience, 1971.

Matheson, J. E. Decision Analysis Practice: Examples and Insights. In J. Lawrence (Ed.) Proceedings of the Fifth International Conference on Operational Research. London: Tavistock, 1969 (1970). pp. 677-691.

Morris, P. A. Bayesian Expert Resolution. Unpublished Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, 1971.

North, D. W. A Tutorial Introduction to Decision Theory, IEEE Transactions on Systems Science and Cybernetics, SSC-4 (1968), 200-210.

Pratt, J. W., Raiffa, H., and Schlaifer, R. Introduction to Statistical Decision Theory (preliminary edition). New York: McGraw-Hill, 1965.

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

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

Spetzler, C. S. The Development of a Corporate Risk Policy for Capital Investment Decisions, IEEE Transactions on Systems Science and Cybernetics, SSC-4 (1968), 279-300.

__________, and Stael von Holstein, C.-A.S., Probability Encoding in Decision Analysis. Unpublished manuscript, Stanford Research Institute, 1972.

Stael von Holstein, C.-A.S. Assessment and Evaluation of Subjective Probability Distributions. Stockholm: The Economic Research Institute at the Stockholm School of Economics, 1970.

__________, and Zamora, R.M., Decision Analysis of a Facilities Investment and Expansion Problem. Unpublished manuscript, Stanford Research Institute, 1971.

7

THE SCIENCE OF DECISION-MAKING

Ronald A. Howard

Department of Engineering-Economic Systems
Stanford University

A recent court case that received national attention indicated how decision making has become a technical concern in our modern society. The case concerned an elderly lady who was knocked down and robbed of her handbag by a young woman. The attack was witnessed by a man who was certain only that the young woman wore her blond hair in a pony tail and jumped into a yellow car driven by a bearded, moustached man. Inquiries in the local neighborhood produced a couple who fitted this description, but neither the victim nor the witness could make a positive identification of the assailants. Because of this lack of positive identification and the fact that the purse was never found after the attack, experienced legal opinion was that no conviction would be possible.

However, the resourceful prosecutor called as an expert witness a mathematics professor who was an authority on probability. At the trial he placed this expert witness on the stand and asked him questions and answers that went something like this:

Q. What is the probability of a man having a moustache?

A. One in three.

Q. What is the probability of a young woman having blond hair?

A. One in four.

The questioning continued in this manner until a probability had been assigned to each of the factors in the description of the crime. Then the witness multiplied the probabilities of the factors together and stated that the chance that the crime was committed by some other couple who happened to meet the same description was one in twelve million. Although this statement rested on many assumptions that we would consider questionable, it was sufficient along with other circumstantial evidence to convince the jurors--they returned a guilty verdict. When the defense attorney objected to this type of expert witness, the presiding judge replied: "Probability has a sound and proper basis. The law provides experts in any field may be used where they have knowledge that is not known to the average person."

Not often does the modern theory of probability used in decision-making come to our attention in such a dramatic setting. Yet it is symptomatic of a minor revolution in the way decisions are being made in this country. The revolution is partly the result of the progress in computation that has occurred over the past twenty years. As the computers have increasingly shown their ability to handle the routine tasks of business, the desire has grown to apply them to the higher functions of management and, in particular, to decision-making. However, the more direct causes of the revolution are found in the academic community. The theory of probability, which was formerly thought to be of interest only to technical specialists, has been shown to be of fundamental relevance to everyone who makes decisions. In what follows we shall explore the emergence of this view and its implications for modern decision-making.

Assertion 1: <u>The process of decision-making is at the heart of most technical, business and governmental problems</u>. Engineers must make decisions when they consider a change in the design of a manufactured item. Marketing executives must decide on the territories and intensities for sales campaigns. Governmental officials must decide on the apportionment of funds for area redevelopment. We could all name many more examples and we might conclude that all real problems are decision problems. Even in our personal lives we face decisions like where to go to college or where to take a vacation or perhaps whom to marry. Although these examples vary in their susceptibility to quantitative analysis, they all fit within the structure of decision theory.

Assertion 2: <u>Decision-making requires the study of uncertainty</u>. The importance of uncertainty is revealed when we realize that decisions in situations where there is no random element can usually be made with little difficulty. Only when we are uncertain about which of a number of possible outcomes will occur do we find ourselves with a real decision problem. For example, suppose that we are planning to take a trip tomorrow and that bad weather is forecsst. We have the choice of flying or taking a train. If someone told us the consequences of each of these acts, then our decision would be very simple. Thus, if this person said that the train would depart at 9:13 and arrive at 5:43, if he described in detail the nature of the train accommodations, the menu in the dining car, the people whom we would meet as travelling companions, then we would have a very clear idea of what taking the train implied. If he further specified that the plane would leave two hours late and arrive two and half hours late, that during a certain portion of the trip the flight would be especially bumpy, and in addition described the meals that would be served and the acquaintances we would meet, then the flying alternative would be described as well. Most of us would have little trouble in making the decision about our means of travel when we considered these specified alternatives in the light of our tastes and desires. Thus the decision problem is difficult because of the uncertainty of departure and arrival times and, in the case of the plane, about whether the trip will be possible at all. The factors of personal convenience and pleasure will be more or less important depending upon the urgency of the trip and consequently so will the uncertainties in these factors. Thus we see that we cannot make a meaningful study of decision-making unless we understand how to deal with uncertainty.

Assertion 3: Uncertainty can only be studied formally through probability theory. Suppose that we desire to create a theory of uncertainty with the following properties:

1. The theory will deal only with unambiguous events so that we understand uniquely what is meant by any statement within the theory.

2. Uncertainty is to be measured by a number so that the uncertainties of different events can be compared.

3. The theory never introduces any assertions into the argument unless they have been explicitly introduced by the decision-maker.

Then it is possible to show using mathematics that the only theory consistent with these requirements is the theory of probability initiated by Pascal, Fermat and Bernoulli three hundred years ago, developed by Bayes and Laplace two hundred years ago, and studied seriously until today by mathematicians the world over. This theory of probability is the only theory of uncertainty that has this important property: the likelihood of any event following the presentation of a sequence of points of data does not depend upon the order in which those data are presented. So basic is this property that many would use it as the defining basis for the theory. If it were not inevitable we would consider ourselves fortunate that a theory so extensively investigated and developed turns out to be at the heart of all decision-making.

Assertion 4: Probability is a state of mind, not of things. A reasonable question at this point is the following: If probability is so essential to decision-making, why has its importance not been more widely appreciated until now? The answer is that many users of probability theory (but certainly not the original developers) considered probabilities to be physical parameters of objects just like weight, volume or hardness. For example, there was much mention of "fair" coins, and "fair" dice with the underlying notion that the probabilities of events associated with these objects could be measured in the real world. For the past fifteen years, however, an important minority of experts on the subjects have been advancing the view that probabilities measure our state of knowledge about phenomena rather than the phenomena themselves. They would say, for example, that when we describe a coin as "fair", we really mean that on the basis of all evidence presented to us we have no reason for asserting that the coin is more likely to fall heads than tails. This view is modern, but not a product of modern times. It was stated clearly and convincingly two hundred years ago by both Bayes and Laplace. We can only regret that such a powerful and fundamental notion remained buried for such a long time.

A colleague of mine has a cogent and entertaining example for driving home the point of this assertion. An astronaut is about to be fired into space on a globe-circling mission. As he is strapping himself into his capsule on top of a gleaming rocket he asks the launch supervisor, "By the way, what's the reliability of this rocket?" The launcher supervisor replies, "99%--we expect only one rocket in 100 to fail." The astronaut is reassured, but still has some doubts about the success of his mission. He asks, "Those rockets around the edge of the field, are they the same type as the one I'm sittting on?" The supervisor replies, "They're identical." The astronaut suggests, "Let's shoot up a few just to give me some courage." A rocket is fitted with a dummy pay load, prepared for launching and fired--it falls in the ocean a complete failure. The supervisor says, "Unlucky break, we'll try another." Unfortunately, that one also fails by exploding in mid-air. A third is tried with disastrous results as it disintegrates on its pad.

We can imagine what all this has done to the courage of our astronaut. By this time he has probably handed in his resignation and headed home. No power on earth could convince him that the reliability of the rocket he was to use is still 99%. And yet, what has changed? His rocket is physically unaffected by the failure of the other rockets. If probability were a state of things, then the reliability of his rocket should still be 0.99. But, of course, it is not. After observing the failure of the first rocket, he might have evaluated the reliability of his rocket at say 0.90; after the second failure, at 0.70; and finally after the third failure, at perhaps 0.30. What happened was that his state of knowledge about his own rocket was influenced by what happened to its sister ships and therefore his estimate of its reliability must decrease. His final view of its reliability is so low that he does not choose to risk his life.

The view of probability as a state of things is just not tenable. We should consider probability as the reading of a kind of mental thermometer that measures uncertainty rather than temperature. The reading goes up as data accumulates that tends to increase the likelihood of the event under consideration. The reading of 1 corresponds to certainty that the event will occur; the reading of 0 to certainty that it will not occur. The inferential theory of probability is concerned with the question of how the reading ought to fluctuate in the face of new data.

Assertion 5: All prior experience must be used in assessing probabilities. Most of us will agree that it would be unwise to make a decision without considering all the knowledge we had obtained prior to making the decision. If we were offered an opportunity to participate in a game of chance by our best friend, a tramp, and a business associate, we would generally have different feelings about the fairness of the game in each case. The major problem is how to encode the knowledge that we have in a usable form. This problem is solved by our observation that probability is the appropriate way to measure uncertainty. And, of course, a probability is a number that we can use in computations.

The difficulty in encoding our prior knowledge as probabilities is that prior information available to us may range in form from a strong belief that results from many years of experience to a vague feeling that arises from a few haphazard observations. Yet I have never met a person who had "no" information about an event that was important to him. People who start out saying that they have "no idea" about what is going to happen can always, when pressed, provide probability assignments that show considerable information about the event in question. The problem of those who would aid decision-makers is to make the process of assigning probabilities as simple, efficient, and accurate as possible.

Assertion 6: Decision-making requires the assessment of values as well as probabilities. We said that the problem of the traveler that we discussed earlier became simple when uncertainty as to the modes of travel was eliminated. More precisely we said that it became a question of taste and preference. One of the key factors in the decision-making process is the establishment of the value to be attached to each of the various outcomes of a decision. When faced with two completely specified future sequences of profits, costs and other consequences, the decision-maker must be able to say which he prefers and to state his preference in quantitative terms. In business problems the desirability of any outcome will usually be measured in terms of dollars, either directly measured as costs or revenues or implicitly assigned as in valuing customer goodwill and employee satisfaction.

The mathematical theory concerned with the assessment of value is called utility theory. Although this theory is not so widely known as probability theory, it is based upon probability theory and on some additional axioms. One of these axioms, for example, is the axiom of transitivity. This axiom states that if the decision-maker prefers outcome A to outcome B and if he prefers outcome B to outcome C, then he must prefer outcome A to outcome C. The theory will not be useful to a person who does not subscribe to this tenet. The other axioms are similar in kind and equally logical.

We all know that we may from time to time behave in a way that is inconsistent with these axioms or with the axioms of logic in general. The point is not whether we do act logically, but rather whether we want to act logically. That is, we are constructing a normative theory that will aid us in making more consistent and logical decisions rather than a descriptive theory that merely specifies our current decision practices.

We might summarize by saying that we haven't specified a decision problem until we have said what it costs us to be wrong. Decision in the absence of value is speculation rather than accomplishment. The computation of values may require extensive staff work and discomfitting executive soul-searching but it is a necessary function.

Assertion 7: Decisions can only be made when a criterion is established for choosing among alternatives. Suppose that probabilities have been assigned to various outcomes and that a value has been attached to each outcome. When this has been done for all alternatives, which of these alternatives should be selected? Should it be the alternative with the highest expected profit? The alternative with the minimum maximum loss? The alternative with the highest probability of the highest gain? The question is a difficult one that few decision-makers have faced squarely.

We can understand the difficulty when we consider the apocryphal problem of William Tell in shooting the apple off his son's head. As an experienced marksman, Tell had a good measure of the uncertainty in his impact point, so the encoding of his previous knowledge was relatively simple. Next, however, he had to construct the loss function. If he shot too low, he would kill his son; if he shot too high, he would lose prestige as a rebel leader and probably be imprisoned. Only if he hit the apple would they both be freed. Tell thus had to evaluate the outcomes: "son dead, Tell free"; "son alive, Tell imprisoned"; and "both free." Let us assume that Tell's arrow was equally likely to vary in all directions from the aim point. Then if Tell rated the outcome of "son dead, Tell free" equal to the outcome "son alive, Tell imprisoned," he would have aimed at the exact center of the apple. If he valued his position as a rebel leader more than his son's life, then he would have aimed slightly low but still, of course, at the apple so that he would be less likely to be imprisoned than to kill his son. On the other hand if he wanted to maximize the probability of the most favorable event, namely, their both being set free without regard to any other considerations, then he should once more have aimed at the exact center of the apple.

We thus see that the assignment of probability and assessment of value are merely the first two steps in formalizing the decision problem. The establishment of the decision criterion plays an equally important role in the decision-making process. Experiments with small groups and large organizations have shown that the establishment of decision criteria is not a simple task. Individuals at different levels in organizations have different propensities for taking risks. They behave differently when using their own money from the way they do in making decisions regarding the company's money. Indeed, it is a real danger if the individual cannot look at the company's decision problems from the company's point of view rather than from his own.

Assertion 8: <u>The implications of the present decision for the future must be considered</u>. The influence of present actions upon the future is a point often disregarded by decision-makers. Unfortunately, a decision that seems appropriate in the short run may in fact place the decision-maker in a very unfavorable position with respect to the future. For example, a novice taxi driver may be persuaded to take a customer on a long trip to the suburbs by the prospect of the higher fare for such a trip. He might not realize, however, that he will have to return in all likelihood without a paying passenger, and that when all alternatives are considered, it could be more profitable for him to refuse the long trip in favor of a number of shorter trips that could be made within the city during the same time. Fortunately, we have at our disposal powerful techniques for handling just this type of problem.

Assertion 9: <u>We must distinguish between a good decision and a good outcome</u>. In everyday life we often do not recognize the distinction between good decisions and good results. For example, suppose that a man said that he never wanted to engage in any game of chance in which the odds were against him. If this man then paid $1.00 for a ticket in a lottery with 1,000 tickets outstanding and a prize of $100, we could describe his decision as illogical or, perhaps better, as inconsistent with his avowed goals. (We are assuming, of course, that no other motives like sympathy for the ticket seller influenced his actions.) In the same way a particular decision made by an agent of a company might be inconsistent with the company's official policy--we would also characterize that decision as inconsistent.

However, suppose that the lottery ticket purchaser wins the $100 prize--does that result affect our appraisal of his action as inconsistent? Not at all, because the judgment of inconsistency was based on the nature of his decision-making process, not on the ultimate outcome. In this case we would still regard his decision as inconsistent, but speak of the outcome of that decision as fortunate. We thus should describe decisions as "good" decisions if they are based on a logical evaluation of the information available in assigning probabilities and values, and if they are consistent with the goals of the organization for which the decision-maker is an agent. We should describe an outcome as "good" or fortunate if it represents a situation highly valued by the organization. Thus good decisions can produce bad outcomes, and bad decisions can produce good outcomes. Some people play wisely and lose; others play foolishly and win. We must be careful to reward the logical, wise, and farsighted decision-maker even though he occasionally will incur bad outcomes and to refrain from rewarding decision-makers whose success is due to chance. The other course is to place the decisions of the organization in the hands of individuals who are "lucky"--a course that would not be too costly if we could tell when the run of luck was going to end.

However, nothing we have said should be construed to mean that the only good decisions come from a formal, mathematical decision procedure. Although such a procedure is a real help for most of us in our desire to be logical and consistent, we may know some individuals who are always capable of arriving intuitively at the same decision the rest of us reach after much labor. Such individuals should be highly valued, for they are rare. The people in which we must not place our confidence are those who make decisions without either deep insight or a formal procedure. The point is this: Suppose you learned that a man with a string of 10 successful decisions to his credit had made those decisions by flipping a coin before even considering the merits of alternate proposals.

Could you be assured of his future success? (According to the joke, you should buy his coin and then fire him.)

THE TOOLS OF DECISION

The tools of decision-making are, as we have said, the theory of probability and the theory of utility. Each of these requires and is worthy of long and serious study by the decision-maker. However, there have also arisen certain techniques based on these theories that aid us in visualizing decision problems.

The most important and simple of these techniques is called the decision tree. A typical decision tree is shown in Figure 1 for a problem we shall call the Judge's Problem. This is the problem of a judge who must decide whether to convict or free a man who may be guilty of a crime. We assume that no jury is involved and that the judge is therefore fully responsible for the decision. The decision tree is merely a pictorial representation of the sequential steps in the decision problem. In this case the essential decision is to convict or to free and we have branches with these labels emanating from a node indicated with an X. The node is called a decision node because the choice of which branch to take is the province of the decision-maker. After the decision is made and regardless of which way it is made we come upon a node from which emanate two branches marked "guilty" and "innocent" corresponding to the possible states of the defendant.

The defendant knows whether or not he is guilty, but of course the judge can only assign a probability to the defendant's guilt. In this Figure, we indicate that the judge has assigned a probability p to the defendant's being guilty on the basis of the evidence presented and therefore has assigned a probability $1-p$ to his being innocent. The node from which emanate the guilty and innocent branches is therefore a chance node. The branch that will actually be taken at this point is governed not by the decision-maker but, as far as the decision-maker is concerned, by chance.

At the tips of the tree are recorded the economic values to society of having each of the four possible outcomes represented by the tree, i.e., convicted and guilty; convicted and innocent; free and guilty; free and innocent. The number A is the contribution to society of one year's income which we shall take as $7,000; B is the cost of keeping a man in prison for one year which we shall take as $2,000. The amount C is the cost of imprisoning an innocent man, a very high cost for a society respecting justice, but one which must be evaluated since it

Figure 1

Decision Tree for the Judge's Problem

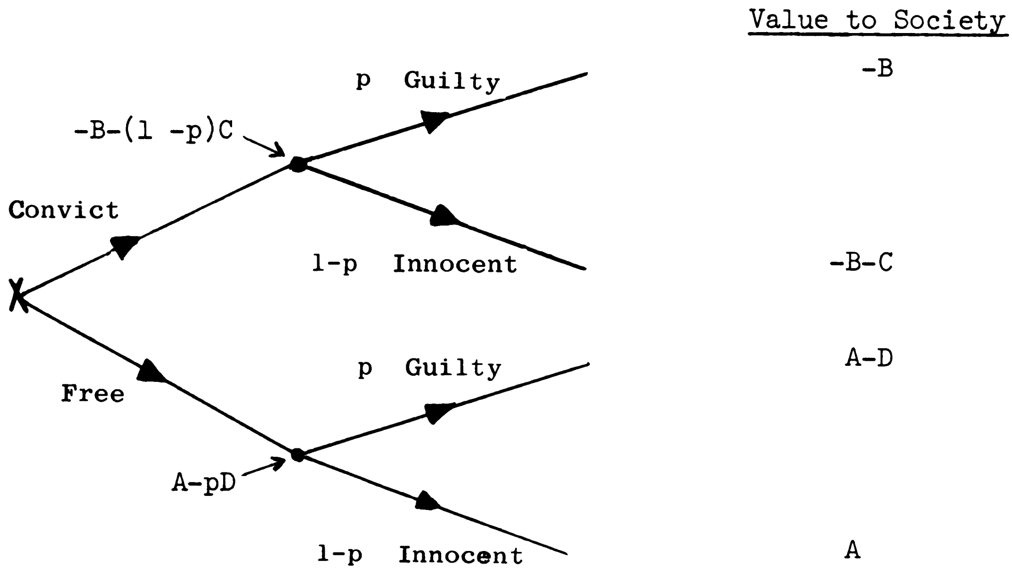

A: Contribution to society of one year's income

B: Cost of keeping a man in prison for one year

C: Cost of imprisoning an innocent man

D: Cost of letting a guilty man go free, perhaps the expected cost of additional crime

will sooner or later be faced. Let us assign this cost as $100,000. The quantity D is the cost of letting a guilty man go free. Perhaps this is the expected cost of additional crimes he may commit during the one year period. We shall say that this is $10,000. Thus, if the man is convicted and is guilty, the cost to society is the cost of imprisoning him. If he is convicted and is innocent, to the cost of imprisoning him must be added the cost of imprisoning an innocent man. If the man is freed and is guilty, then society receives his income A but loses through his crimes an amount D. If the man is freed and is innocent, then he makes his normal contribution of A to the society.

Now that the values and probabilities have been assigned, the problem for the judge is to establish the decision criterion. Let us suppose that he chooses to follow the route that maximizes the expected value to society. We compute the expected value to society under a given alternative by multiplying the value to society for each outcome under that alternative by the probability of that outcome and summing over all outcomes. Thus if the defendant is convicted, the expected value to society is given by $-B - (1-p)\,C$, while if he is freed, it is given $A - pD$. The convict alternative will have a higher expected value if

$$-B - (1 - p)\, C > A - pD \tag{1}$$

or if

$$p > \frac{A + B + C}{C + D} \tag{2}$$

for the numbers we have assigned the criterion on p for conviction is

$$p > \frac{109}{110} \tag{3}$$

Thus in terms of the expected value criterion and the costs that we have assigned the judge should convict the defendant if he is more than approximately 99% sure that the defendant is guilty. Under these very special assumptions, we have therefore established a quantitative meaning for the term "reasonable doubt."

This example could be made considerable more realistic by separating sentencing from conviction, for example. We could then build a more detailed model for the defendant's future behavior in order to determine what an appropriate sentence might be. However, rather than to continue further with this example, let us consider a problem from an industrial context.

Figure 2 shows a decision tree for a company trying to decide whether to introduce a new product. The decision node shows that there are three alternatives: produce the new product, test market the new product or forget about the new product. If the new product is produced, then it will either be a success or failure and we could spell out those consequences in more detail. If the decision is to forget about the new product, then there will probably be no future consequences except possible loss of future profit. The decision of test market, however, is especially interesting. The result of the test market alternative will

Figure 2

Decision Tree for a Company Problem

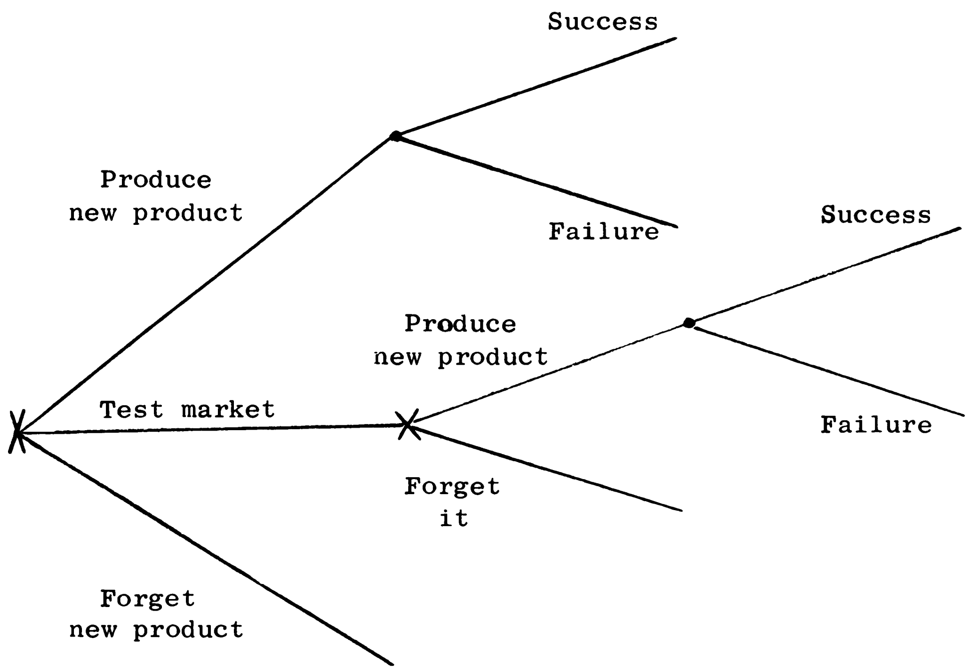

have to be either a decision to produce or to forget about the product when the test market results are available. Of course, there is still another alternative and that is to continue test marketing but for simplicity we have not included it in the figure. When the tree is drawn, then management must assign the probabilities for success and failure contingent on each of the alternatives and then establish the values to be assigned to each outcome. We should notice that the cost of test marketing must be included in assigning the value to the test marketing alternative. Of course the reason that this test marketing is conducted is that it is hoped that the ultimate probabilities of failure and success will be more clearly indicated by the results from the test market. It is consideration of such alternatives as test marketing that decision theory can play an especially valuable role. Decision theory can tell us first of all whether test marketing is worthwhile, whether the information that could be gained from it is expected to be as valuable as its cost. If test marketing is profitable in principle, then decision theory will tell us how much test marketing should be done and can even aid in the selection of test markets and in establishing the extent of the testing in each market.

Thus the evaluation of experimental programs in test marketing is merely another alternative to be evaluated in a decision tree. The techniques for performing this evaluation are straightforward and readily adaptable to automatic computation. When we do employ automatic computation, the trees that we consider can be very large, with many alternatives considered and many outcomes evaluated. Thus we see that though simple in concept the decision tree is a very valuable tool in decision-making.

But is the construction and solution of trees the only contribution of decision theory? The answer is no. The decision tree is simply the most easily seen part of an iceberg of decision theory whose larger part is based upon more complex probabilistic structures. We shall illustrate this with the problem space shown in Figure 3. We can characterize a decision problem by the type of probabilistic structure that underlies it. One dimension of the problem is the degree of uncertainty in it. How many of its elements are probabilistic, i.e., how many must have probabilities assigned to them? How many are deterministic, i.e., known with great precision? Another dimension of the problem is the number of variables, a number that may range from one to several hundred. Some of the variables may be known deterministically, others only in a probabilistic sense. The final dimension of the problem space is time. Is the problem a static one, like perhaps the travel problem we discussed earlier where the decision once taken will have implications only into the very near future, or is the problem a dynamic one where the effects of the present decision must extend over several years? We note that if time is important in a decision, then we must use the principles of discounting future income and costs to reflect the true economic nature of the problem.

As we might expect, the simplest decision problem would be that in which there was one variable known deterministically and for which the time factor was unimportant. Most such problems would in fact be trivial. However, as we move away from the origin in this problem space, the problems become increasingly difficult. As the number of variables increases, as more of them must be described probabilistically, as the effect of time becomes increasingly important, we arrive at decision problems that are difficult not only from the point of view of computation, but even from the point of view of formulation.

We can identify each of the corners of the problem space with the fields of mathematics. We have already dismissed corner 1 as trivial. Corner 2, the deterministic, dynamic, one-variable problem, is treated in college courses on differential equations. Corner 3, the probabilistic, static, one-variable problem, is covered by the elementary probability theory for individual random variables. Corner 4, the probabilistic,

Figure 3

The Problem Space

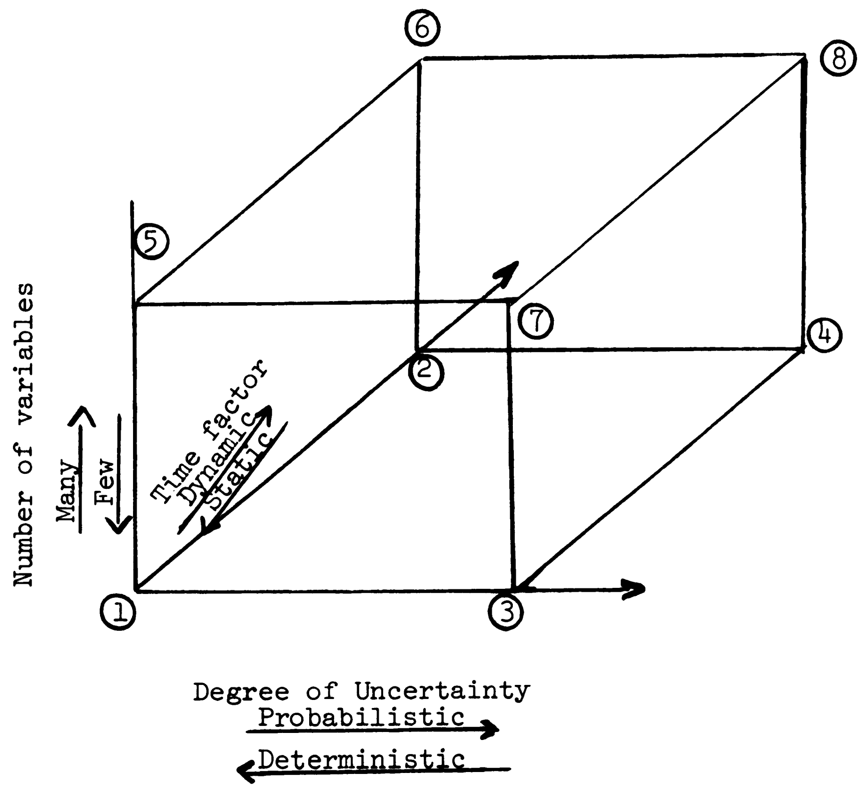

dynamic, one-variable problem, is the province of the theory of stochastic processes. Corner 5, the deterministic, static, many-variable problem, is treated by the theory of matrices and of multi-variable calculus in general. Corner 6, the deterministic, dynamic, many-variable problem, is the primary concern of the modern theory of control practiced by control engineers. Corner 7 is the probabilistic, static, many-variable problem that we study as the theory of joint distributions in probability theory of joint distributions in probability theory. Corner 8, the most complicated corner of all, is the probabilistic, dynamic, many-variable problem for which the Markov process and its relatives are helpful models.

Thus we see that the technology exists for analyzing problems in almost any area of the problem space. The depth of the coverage varies, but the basic structure is there. To get an idea of the kind of models available, let us consider the simple Markov process shown in Figure 4. The process has two states, state 1 and state 2, that it may occupy. From each state it may make a transition back to that state or a transition to the other state. The arrows in the Figure indicate the possible transitions. The numbers appended to the arrows are the probabilities that if the process is in a certain state, it will make the transition

Figure 4

A Markov Process

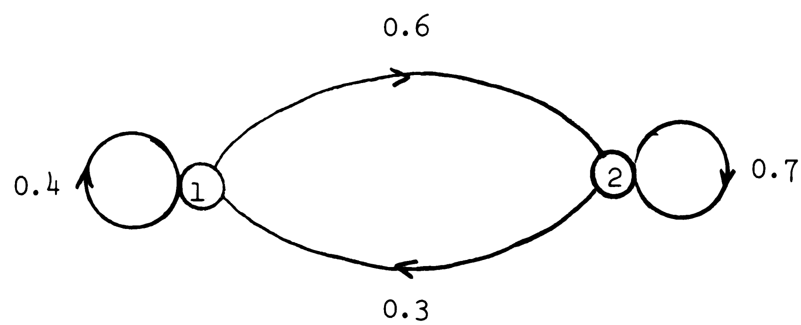

State 1: Customer bought brand A last

State 2: Customer bought brand B last

indicated by the arrow. Thus when the process is in state 1 we say that there is a 0.4 probability of its returning to state 1 and 0.6 probability of its moving to state 2 on its next transition. Similarly, when the process is in state 2, we say there is a 0.7 probability of its returning to state 2 and 0.3 probability of its making a transition to state 1.

We can consider this Markov process as a model of purchasing behavior. We let state 1 represent the state of a customer who last bought brand A and state 2 that of a customer who last bought brand B. Thus the Markov model says that a customer who bought brand B last is more likely to follow that purchase with the purchase of Brand B than he is with a purchase of brand A. The reverse is true for a customer who bought brand A last: he has a 0.6 probability of buying brand B on his next purchase. From the theory of Markov processes we can calculate the probability that if a customer bought brand A on his last purchase he will also buy brand A three purchases from now, or five, or ten, or one hundred. A moment's reflection makes us realize that if we let this process operate for a long enough time, knowledge of a purchase a long time in the past should have no influence on the probabilities of present purchase. This is in fact the case and the theory shows that if a customer is operating as the process indicates, then he is twice as likely to be in state 2 at a time far in the future as he is to be in state 1. In other works, we predict that he will buy brand A twice as often as brand B in the long run.

Many other interesting questions may be asked of Markov processes. Further, the processes we consider can be much larger with perhaps fifty or one hundred states. We can expand the model by allowing each transition to have a time duration drawn from a probability distribution, but perhaps the most valuable flexibility available is in allowing monetary rewards as well as probabilities to be associated with each transition. Then we can talk of profit in a Markov process model. Furthermore, we can superimpose upon the model a decision structure that allows us to calculate the way of operating the system that will be most profitable in the long run. In particular this Markovian system is an explicit model for taking into account the effect of present decisions upon the future, an effect whose importance we stressed earlier.

But the Markov process is only one of many probabilistic models that aid us in solving the decision problems based on the problem space of Figure 3. In attacking decision problems there is no substitute for a fundamental knowledge of the underlying probabilistic structure.

CONCLUSION

We have now seen how a theory of decision-making can and must be based on the theory of probability. The functions of the decision-maker are thus to assign probabilities, assess values, and establish a decision criterion. When this has been done, the solution of the problem is an exercise in logic and therefore the province of the digital computer, if necessary.

Perhaps the best way of ending is with the statement of J. Clerk Maxwell, the father of electro-magnetic theory: "The true logic for this world is the calculus of probabilities which takes account of the magnitude of the probability which is, or ought to be, in a reasonable man's mind."

8

AN ASSESSMENT OF DECISION ANALYSIS

Ronald A. Howard

Department of Engineering-Economic Systems
Stanford University

Reprinted from Operations Research, Volume 28, Issue 1,
January-February 1980.

An Assessment of Decision Analysis

RONALD A. HOWARD
Stanford University, Stanford, California
(Received July 1979; accepted October 1979)

Making decisions is what you do when you don't know what to do. Decision analysis is a process that enhances effective decision making by providing for both logical, systematic analysis and imaginative creativity. The procedure permits representing the decision-maker's information and preferences concerning the uncertain, complex, and dynamic features of the decision problem. As decision analysis has become more accepted and influential the ethical responsibility of decision analysts has increased. Analysts must be sensitive to assuming improper roles of advocacy and to participating in analyses whose means or ends are ethically repugnant. Criticisms of decision analysis are examined at three levels. Application criticisms question how much decision analysis improves actual decision making. Conceptual criticisms argue that the decomposition and recomposition of the decision analysis process may lend to a misshapen framing of the problem or to a suppression of "soft" or "fragile" considerations. Criticisms at the level of principle grant the effectiveness and comprehensiveness of decision analysis but express fear that the process may legitimize decisions otherwise questionable because of their end-state value system or their anthropocentric focus. Decision analysis is the most effective decision methodology yet advanced. Sensitivity to practical and ethical concerns about its use can only increase its effectiveness.

IN THE 10 YEARS since the first special issue on decision analysis (Howard [10]), the profession has grown considerably in number of applications and professionals. With successful establishment of the profession there is an obligation to examine the advantages of its use and the possibilities of its misuse to avoid either a limitation of future growth through the misunderstanding of potential users or worrisome misapplication through the insensitivity of practitioners.

My purpose in this essay is to present views on human decision making, on the nature of decision analysis, and on the usefulness of decision analysis. I shall then examine some practical and ethical issues involved in using decision analysis as public policy analysis.

1. HUMAN DECISION MAKING

To place decision making in perspective we have to return to a controversy among the ancient Greeks (Capra [2]). Heraclites of Ephesus

Operations Research
Vol. 28, No. 1, January-February 1980

0030-364X/80/2801-0004 $01.25

believed in a world of perpetual change arising from the dynamic interaction of opposites. These pairs of opposites were the unity which contained and transcended all opposing forces. Parmenides of Elea believed in a divine principle standing above all gods and men, a belief that ultimately led to the separation of spirit and matter, to the separation of body and soul, and to the distinction between subject and object. The philosophy of Parmenides culminated in the development of Western thought and science while that of Heraclites is remarkably similar to the Eastern world views of Hindus, Buddhists, or Taoists. Within the present century we have seen the advance of physics into the realm of quantum theory question the subject-object world and produce apparent paradoxes that can only be resolved by world-views (Dewitt and Graham [3]) that are remarkably close to those of Heraclites.

The idea of a "decision" is a quintessentially Western idea, an act of hubris to a believer in Eastern philosophy and a joke to the enlightened. (Can you imagine Buddha or Lao-Tzu making a decision?) However, we in the West are captives of our culture and so we are usually strong believers in the idea of making decisions. Yet many of us have had the experience of knowing that certain actions are beyond decision, particularly actions concerning love, the infinite resource which need never be allocated. Here, perhaps, we perceive the world with the undifferentiated gaze of the East.

There remain for most of us many situations where we don't "know" what to do in this sense, situations where we must allocate resources and balance in some way the pros and cons of each alternative allocation. This is what I call the realm of decision making. (I tell my class: decision making is what you do when you don't know what to do.)

There are many approaches one can use in decision making. One is intuitive or holistic. The Gestalt of information, sensations, and impressions gathered by the brain somehow results in the individual choosing a course of action. Another approach is analytic or rational. Here the situation is dissected into its features, and these are then evaluated by some logical process to arrive at a decision. Recent brain research indicates that the right hemisphere of the brain is heavily involved in the first process while the left hemisphere is predominant in the second. I shall discuss these processes in more detail in what follows, but I can say now that decision analysis as a formal methodology is a candidate (and some would say "the" candidate) for a logical procedure.

My personal view is that the analytic and intuitive capabilities of the mind are synergistic and not destructively competitive. For example, there is to my knowledge no synthesis procedure for a color television set or a jet airliner. The creation of each requires imaginative solutions to a variety of problems. However, these problems are often revealed and solutions suggested by the extensive analysis engineers perform in testing their designs. Similarly, I consider decision analysis not simply a logical procedure but in addition an essentially artistic process for achieving the creativity in decision making that can only result when we use all our faculties.

I have focused on individual decision making because in a most important sense all decisions are individual decisions. Individuals making decisions as agents for others must conform to the agreement they have made in accepting the responsibility to act as agent. This category includes all those who act on behalf of organizations. The individual who as principal devolves some decision making authority upon an agent is making a decision in that devolution. Thus, whether as principal or agent, everyone is exercising individual decision making. This view leaves no room for group decision making except that of individuals acting collectively in accordance with an agreement.

2. DECISION ANALYSIS

Decision analysis is the profession concerned with helping individuals make decisions. The profession consists of a theoretical paradigm for decision making and a body of practical experience for using this paradigm to illuminate the decision problem for the decision-maker. Central to the paradigm is the decomposition of a possibly uncertain, complex, and dynamic decision problem into the choices, information, and preferences of the decision-maker: the *decision set* of the decision-maker. If the decision-maker wishes to follow certain normative rules for decision making, logic applied to the decision set reveals the preferred course of action. The process is thus one of decomposition and recomposition with considerable emphasis on the insight to be gathered in both procedures.

I must emphasize the normative nature of the process. Decision analysis will typically be a very poor description of the way people make decisions. In fact, its power derives from the fact that its procedures are not automatically followed: decision analysis can improve upon natural decision making only because our natural decision processes are so deficient when we encounter novel decision problems, as we shall see.

I have described elsewhere the detailed procedures of decision analysis (Howard [11, 12, 17], Matheson and Howard [21]). The decision analysis cycle separates problems into deterministic, probabilistic, and informational phases. Assessment and modeling procedures form choices, information, and preferences into the decision set. The concepts of clairvoyance and wizardry permit calculating what it would be worth to know and to change what is now uncertain and uncontrollable.

The key to decision analysis is the construction of the decision set. To some this is a matter mainly of assessing probabilities; I call these people the "direct assessment" school. The decision-maker is asked to create alternatives (choices) and to provide a set of variables, the outcome vector, on which the outcome will be judged. Then he assesses probabilities on the outcome vector given the alternative (information), and finally a utility function on the outcome vector (preferences). By the normative axioms the preferred alternative is the one with the highest expected utility.

To others creation of the decision set is a more highly structured activity. This school, of which I count myself a member, is the "modeling school." It is based on the premise that few individuals are either comfortable or effective in representing their information and preferences in the form described above except in the simplest of decision situations. Members of the modeling school construct a more explicit representation of the decision set. Relations between outcomes and alternatives are captured using the structural information possessed by the decision-maker or his delegates. The models may be modest or extensive depending on the nature of the decision problem. Likewise, the representation of preference is usually divided into placing values on certain outcomes by means of a value function and then encoding risk preference on an appropriate numeraire. The result of the process, which we may regard as decision engineering in the same sense as electrical or mechanical engineering, is an extrapersonal representation of the decision problem—a representation that can be prodded, tested, and compared with other representations or within itself. The representation usually proceeds through a number of stages, such as pilot, prototype, and production just as do the designs of physical devices [12].

This extrapersonal representation has several advantages. It allows experts to contribute in their field of expertise: lawyers on legal aspects, metallurgists on material technology, and salesmen on marketing prospects. The representation thus serves as a vehicle for focusing all the information of experts that the decision-maker may wish to bring to bear on the problem while leaving the decision-maker free to accept, reject, or modify any of this information and to establish preferences. In the many cases where the decision-maker is acting as the agent of others (stockholders, for example) the extrapersonal model provides a communication tool for demonstrating that the decision-maker is functioning according to his agreement with them.

Certain issues seem to arise whenever a decision analysis is performed. One is the question of whether the analysis is "objective" since it uses "subjective" probabilities. This is mainly a semantic problem that can be avoided by using simply the word "probability" instead of "subjective probability." Since decision analysts believe that the only meaning of a probability is a particular individual's quantitative description of uncertainty, no modifier is necessary. The confusion arises because many people have been led to believe that only "objective" results are valuable, a belief that can be traced back to our pre-twentieth century scientific views since it relies on the subject-object dichotomy. What many people mean when they say "objective" is "impartial"; that is, not influenced by the prejudices of the analyst. The decision-maker must be convinced of this if he is to find the analysis credible.

An initial concern of decision-makers is often whether or not decision analysis is a discipline rich enough to encompass all the factors that may be important in the decision. For example, they may ask whether the decision analyst can treat "intangibles," by which is meant such very palpable outcomes as pain, disability, etc. The answer is "yes." In principle, any outcome can be valued and in practice many have, including scarring (Ginsberg [8]) and death (Howard [13]). The concern about comprehensiveness is misplaced because the real problem in decision analysis is not making analyses complicated enough to be comprehensive, but rather keeping them simple enough to be affordable and useful. We occasionally encounter executives who say that they tried decision analysis and that "it didn't work out." Further questioning reveals that someone who had had a course in decision trees attempted to draw a tree for the problem and got lost in his own jungle. The problem is not creating complexity but retaining informative simplicity.

There is sometimes confusion about the professional role played by the decision analyst. The decision analyst is an elicitor of information and preferences, an engineer of logical models, and an evaluator of alternatives. He is *not*, except by chance, an expert in the field of the decision. He is skilled in constructing the decision set using his imaginary and colorful friends, the clairvoyant (who knows all and who helps with defining variables and events unambiguously) and the wizard (who can do all and who helps with value assignment), but the information and preferences in the decision set must come from the decision-maker and his delegates.

Many decision analysts learn the language of applied fields like electric power generation or polymer chemistry. However, when the analyst feels that he can, by himself, specify structure or, even more significantly assess probabilities in a substantive field, then he has moved beyond the role of decision analyst to one of expert. Since a prime virtue of the decision analyst should be his detachment with respect to the alternative chosen it follows that he must not contend with the decision-maker's experts in any attempt to replace their information with his.

The decision analysis process is not static but iterative and interactive. Although we speak of assessing probabilities or preferences the process is better described as the formation, encoding, and verification of these quantities. Verification means not merely pointing out the consequences of the process in specific situations and receiving confirmation but also presenting the generic properties of the resulting assessment to make sure that they are understood and accepted. As a simple example, a decision-maker who showed a region of risk preference in his risk attitude should be informed of the practical implications of this preference.

The overall aim of decision analysis is insight, not numbers. If the decision-maker does not feel that the analysis has captured his knowledge and concerns and that it has produced a course of action he believes in then the decision analyst has failed. But this is rarely the case. In a recent study the decision analyst presented his final conclusions to the entrepreneur who had hired him. At the conclusion of the presentation the decision analyst asked about the amount of written reporting that would be required. The entrepreneur replied, "I believe the results of the analysis and I am going to act in accordance with the recommendation. Why should I pay more for a report?"

3. THE USEFULNESS OF DECISION ANALYSIS

Upon learning about decision analysis some will say, and many have, "Why should I bother with decision analysis. I make excellent decisions anyhow." Perhaps this is so. However, one could ask how the quality of a decision is measured in the absence of decision analysis, the field which enabled the quality of a decision to be defined as logically distinct from the quality of the outcome that follows it. But leaving this point aside, are people good natural decision-makers? There is now considerable evidence that they are not. To quote Slovic, Fischhoff and Lichtenstein [25]:

> The major advance in descriptive research over the last five years has been the discovery that people systematically violate the principles of rational decision-making when judging probabilities, making predictions, or otherwise attempting to cope with probabilistic tasks. Biases in judgments of uncertain events are often large and difficult to eliminate. The source of these biases can be traced to various heuristics or mental strategies that people use to process information . . . In the final discussion, a strong case is made that judgmental biases affect important decisions in the real world; numerous examples are provided.

These conclusions, based on the pioneering work of Tversky and Kahneman [42] as well as on the findings of numerous other psychologists, seem to hold quite generally as an assessment of how humans behave in probabilistic and decision making tasks.

In a study on the effects of stress on decision making Janis and Mann [18] show how the setting of the problem changes the kinds of behavior exhibited. Their findings are briefly summarized as follows. If the problem is seen as unimportant individuals exhibit either "unconflicted adherence"—complacent continuation of present behavior—or "unconflicted change"—uncritical adoption of a new course of action. If the individuals now perceive that risks are involved, they move to "defensive avoidance"—procrastination, shifting responsibility, and selective inattention to corrective information. If they, in addition, perceive hope of finding a better solution they exhibit "hypervigilance" or panic, characterized by frantic search and impulsive adoption of proposals on superficial grounds. Only if they perceive that there is ample time available to make the decision, do they exhibit "vigilance"—painstaking search for relevant information, unbiased assimilation of relevant information, and careful appraisal of alternatives. What Janis and Mann call "vigilance" I would call decision analysis when the decision problem under consideration is worthy of professional assistance. Thus, the way I view the results of Janis and Mann is that if an individual with an important responsibility to act for others does not use a procedure essentially equivalent to decision analysis when making a major decision he is likely to exhibit one of the pathologies of decision making described; the only remaining question is which one. I have made these mistakes enough myself to recognize the accuracy of the Janis-Mann analysis.

Some critics of the usefulness of decision analysis strike at the heart of the procedure, at the idea that a decision problem can be divided into its components and then recomposed. Critics would call this "reductionism" and contrast it to their "holistic" view.

An extremely articulate critic, whose views we shall examine at some length, is Lawrence H. Tribe. (Though some of Professor Tribe's views have changed over time, I have selected quotes from his work over a period of years because he has stated so well positions that one hears expressed continually.) As an analogy to describe his concern about reductionism Tribe [31] considers art:

> To offer a crude but instructive analogy, the comparison of a particular painting by Rembrandt with one by Picasso (to help decide, for example, whether it would be desirable to sell one in order to buy the other) in terms true to the objectivist ideal might proceed first by disregarding the history of each work (the "process" reduction), in order to focus exclusively on what appears on the canvas; and second by considering each work (the "substance" reduction) as just so much paint of various specifiable colors, in order to focus on features that can be impersonally compared (e.g., the Picasso might contain more of certain pigments than the Rembrandt). Such "structural" features as balance, movement, composition, and the like would be left out of account; for how could one "objectively" compare or even "analyze" them.

I would not say that there exists no analyst foolish enough to carry out the analysis that Tribe describes, but they are rare. To suggest that what gives value to the Rembrandt is its spectrum is to suggest that one would value a meal solely by its composition in protein, fat, and carbohydrate. The question is not whether one can perform foolish analyses—that is conceded; but, rather, whether one can gain insight from a proper analysis, and here I submit that the answer is "yes." For example, in the painting example, the owner would do well to consider the extent to which the value of ownership is intrinsic in the possession and enjoyment of the painting as opposed to its investment value.

The reductionist criticism lies at the center of perpetual controversies between artists and scientists. The circuit diagram of a television set is seen by the engineer as a more fundamental description of the set than its physical embodiment in terms of circuit boards, wires, picture tube, etc., but the artist may see it as a graphic composition on a piece of paper and obviously not a television set. Those who criticize analysis are more concerned by what may be lost in the process than by what is gained. The trick, of course, is to keep the gains without incurring the losses.

Some people believe that decision-makers are effective because they learn on the job, and that, in fact, management science is ruining the manager (Levitt [20]):

> Still, as a corporation gets better managed and more concerned with the quality and practice of management itself, its top people develop a powerful propensity to manage differently. They are encouraged in this by a rapidly expanding retinue of eager sycophants, equipped with new "scientific" tools and decision-making models, who promise to free the manager from the inescapable uncertainties, risks, and traumas of running an enterprise. "Experts," trained to the teeth in the techniques (but not necessarily the practice) of management, are enlisted to do even better what people of native shrewdness, sound good sense, and abundant energy did quite beautifully before.

Some of this criticism may be due to defects in application and to the limited perspective of some areas of management science, however, there are those who question whether decision analysis itself is useful to decision-makers. Dreyfus and Dreyfus [4] claim that attempts of a decision analyst to improve decision making are more likely to hurt than to help:

> While the formal model has the attractive feature, desired by advocates of "scientific decision-making" that it lays bare and arguable the complete explanation of a decision, it in no way represents the actual process through which expert planners decide. Since the similarity-based process actually used by experienced human beings ultimately leads to better performance in all areas than does the formal approach often practiced by beginners, decision-making based on proven expertise should neither be replaced by formal models nor should proven experts feel any obligation to explain their decisions in that way.

There is much here to comment on. First, few would argue with the statement "the formal model . . . in no way represents the actual process through which expert planners decide"; in fact, that is exactly what psychologists have confirmed. However, the claim "the similarity-based process actually used by experienced human beings ultimately leads to better performance in all areas than does the formal approach often practiced by beginners" is a statement with which I can heartily disagree when it applies to experienced decision-makers making difficult decisions. There is little evidence adduced by Dreyfus and Dreyfus, or anyone else for that matter, to support the claim. In fact, there is considerable evidence to refute it. Eddy [6] has studied extensively the quality of actual medical decision making. He has found not only that doctors make gross probabilistic errors, such as mistaking a conditional probability for its converse, but also that these errors lead to serious mistakes in selecting policies for medical treatment.

The "similarity-based" process seems close to the representativeness heuristic described by Tversky and Kahneman [42]. This heuristic, which may be useful in situations with little uncertainty, can lead to serious error in probabilistic settings. For example, in coin-tossing with heads represented by H and tails by T, people often regard the sequence $H\ T\ H\ T\ T\ H$ as more probable than $H\ H\ H\ T\ T\ T$ because it is "more random." The reference cites many examples of this behavior.

Why then do the Dreyfuses think so highly of "similarity-based" processes? Primarily, it appears, because the examples they have considered are those with little uncertainty, with repetitive opportunities to practice, and with immediate feedback of results. "The chess player has a 'feel for the game,' the language learner becomes fluent, the pilot stops feeling that he is flying the plane and simply feels he is flying" [5]. We have all experienced this type of learning—it is characterized by a high predictability of outcome given our decision. We could have added bicycle riding or playing tennis.

People behave quite differently in probabilistic situations. (See Slovic and Fischhoff [27] and Slovic, Fischhoff and Lichtenstein [26, 28].) As the degree of uncertainty goes up experimental subjects begin to form false hypotheses and to retain them in the face of contrary evidence. It is a case of "the burned cat fears the hot stove—and the cold one, too." Perhaps this is the reason for the growth of superstition in our species. One could easily believe that human beings have very little inherent ability to handle uncertainty, that it is a blind spot just like our inability to sense radioactivity. People seem to have no intuitive idea of how to update their beliefs in the face of new evidence or of how the size of an experiment affects the inference that may be drawn from it. It is a source of wonder among lay people that the relatively small samples considered by TV rating services, pollsters, or the census have any value.

Perhaps the pilot example cited above is the most instructive. One of the lessons learned by all pilots when flying in bad weather is "trust your instruments." The pilot in the cloud may "feel" he is upside-down, or whirling in a spin, but the instruments will show the actual flying condition. No matter how experienced a pilot becomes he still uses his instruments in bad weather—he never outgrows them. It would be a mark of foolhardiness, and not maturity, for a pilot to fly in bad weather without instruments. Where is the "similarity-based" process here?

I believe that in dealing with uncertainty the human being needs an instrument—probability theory, and that he will never be able to perform well in an uncertain environment without his instrument. Once a student approached me after a probability class and said, "I can see that this is a subtle subject and that my intuition is lousy, but won't I have developed an excellent intuition about probability when I have completed the course?" I had to answer "no," and assure him furthermore that such a facility still eluded me after teaching probability for many years. In fact, one of the highlights of the course is to show how many great minds of history became foggy when they encountered probability questions. For example, a book published in 1686, bearing Isaac Newton's imprimatur, advanced the idea that the age of the purchaser should have no effect on the price of a life annuity (Hacking [9]).

Of course the representation of uncertainty is only one feature of decision making, even though it is usually the most perplexing feature. When we ask about the ability of people to make choices in uncertain situations we find pathologies of preference as well as of inference behavior (Kahneman and Tversky [19]). For example, people tend to be risk-preferring for losses and risk-averse for gains even when they are highly unsure about where the zero point is. Thus formal methods are not the dangers that the Dreyfuses foresee but rather the only hope for compensating for the defects of the human mechanism when it faces highly uncertain decision problems. To use a less effective procedure in problems of great moment would be the real danger.

Someone may point out that the studies of the inadequacy of human decision making were performed in laboratory settings, not in the executive suites of America. That is true, but if the natural decision making of executives is to be excellent then some magical change must come over them when they put on a three-piece suit and sit behind a desk. Pending a major study of *in situ* executive behavior each of us will have to base our assessments of executive decision making on our own observations. It has been my experience that the highest level executives are those most willing to face the shortcomings of their decision processes and to seek help in improving them.

As an illustration of how executives can benefit from consorting with decision analysts, I am reminded of a project several years ago where a firm of lease brokers worked with decision analysts in an attempt to improve bidding effectiveness. In this business the low bid wins. Our clients were very proud of their "aggressive, risk-taking, low-bid" attitude. We were able to show that their bidding strategy was, in fact, one that would be favored by a highly risk-averse firm because to such a firm the consequences of not winning the bid were of more concern than the low profits obtained given that they won it. It was quite a shock for these executives to reassess their image but they did so and profited from the experience.

Decision analysis is not immune to the human foibles that psychologists find in intuitive decision making. Particularly in the process of probability and risk assessment, the decision analyst must be sensitive to the heuristic biases and must develop a methodology and professional practice that minimizes their effect. The decision analysts in my acquaintance are, in fact, very sensitive to this issue; they go to great lengths to avoid creating artifacts in their analyses (Spetzler and von Holstein [30]). One would say that to the extent individuals are not excellent processors the analysts are building the instruments to assist them. The happy fact that lets us all relax somewhat is the practical rarity of analyses in which the recommendations are highly sensitive to small changes in either probability or risk attitude assessments.

The use, as opposed to the usefulness, of decision analysis will depend on the value that decision-makers place on the virtues and drawbacks of the procedure. One value is what we might call the "immersion" value, the value of any procedure that entails a careful scrutiny of the factors that influence the decision. The immersion value of decision analysis is high because of its aggregative nature: each question answered leads to another question asked. The questions all fit within one structure because the process is comprehensive—no factor influencing the decision need be omitted because of conceptual limitations. This complete investigation of a decision problem is particularly valuable because it is examinable: the elements are quantitative and explicit. The input, the model, and the results are checkable by any interested party. Because the representation can be exercised, sensitivities to each feature of the problem may be calculated for the insight they provide.

The drawbacks of the decision analysis process can be large. Foremost is the fact that the quality of the analysis depends critically on the quality of the decision analyst. In perhaps no other form of analysis is it so easy for the analyst to produce any result he likes by taking advantage of his knowledge of biases and his modeling choices. Even if the analyst is well-intentioned, he can produce very misleading results by sheer incompetence. I have seen more than one multiattribute study where the decision-maker was not apprised of the tradeoffs between outcome variables that the utility function implied. Some of these tradeoffs seemed very strange in view of the setting of the problem.

One could ask whether the difficulty of performing "good" decision analysis is not itself a criticism of decision analysis. The analogy I like here is brain surgery. Because effective brain surgery is difficult does not mean that there is anything wrong with it. I would no more expect a person with little training to complete an effective decision analysis than

I would expect him to perform a successful brain operation. A distinguished nuclear engineer, upon being exposed to decision analysis, said, "I see, it's easy. All you have to do is knock a few probabilities together." I said, "You're right. It's just like nuclear engineering. All you have to do is hook a few pipes together."

An important and fundamental difficulty of decision analysis is that it is expensive. The level of skills required assures that it will always be relatively expensive even in a world where computation is cheap. I find that decision analysts require 3–4 years of graduate education and at least 2 years of practical experience before they can be considered fully trained. Only a tiny fraction of decisions will ever be aided by professional decision analysis; we can hope only that among them will be the most important decisions.

Finally, every analysis is necessarily limited—it can treat only a limited number of possibilities no matter how large the number or how important the possibilities. This necessary limitation is not really confining unless we view decision analysis too narrowly. Those who see decision analysis as a creative medium in which to apply the totality of their consciousness will build on its strengths and transcend its weaknesses. The only caution to add to this statement is that at the moment the skilled, sensitive, creative decision analyst is even more rare than the skilled, sensitive, creative decision-maker.

Decision Analysis as Policy Analysis

The usefulness of decision analysis in making a wide variety of both private and public decisions has now been established. However, with this growth has come increased concern about the use of decision analysis as policy analysis to support governmental decisions. We shall explore first the practical and then the ethical concerns.

Practical Concerns

The practical concerns focus on whether it is possible to produce a useful policy analysis using decision analysis. Such concerns are well-illustrated by cost-risk-benefit analysis, which at its best is an attempt to apply decision analysis to social decisions. "The most general form of cost-benefit analysis is *decision analysis* in which the role of uncertainty, the subjective nature of costs and benefits, and the existence of alternative actions are made most explicit" (Fischhoff, Slovic and Lichtenstein [7]). The problems the analyst will face in assessing benefits are usually perceived as the most difficult, and they are difficult. But what about the equally difficult problems of assessing cost in a society that is more than one-third government? Every price in our society is affected by local, state, national, and even international regulations, laws, subsidies, duties, and taxes. In computing the investment for a new energy development, should one use the cost of domestic steel or perhaps the lower cost of imported "dumped" steel? In buying tires for the trucks, should the U.S. excise tax be included in the cost? If federal law requires that union wage rates be paid for construction, should these rates be used as the actual opportunity cost for labor in the calculation even when there are unemployed nonunion laborers of equal competence available? In fact, is cost not just as uncertain as benefit? Think about any cost benefit analysis you have ever seen and determine if it still makes sense after such effects are included.

Note that this difficulty does not arise in nongovernmental calculations. A company does not care in principle whether the higher prices it faces are due to a tax or a drought. But the government should care since it can control the tax.

The problem is even worse if the decision analysis concerns whether a given function should be performed by the government or by private enterprise. If one assumes that the government is virtually risk indifferent, that it can borrow at the lowest rates in the society, and that the tax payments by private enterprise are costs of private enterprise but not costs of government entities, then virtually every analysis will show that the government can carry out the same activity more cheaply. Of course, then we think about the example of the post office. There is apparently something about the difference between government and private undertakings that may well not have appeared in the model. Lest you think that this issue is hypothetical, a recent American Physical Society study of nuclear reprocessing (Several Editors [24]) was carried out under the assumptions stated above and arrived at the predicted conclusions.

At a more subtle level, Tribe [32] believes that reductionism is a major concern in policy analysis because of a tendency toward the "dwarfing of soft variables":

> Whenever [certain kinds of values] appear to be involved, at least potentially, in a given problem, one should recognize that the tecnhiques of policy analysis as currently conceived will tend either to filter them out of the investigation altogether or to treat them in ways inconsistent with their special character ... But the problem here goes deeper. It relates not merely to undervaluing certain factors but to *reducing entire problems to terms that misstate their underlying structure*, typically collapsing into the task of maximizing some simple quantity an enterprise whose ordering principle is not one of maximization at all [italics Tribe's].

In another paper, he continues this charge [33]:

> Thus, because policy-analytic techniques prove most powerful when the various dimensions of a question are reduced to a common denominator, or at least to smoothly exchangeable attributes, the continuing tendency that accompanies analytic techniques is to engage in such reduction whenever possible, with the

result not only that "soft" variables tend to be ignored or understated but also that *entire problems tend to be reduced to terms that misstate their underlying structure and ignore the 'global' features that give them their total character* [italics Tribe's].

This is indeed a serious gravamen, and one not lightly dismissed by saying that the problem arises only with incompetent analysts. The problem is least likely to arise in situations where the analyst is working with a decision-maker to assist him in making his own personal decisions or where he is acting under a clear agreement with a group. It is most likely to arise in governmental decision making and is thus one more reason why I personally feel uncomfortable about other people making decisions on my behalf and without my concurrence, regardless of the methodology they employ.

One particular form of the reductionist criticism that is of concern to Tribe in the policy area is the separation of information and preference [34]:

> To offer one concrete if limited illustration, I would focus on the frequently stressed tenet of decision theorists that one of the analyst's main functions is to help the decision-maker separate clearly (1) how he feels about various possible outcomes of his decision (this preference being a matter of personal value) from (2) his best assessment of the probability of each such outcome (this probability being a matter of impersonal fact).

Tribe points out in a footnote: "The 'correct' probability is thought to be a matter of impersonal fact even though any particular assessment of it will invariably be personal and subjective." Tribe goes on to say [35]:

> . . . convicting an innocent person should be deemed a worse outcome when the jury feels very unsure of the person's guilt (but chooses to convict anyway) than when the jury feels fully confident of guilt (but simply happens to be mistaken). What is being done to the accused in the two cases differs just as surely as kicking a child does from tripping over it, and the consequences for society of permitting each of these practices differs as well . . . Similarly, destroying a species of wildlife should probably be regarded as a worse outcome when it results from the disregard of a high known risk than when it results from the materialization of a highly unlikely contingency. The tradition in many legal systems of distinguishing among acts in terms of the mental state accompanying them (treating murder differently from manslaughter, for example) rests on this sort of proposition. Yet the objectivist's fact-value dichotomy, leading to an insistence on separating assessments of probability from the valuation of outcomes, tends to exclude this important dimension of human choice.

Tribe explains in a footnote, "I say 'tends to exclude' rather than 'excludes' because a sufficiently careful policy analyst could define the 'outcomes' in question so as to include information about the associated probability assessments. The objectivist perspective does not preclude such a step, but does make it far less likely." This footnote shifts the criticism somewhat to the level of application, but there remains the implication that analysis encourages overlooking important value features of the process.

Perhaps Tribe does not understand that there is no difficulty in assessing preference on the combination on decision and outcome—the value to be placed on intentionally driving a car at a person and killing him can be quite different from the value of killing a person in a traffic accident. Similarly, in judging the behavior of a jury, the decision to convict a person in the face of grave doubts about his guilt can be valued differently from the decision to convict him when there is little doubt. In fact, decision analysis should be the test of whether the jury made a good decision. Far from being a criticism of analysis, Tribe's discussion merely reveals to me the clarification that results when analysis is used to illuminate what can otherwise be a perplexing situation.

These practical criticisms do not seem to me to pose serious objections to decision analysis, per se. They do show, however, the difficulty of performing policy analysis using decision analysis and the necessity of considering decision analysis as a profession rather than as a technique.

4. ETHICAL CONCERNS

As the practical impact of decision analysis increases, so does my concern about the ethics of its use. Since the formalism of decision analysis is amoral, like arithmetic, any moral considerations must come from the people involved in the application. By analogy, suppose that a maker of fine rifles that he thought were used for target competition should find out one day that his customers were, in fact, assassins. To continue to make the rifles is to be the *de facto* accomplice of assassins, even though that is not his intent. The decision analyst, and the educator of decision analysts, faces the same question. If the decision analyst or his analysis becomes the means to an end that he finds morally unacceptable, should he not withhold his labor?

I have been asked why ethics is of more concern to decision analysts, or to management scientists in general, than it is to other members of our society. I believe that analysts (management scientists and operations researchers) as advisors on resource allocation are more likely than most professionals, and perhaps as likely as doctors and lawyers, to face important ethical questions. The scale of the problems they work on may make the consequences of their activities even more extensive than they are in the medical and legal cases. (Many will recall the ABM controversies of the last decade [43, 44].)

For the decision analyst the ethical question is even greater. According to Boulding [1], any decision is a process "affected with the ethical interest." The decision analyst holds himself out as being able to analyze the most complicated of decisions, including those affecting human life and the environment—I have personally been involved in several such studies (Howard, Matheson and North [14], Howard, Matheson and Owen [15], also [45, 46]). Particularly in social decisions, the public could well believe that when a "decision analysis" recommended a course of action the recommendation was "scientifically based" rather than a consequence of the information and preference inputs of the decision-maker. Such a belief would be sadly in error.

While decision analysts like to think that they are hired because of a desire for more systematic logic in making decisions, there is usually a more direct cause. I have recently been asking decision analysts how many of their studies were prompted by a belief in systematic analysis per se rather than by a desire of some party to the decision to advance either his own proposals or to defend them against attacks. A very large proportion of responses has been in favor of the self-interest hypothesis. While there is no reason why a useful decision analysis could not be promoted by an interested party, there is in these situations an increased potential for ethical dilemmas. One has only to participate in a few studies where the logical, systematic answer differed from the one the sponsor expected to see the strength of this potential.

Ethical problems could arise in business applications of decision analysis. Although I have never heard of a situation where the client was trying to use the decision analyst or his work in an attempt to break the law, there may be cases where the decision analyst begins to wonder whether the behavior of his client is in accordance with the agreement that the client has made with others, for example, the stockholders of the company. The decision analyst could also be faced with being asked to work on a problem where the nature of the problem is itself ethically unattractive; for example, a Catholic decision analyst confronting someone's abortion decision or a pacifist decision analyst facing a decision involving the manufacture of arms. Here the principle that no one should participate as the means to an end that he feels is ethically unacceptable will guide the analyst.

The most serious ethical problems arise when the analyst works in the public arena. No one who participates in public affairs, however remotely, can escape the moral responsibility for his actions. Here more than elsewhere the analyst will be tempted to slip over into the role of advocate of some position that he personally favors. To be an advocate while posing as analyst is, of course, professionally unethical even if the position advocated is morally excellent to the individual. The problem is that one is using morally reprehensible means, fraud, to achieve the end. Unless the person advocates fraud as a moral virtue he is in an ethical contradiction. While the end must justify the means in the sense that one chooses effective means to attain goals, one cannot ethically use a means that is inconsistent with the moral system that validates the end.

Is an activity ethical because it is the legal result of even a democratic process? For example, consider the distributional questions posed by cost-risk-benefit analysis. Typical practice is to compute the total benefits and total costs to society for each alternative and to recommend the alternative with the greatest difference. This method assumes that there is some mechanism to redistribute the "profit" so that everyone will be at least as well off as before. However, the past record of government programs shows that this ideal is almost never met. In the case of a taxation measure where property will be removed from some and given to others, I have never heard of a justification that it was a redistribution of social profit resulting from previous "cost-benefit" decisions. The question of social efficiency addressed by cost-risk-benefit analysis (inadequately, because of the earlier problems we mentioned) has nothing to do with the ethics of the situation. The primacy of the individual on which our system is based is inconsistent with the idea that you can hurt one to help many. Therefore, unless the redistribution is, in fact, going to take place, is not the adoption of an alternative that benefits some and hurts others not just a case of theft from those who have lost?

Other ethical considerations may be more subtle. In a paper on environmental law, Tribe [36] used the example of plastic trees planted in a Los Angeles median strip—the only trees that could survive in this environment:

> Consider again the plastic trees planted along a freeway's median strip by Los Angeles County officials. If the most sophisticated application of the techniques of policy analysis could unearth no human need which would, after appropriate "education," be better served by natural trees, then the environmental inquiry would be at an end. The natural trees, more costly and vulnerable than those made of plastic, would offer no increment of satisfaction to justify the added effort of planting and maintaining them.
>
> To insist on the superiority of natural trees in the teeth of a convincing demonstration that plastic ones would equally well serve human purposes may seem irrational. Yet the tendency to balk at the result of the analysis remains. There is a suspicion that some crucial perspective has been omitted from consideration, that the conclusion is as much a product of myopia as of logic.

We sense here that the focus of the attack has shifted. We indeed find [37]:

> . . . one must concede that there is nothing in the structure of the techniques themselves, or in the logical premises on which they rest, which inherently precludes their intelligent use by a public decision-maker in the service of these "intangible," or otherwise "fuzzy," concerns.

This does seem like a retreat on the soft variable issue and leads to an attack not on policy analysis, but on the ideology within which it is used:

Thus the distortion results not from a logical flaw in the techniques of policy analysis but rather from what I have elsewhere described as the ideological bias of the system in which such analysis is imbedded, a system that has come to treat the human will and its wants as the center around which reason as calculation must revolve [38].

In most areas of human endeavor—from performing a symphony to orchestrating a society—the processes and rules that constitute the enterprise and define the roles of its participants matter quite apart from any identifiable "end state" that is ultimately produced. Indeed, in many cases it is the process itself that matters *most* to those who take part in it [italics Tribe's] [39].

The only entities that can "count" in a calculus of end-maximization, whether utilitarian or contractarian, are those entities that possess their own systems of ends or at least the capacity to experience pleasure and pain, and nothing outside the private ends and pleasures of such beings can come to the rescue of a philosophy devoted solely to their pursuit [40].

Thus, Tribe criticizes not only the use of policy analysis to attain end-states, regardless of process, but also the anthropocentric value system it incorporates. The concern with a world beyond man takes us to the question of animal rights [41]:

What is crucial to recognize is that the human capacity for empathy and identification is not static; the very process of recognizing *rights* in those higher vertebrates with whom we can already empathize could well pave the way for still further extensions as we move upward along the spiral of moral evolution [italics Tribe's].

A careful view of Tribe's criticisms shows that they are not criticisms of decision analysis so much as criticisms of our social decision making processes regardless of the procedures employed. End-state versus process ethics and animal rights are issues that would exist even if decision analysis did not. The ethical danger for decision analysts lies in attempting to use decision analysis to overwhelm with technology what are really ethical problems.

5. THE FIRST DECISION ANALYST

There are those who believe that there is something "cold and inhuman" about rational analysis. I believe that to be human is to be reasoning as well as compassionate. My ideal here is Buddha:

Perhaps the most striking thing about him, to use the words of J. B. Pratt, was his combination of a cool head and a warm heart, a blend which shielded him from sentimentality on the one hand and indifference on the other. He was undoubtedly one of the great rationalists of all times, resembling in this respect no one as much as Socrates. Every problem that came his way was automatically subjected to the cold, analytical glare of his intellect. First, it would be dissected into its component parts, after which these would be reassembled in logical, architectonic order with their meaning and import laid bare (Smith [29]).

Perhaps Buddha was the first decision analyst.

6. SUMMARY AND CONCLUSION

The concept of a decision seems very natural to Western thought but is strange to Eastern philosophy. Almost everyone even in the West has experienced situations where he "knew" what to do, not in the sense of selecting one among many possible courses or action, nor in the sense of having a hunch, but rather in the conviction that this was the only "right action." Furthermore many people, upon reflecting on their lives, see that there were no "good" or "bad" outcomes, even though outcomes were perceived as such at the time. These outcomes were simply necessary lessons in the game of life.

However, many of us at all times, and most of us at some times, will see ourselves facing a decision. For example, the question of which car to buy will usually be seen as relatively important to the individual and not the sort of question evocative of inner knowledge. Most people will make this decision by some combination of intuition and logical procedure.

Decision analysis, as I have described it, is, as a formalism, a logical procedure for decision making. When decision analysis is practiced as an applied art the formalism interacts with the intuitive and creative faculties to provide understanding of the nature of the decision problem and therefore guidance in selecting a desirable course of action. I know of no other formal-artistic approach that has been so effective in guiding decision-makers.

We must realize, however, that one of the arts of the decision analyst is the art of knowing how much and what kind of decision analysis to do. The degree of analysis can range from making simple lists to constructing giant interactive computer models. To be effective decision analysis must be "appropriate": the extent of the analysis must be suitable to the means and ends of the decision-maker. Thus simple problems of little consequence should receive very modest analysis and complicated problems of great importance should be extensively studied. The question of whether the analysis was appropriate to the decision-maker and his problem is one that should always be raised in judging effectiveness.

When I think about possible reasons why apparently appropriate decision analyses might not be judged as effective now and in the future, I find that my greatest concern is the improper treatment of probabilistic dependencies. At every level from the academic teaching of probability theory to the most extensive applications, the issue of characterizing probabilistic dependence seems to be the source of the greatest errors. The belief in the unsinkability of the Titanic was a belief that the flooding of the various holds constituted a set of weakly dependent, if not independent, events. The analysis of "common-mode" failures in nuclear power plants is a modern expression of the dependency issue. So much probabilistic instruction and probabilistic modeling is based on independence assumptions for convenience that making strong independence assumptions is rarely questioned. Yet, grievous error can be the result. One development that has proved helpful in making both analysts and decision-makers sensitive to the issue of dependence is the influence diagram (Howard and Matheson [16] and D. Owen [22]). The decision-maker and his experts, even when not technically trained in probability, can review the assumptions of the analysis and question them where necessary.

I have emphasized that decision analysis is a paradigm suitable only for the individual decision-maker, whether he is principal or agent. I know of no decision analysis paradigm suitable for groups that are not acting within a defined agreement. For example, the question of nuclear power plant licensing involves many parties: several branches of the government, utilities, suppliers, consumers of power, environmentalists, and the general public. I can see how to build an extrapersonal decision model that would show the recommended alternative when various decision sets relevant to each party were used, and this model might be very helpful in understanding the various positions. However, I can see no way that the problem can be "solved using decision analysis" in the sense that would be correct for an individual decision-maker. I believe that the principles of decision analysis may be very important in designing agreements among individuals, but the paradigm is limited in the absence of such agreements by the very personal, normative, and judgmental characteristics that are its strength in individual decision making.

When we proceed to questions of ethics, we find that these same characteristics are the source of ethical concern. Since a decision analysis may have the appearance of impartiality even though it is based on some person's information and preferences, there is a grave danger of misrepresentation. The fact that the analysis is explicit is of great advantage in avoiding misrepresentation, but everyone who has done such analyses knows the great influence that specific assessments have over the final result. Analysts who are advocates rather than impartial analysts can have major impacts on conclusions because of conscious or unconscious choices made in assessment and modeling.

To the ethical problems of analysis we must add the ethical problems of the decision context. The analyst must question whether he wishes to participate in the process of analyzing the decision. His judgment should be based on such factors as whether the decision context is ethical to him, whether the decision-maker has title to the resources he is allocating, and whether the decision analysis itself will be marred by technical or exogenous limitations. As an example of the last category, I remember a government study in which I participated where a decision analysis was to be done to determine the national interest but subject to the condition that no alternative could be considered that was contrary to the current administration's policy. Furthermore, changes in this policy could not be considered as a possibility in the future, much less as a probability. These conditions made meaningful decision analysis impossible, and so the government report contained not a word of our extensive work before these limitations were imposed.

Finally, we turn to those issues that are normally not treated in even the most advanced analyses. I have already discussed animal rights; another similar issue is the effect of decisions on future generations. When I first became interested in this issue, I started to ask people about a proposition I call the "galactic bargain." Suppose that we were approached by galactic travelers with very advanced technology who offer us the following opportunity. They guarantee to double the standard of living of *everyone* on earth for one thousand years. At the end of that time they will arrange that humans will become painlessly sterile; when the last human is gone, the galactic travelers will occupy the planet. We check out their offer (with their help) and find that their references are impeccable: they have made similar offers to other planetary systems and have always fulfilled their bargain. The question I now ask each person is whether they would accept this deal. I have not yet had a positive response. Yet if we apply standard cost-benefit thinking to this proposition we find it is very attractive. If our standard methodology is in disagreement with common sense in this problem, how can we have confidence in it in the more mundane settings where it is applied?

As a less extreme example, consider the case of helium conservation. The primary source of helium is natural gas. When the gas is burned, the helium is lost to the atmosphere. Some scientists think that there will be a large demand for helium in the early 21st century because of increased use of superconductive devices. If there is little natural gas left to provide helium, as seems likely, the alternative route would be centrifugal extraction from the atmosphere, a very expensive process. In a recent thesis, Owen [23] found that at market interest rates there was no incentive to conserve helium and that if it were to be conserved people alive today would essentially have to make gifts to future generations. She suggests a methodology for accomplishing just such transfers. I expect that many similar studies will be necessary to illuminate adequately the question of what we owe to the future.

I have now explored both the promise of decision analysis and the challenges that must be faced if the field is to continue its growth. If correctly practiced at appropriate levels, in suitable problems, with ethical sensitivity, decision analysis can be, like technology, a great force for the growth of human potential.

ACKNOWLEDGMENT

This research was supported in part by Advanced Research Projects Agency, Contract No. N00014-79-C-0036, and National Science Foundation Grant No. ENG 76-81056.

REFERENCES

1. K. E. BOULDING, "The Ethics of Rational Decision," *Mgmt. Sci.* **12,** No. 6 (February 1966).
2. F. CAPRA, *The Tao of Physics*, Shambala, Boulder, Col., 1975.
3. B. DEWITT AND N. GRAHAM, *The Many-Worlds Interpretation of Quantum Mechanics*, Princeton University Press, Princeton, N. J., 1973.
4. H. DREYFUS AND S. DREYFUS, "Uses and Abuses of Multi-Attribute and Multi-Aspect Models of Decision-Making," University of California, Berkeley, unpublished report, p. 28.
5. H. DREYFUS AND S. DREYFUS, "Uses and Abuses of Multi-Attribute and Multi-Aspect Models of Decision-Making," University of California, Berkeley, unpublished report, pp. 12–13.
6. D. M. EDDY, "Probabilistic Reasoning in Clinical Medicine: Problems and Opportunities," in *Judgment under Uncertainty: Heuristics and Biases*, D. Kahneman, P. Slovic and A. Tversky (eds.), Cambridge University Press, 1980.
7. B. FISCHHOFF, P. SLOVIC AND S. LICHTENSTEIN, "Acceptable Risk," unpublished manuscript, p. 2, April 1979.
8. A. S. GINSBERG, "Decision Analysis in Clinical Patient Management with an Application to the Pleural Effusion Problem," Ph.D. dissertation, Department of Engineering-Economic Systems, Stanford University, January 1970.
9. HACKING, *The Emergency of Probability*, p. 5, Cambridge University Press, Cambridge, 1975.
10. R. A. HOWARD (ed.), "Special Issue on Decision Analysis," *IEEE Trans. Syst. Sci. Cybernet.* **SSC-4,** No. 3 (September 1968).
11. R. A. HOWARD, "Decision Analysis: Applied Decision Theory," in *Proceedings of the Fourth International Conference on Operational Research*, pp. 55–71, Wiley-Interscience, New York, 1966.
12. R. A. HOWARD, "The Foundations of Decision Analysis," *IEEE Trans. Syst. Sci. Cybernet.* **SSC-4,** 211-219, (1968).
13. R. A. HOWARD, "Life and Death Decision Analysis," in *Proceedings: Second Lawrence Symposium on Systems and Decision Sciences*, pp. 271–277, Lawrence Hall of Science, Berkeley, Calif., October 1978.
14. R. A. HOWARD, J. E. MATHESON AND D. W. NORTH, "The Decision to Seed Hurricanes," *Science* **176,** 1191-1201 (1972).
15. R. A. HOWARD, J. E. MATHESON AND D. L. OWEN, "The Value of Life and Nuclear Design," Presented at the American Nuclear Society Topical Meeting, Los Angeles, May 1978.
16. R. A. HOWARD AND J. E. MATHESON, "Influence Diagrams," Department of

Engineering-Economic Systems, Stanford University, July 1979, manuscript submitted for publication.

17. R. A. HOWARD, "Decision Analysis in Systems Engineering," *Systems Concepts: Lectures on Contemporary Approaches to Systems*, Ch. 4, pp. 51–85, R. F. Miles, Jr. (ed.), Wiley-Interscience, New York, 1973.
18. I. L. JANIS AND L. MANN, "Coping with Decisional Conflict," *Am. Sci.* pp. 657–667 (November-December 1976).
19. D. KAHNEMAN AND A. TVERSKY, "Prospect Theory," *Econometrica* **47,** No. 2 (March 1979).
20. T. LEVITT, "A Heretical View of Management 'Science,'" *Fortune* p. 50 (December 18, 1978).
21. J. E. MATHESON AND R. A. HOWARD, "An Introduction to Decision Analysis," in *Readings in Decision Analysis*, Ed. 2, pp. 9–43, Decision Analysis Group, SRI International, Menlo Park, Calif., 1976.
22. D. L. OWEN, "The Concept of Influence and Its Use in Structuring Complex Decision Problems," Ph.D. dissertation, Department of Engineering-Economic Systems, Stanford University, June 1979.
23. P. A. OWEN, "Decisions that Affect Outcomes in the Distant Future," Ph.D. dissertation, Department of Engineering-Economic Systems, Stanford University, June 1979.
24. Several Editors, "Report to the American Physical Society by the Study Group on Nuclear Fuel Cycles and Waste Management," *Rev. Modern Phys.* **50,** No. 1, Part 2 (January 1978).
25. P. SLOVIC, B. FISCHHOFF AND S. LICHTENSTEIN, "Behavioral Decision Theory," Oregon Research Institute, Technical Report DDI-7, pp. ii–iii, September 1976.
26. P. SLOVIC, B. FISCHHOFF AND S. LICHTENSTEIN, "Behavioral Decision Theory," *Annu. Rev. Psychol.* **28,** 1–39 (1977).
27. P. SLOVIC AND B. FISCHHOFF, "Determinants of Perceived and Acceptable Risk," in *Too Hot To Handle: The Management of Nuclear Wastes*, L. C. Gould and C. A. Walker (eds.), Yale University Press, 1979.
28. P. SLOVIC, B. FISCHHOFF AND S. LICHTENSTEIN, "Rating the Risks," *Environment* **24,** No. 4, M 14–20 and 36–39 (1979).
29. H. SMITH, *The Religions of Man*, p. 86, Harper & Row, New York, 1958.
30. C. S. SPETZLER AND C. S. STAEL VON HOLSTEIN, "Probability Encoding in Decision Analysis," *Mgmt. Sci.* **22,** 340–358 (1975).
31. L. TRIBE, "Policy Science: Analysis or Ideology?" *Philosophy Public Affairs* **66,** 84 (1972).
32. L. TRIBE, "Policy Science: Analysis or Ideology?" *Philosophy Public Affairs* **66,** 97 (1972).
33. L. TRIBE, "Technology Assessment and the Fourth Discontinuity: The Limits of Instrumental Rationality," *South. Calif. Law Rev.* **46,** 617 (1973).
34. L. TRIBE, "Policy Science: Analysis or Ideology?" *Philosophy Public Affairs* **66,** 100 (1972).
35. L. TRIBE, "Policy Science: Analysis or Ideology?" *Philosophy Public Affairs* **66,** 101 (1972).
36. L. TRIBE, "Ways Not to Think about Plastic Trees: New Foundations for Environmental Law," *Yale Law J.* **83,** 326 (1974).

37. L. Tribe, "Ways Not to Think about Plastic Trees: New Foundations for Environmental Law," *Yale Law J.* **83,** 1319 (1974).
38. L. Tribe, "Ways Not to Think about Plastic Trees: New Foundations for Environmental Law," *Yale Law J.* **83,** 1332 (1974).
39. L. Tribe, "Policy Science: Analysis or Ideology?" *Philosophy Public Affairs* **66,** 83 (1972).
40. L. Tribe, "Ways Not to Think about Plastic Trees: New Foundations for Environmental Law," *Yale Law J.* **83,** 1336 (1974).
41. L. Tribe, "Ways Not to Think about Plastic Trees: New Foundations for Environmental Law," *Yale Law J.* **83,** 1345 (1974).
42. A. Tversky and D. Kahneman, "Judgment under Uncertainty: Heuristics and Biases," *Science* **185,** 1124–1131 (1974).
43. ORSA Ad Hoc Committee on Professional Standards, "Guidelines for the Practice of Operations Research," *Opns. Res.* **19,** 1123–1258 (1971).
44. "Reaction to the Guidelines for the Practice of Operations Research," *Opns. Res.* **20,** 205–246 (1972).
45. Synfuels Interagency Task Force, *Recommendations for a Synthetic Fuels Commercialization Program, Vol. I,* November 1975.
46. Synfuels Interagency Task Force, *Recommendations for a Synthetic Fuels Commercialization Program, Vol. II,* Cost/Benefit Analysis of Alternate Production Levels, November 1975.

APPLICATIONS

Preface

The papers that follow contain discussions of applications in investment and strategic planning, research and development, and social policy.

Investment and Strategic Planning

This section contains several examples showing the applications of decision analysis to corporate investment and strategic decision problems. These papers range from an early survey to discussions of the most recent developments.

"Decision Analysis Practice: Examples and Insights" is an early survey commenting on the principles of practice, which draws from examples in new-product development, electrical power system capacity expansion, and space program planning, and which includes a discussion of the first computer system for analyzing very large decision trees.

"Decision Analysis of a Facilities Investment and Expansion Problem" presents a typical early application. It clearly illustrates the concept of the value of perfect information and demonstrates the importance of two critical value measures: time and risk preference.

"Strategic Planning in an Age of Uncertainty" shows how decision analysis deals directly and effectively with uncertainty. It discusses a 1967 probability assessment on the reopening date of the Suez Canal, revealing how a major integrated mining, refining, and marketing company used decision analysis to invest in an ocean shipping system for transporting its ore.

"A Decision Analysis of a Petrochemical Expansion Study," which presents one of the few available undisguised corporate applications, describes Gulf Oil Chemical Corporation's decision to build a one-billion dollar major capacity addition to its olefins system. The paper traces the evolution of the thinking of top management, from its expectation that decision analysis would confirm its initial judgment, through its disbelief when early indications showed that new alternatives might be better, to its final acceptance of the soundness of a totally new alternative.

"The Dangerous Quest for Certainty in Market Forecasting" shows how ineffective methods of dealing with uncertainty can lead to costly mistakes. It presents a probabilistic development of a market forecast drawn from experience at the Pharmaceutical Division of CIBA-GEIGY and shows how the forecast was used to determine production capacity.

"An Inside View: Analyzing Investment Strategies" reveals how another Swiss pharmaceutical and chemical company, Hoffmann-LaRoche, has developed a comprehensive decision analysis process for analyzing important investment decisions. The paper lists eight key reasons for the success of decision analysis at Roche and describes the evolution of decision analysis into an analytical process whose results are routinely presented to top management.

"Managing the Corporate Business Portfolio" was written especially for this collection to present some very recent developments in strategic analysis. It develops measures of risk and return, soundly based in decision theory, that appeal to businessmen. Then, it shows how multiple business units can be combined to produce a portfolio of businesses that compensates for risk and, therefore, lowers the overall risk/return ratio.

Research and Development

Probably over one-half of the decision analysis applications have been for research and development decisions. This is appropriate because R&D is inherently an uncertain venture that produces results for the distant future.

"Overview of R&D Decision Analysis" is a short description of the basic paradigm and rationale for applying decision analysis to research and development.

"Using Decision Analysis to Determine R&D's Value," a reprint of a recent two-part summary article, delineates the unique characteristics of R&D that must be treated in the decision analysis process.

"Selecting Projects to Obtain a Balanced Research Portfolio" presents a new concept for addressing the question of how to balance the whole portfolio -- the research portfolio grid. The paper discusses the trade-off between projects having a high probability of a moderate level of success and those having a low probability of an extremely high level of success. This approach provides a framework for communication among researchers, marketers, planners, and managers.

"Calling the Shots in R&D" describes Sandoz's experience in assessing uncertainty in R&D projects over a seven-year period. It concludes that probability judgments are reliable and provide a guide to readjusting R&D priorities -- particularly the timing of projects -- to deliver more predictable research results.

"Quantifying and Forecasting Exploratory Research Success" focuses the ideas of decision analysis on the exploratory phase of the research and development process. It shows how dynamic assessments of research success probabilities can be combined with generic product life cycle models to evaluate individual research projects and to combine them into a picture of the total portfolio.

"Evaluating Basic Research Strategies" shows how basic research results can be described in generic terms by using generalized models of how new ideas are developed. The example presented here focuses on CIBA-GEIGY's strategy for directing research in infectious disease chemotherapy.

Social Policy

This group of papers shows how decision analysis can aid in solving social decision problems; they introduce basic concepts and demonstrate them in three areas of social concern.

"Social Decision Analysis" shows how the principles of decision analysis can be applied to social choice problems, which illustrates the calculation of net social benefit or social surplus, which provides a measure parallel to profit in a commercial enterprise. It includes examples in emission control, weather modification, and nuclear safety.

"The Decision to Seed Hurricanes" is a fascinating discussion of a government decision to control the forces of nature. It analyzes whether the government should seed hurricanes with silver iodide crystals to mitigate their destructive effects. The concept of "government responsibility cost" is used to assess the non-economic impact of the government's decision, and important issues in social decision analysis are explored in subsequent correspondence.

"Decision Analysis of the Synthetic Fuels Commercialization Program" documents a decision analysis performed for the White House. In 1975, President Ford announced a major government initiative to commercialize synthetic fuels. By showing the structure of the problem and the actual probability assessments made by an interagency task force, the analysis clearly revealed why synthetic fuels were not a good social investment at that time.

"Decision Analysis of Space Projects: Voyager Mars" shows the development of methodology for planning an entire program of Mars exploration missions. It illustrates the advantages of a staged series of successively refined decision models and the use of large-scale decision tree analysis.

9

DECISION ANALYSIS PRACTICE: EXAMPLES AND INSIGHTS

James E. Matheson

Strategic Decisions Group

Reprinted with permission of the author from Proceedings of the Fifth International Conference on Operational Research, Venice, 1969. (John Lawrence, ed.) London: Tavistock Publications, 1970. pp. 677-691

Decision Analysis Practice: Examples and Insights

JAMES E. MATHESON

INTRODUCTION

Decision analysis is a discipline that merges the logical foundations of statistical decision theory with the capabilities of modeling and solving complex problems developed in the fields of systems analysis and operations research [1, 2]. Statistical decision theory forms both a logical structure for describing the uncertainties, values, and preferences that are relevant to a decision and a set of mathematical techniques for treating problems in which uncertainty is a factor. The fields of systems analysis and operations research provide the methodology for applying abstract models to complex, real-world situations. Together these foundations yield the new discipline of decision analysis. Using the decision to be made as the focal point of the analysis, the analyst tailors his modeling and information gathering efforts to the specific decision. In this paper I will describe the professional practice of decision analysis and will present several applications of it that are familiar to me.

BOUNDING THE DECISION PROBLEM

In approaching a problem, the decision analyst's first responsibility is to define clearly the decision to be made. Since most, if not all, decision problems are subordinate to some higher-level system, it is vitally important to bound the decision problem; that is, to establish who has the responsibility for making the decision, to determine what resources are to be allocated, and to set out which values and preferences are to be delegated explicitly by the higher-level system and which ones are to be specified by the decision-maker. For example, if the decision calls for allocating funds for new capital investments, the analyst might decide to use interest rates derived from a higher-level financial system and to use present worth of profits as the measure of value. However, if the decision calls for securing financing, considering the characteristics of each method of financing might well be within the bounds of the problem. Many times a problem is 'difficult' because of the way in which the boundaries of the problem have been specified. In many cases, the analyst can transcend such difficulties by changing the specification of the bounds.

ESTABLISHING THE EXTENSIVENESS OF THE ANALYSIS

The extent of the analysis that should be applied to any decision problem depends on the value of the resources that are at stake and the likelihood that the analysis will improve the outcome of the action taken through the selection of a 'better' decision. In fact, establishing the economic value of the analysis is a decision analysis in itself [3]. However, generally the amount of resources being allocated to the analysis is too small to justify such formal treatment.

In practice, an attempt is usually made to carry out a simplified analysis of the entire decision problem. Techniques such as sensitivity analysis and determination of the value of perfect (and sometimes imperfect) information indicate where the model should be refined and the kind of new information that should be gathered. In many cases, the analysis effort goes through three stages. The first is the pilot stage, in which the conceptual structure of the analysis is created and tested, while many of the detailed features of the problem are suppressed. During the next stage, the prototype stage, a more detailed analysis is carried out in an attempt to capture all of the relevant features of the problem. This stage is likely to involve the development of large computer models. The final stage is the production analysis, in which all aspects of the problem are critically reviewed and a decision is recommended. The decision may, of course, be a decision to gather more information and incorporate it into the analysis before making the final decision.

RELATIONSHIP BETWEEN THE ANALYST AND HIS CLIENT

The decision analyst usually serves a decision-maker or a decision-making body that I will call the client. The decision analyst is expert only in his discipline, while the client holds the resources, and knows the information, the values, and the preference that form the decision problem. If the analyst is to conduct an unbiased analysis, he must be careful to encode only his client's information and avoid biasing his analysis by inserting his own opinions.

To allow the analyst to maintain this division, the client must clearly designate who will be responsible for supplying various kinds of information, values, and preferences. In complex problems, much of the information is encoded in the structure of the model itself. Building and verifying the decision analysis model requires an interaction between the analyst and client that is perhaps the most difficult and challenging part of the task.

EXAMPLES

In the rest of this paper, I will present three examples of applications for purposes of illustrating the practice of decision analysis. The first is a 'typical application' to a new product development decision. The second is the result of decision analysis research on space program planning. The last is a large-scale application to planning for an electrical power system.

NEW PRODUCT DEVELOPMENT

A major manufacturing research company had developed two compounds for a particular market. Compound A was developed and tested to the point where it was well beyond the research stage and one alternative was to develop it into the final product. Another alternative was to develop compound B, which was still in the research stage, but was thought to be more potent than compound A. A third alternative was to abandon the whole effort.

It was thought that the development of the new product would be lengthy and expensive and that the potential market was very uncertain. Since this was a new marketing area for the client, he engaged an outside expert to carry out a market survey for use as one of the informational inputs to the decision analysis.

The analysis followed quite closely the decision analysis cycle displayed in Figure 1. The deterministic phase was begun by laying out the decision tree shown in Figure 2. The first decision was whether to develop compound A or compound B (or both) into a final product. This development determined the production cost of the compound and the concentration of it that would be required in the final product. After this determination was complete the choice of whether to market or abandon the product could be made. There were still uncertainties about the size and growth rate of the market and the action of competitors. These additional facets of the problem were represented in the structural model shown in Figure 3.

Many of the variables in the problem were subjected to sensitivity analysis. The most sensitive variable, international market size, produced changes of 16 million dollars in the present value of profit. Five variables were selected as aleatory variables–variables whose uncertainty is encoded in terms of probability distributions–for the probabilistic phase.

Figure 1 The decision analysis cycle

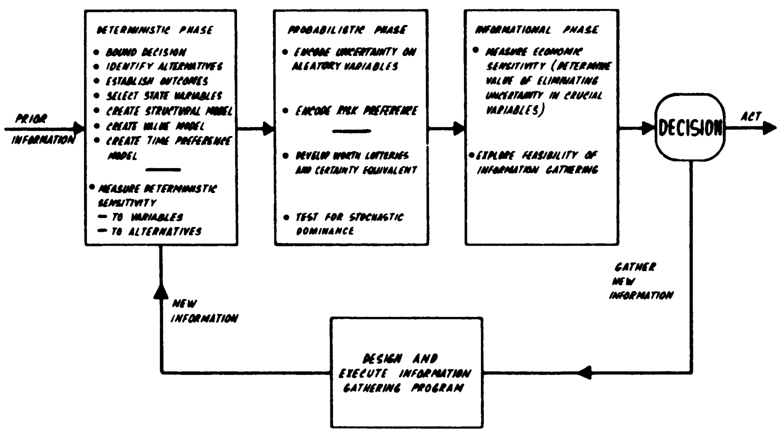

Figure 2 Simplified decision tree for the new product introduction

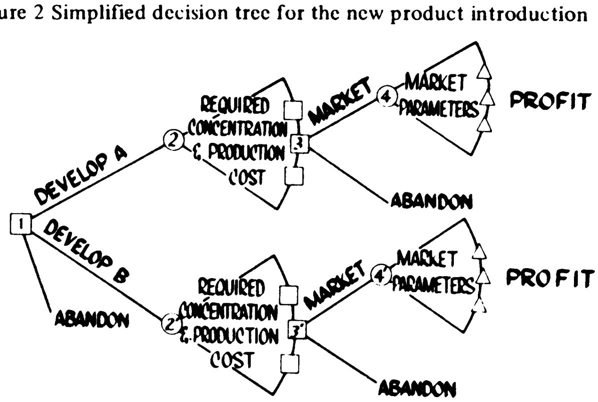

In the probabilistic phase, the simplified decision tree (Figure 2) was developed into a detailed decision tree, assigning actual conditional probabilities to the aleatory variables represented in the structure of the tree. At the tips of this tree, expected profits were assigned by a Monte Carlo simulation of the structural model of Figure 3, which contained the remaining aleatory variables. The decision tree was then evaluated on an expected value basis. The amount of corporate resources to be devoted to this product were small enough so that no significant risk aversion was desirable.

Figure 3 Structural model for the new product introduction

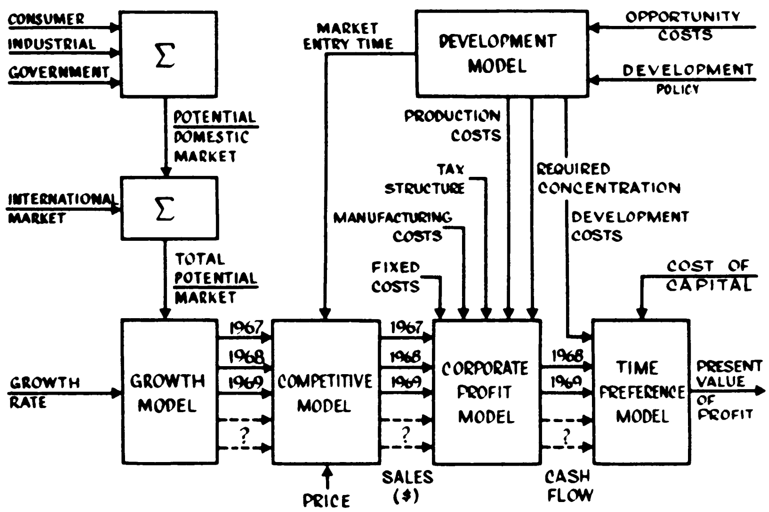

The result of the probabilistic phase was that the profit lottery for development of compound B stochastically dominated that of compound A. However, the profit lottery for the development of compound B, with the cumulative probability distribution shown in Figure 4, had negative expected present value, so the best decision was to abandon the effort.

In the informational phase, the expected value of perfect information (economic sensitivity) was computed on several important aleatory variables. The highest economic sensitivity of $1,415,000 was exhibited by the international market size. The international market size showed such high economic sensitivity because the new information might reveal a very large international market for the product, making it profitable to go ahead with the development of compound B in light of this new information. As a result, the client undertook a more extensive analysis of the international market for his product.

SPACE PROGRAM PLANNING

The space program planning application was conducted for the purpose of developing a methodology that would be useful in approaching technically complex decision problems; the intent was to carry out research on decision analysis itself. Although a very detailed analysis of the U.S. program for the unmanned exploration of Mars was conducted, no attempt was made to recommend specific decisions to the U.S. government. Instead, a large corporation that was quite familiar with the space effort played the role of the decision-maker during the analysis.

Figure 4 Profit lottery for the development of compound B

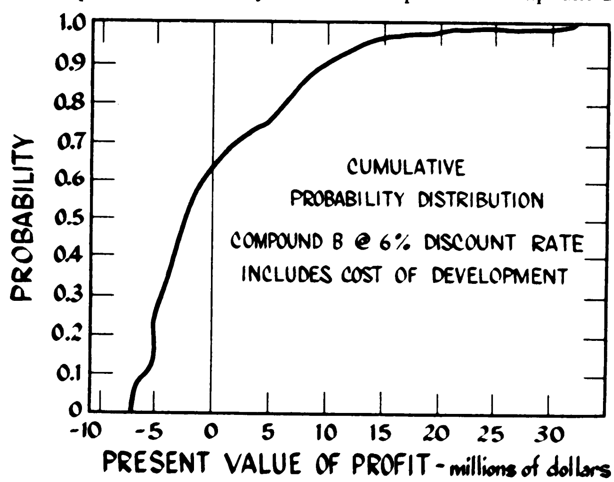

The problem was to determine the sequence of designs of rockets and payloads that should be used to pursue the goal of exploring Mars. It was considered desirable to place vehicles in orbit around Mars as well as to explore its atmosphere and to land vehicles on the surface of the planet to collect scientific data.

For purposes of obtaining sufficient information to encode properly the complex structure and information required to analyze this problem, a decision analyst resided with the client for a period of about one year. The client and the decision analyst worked as a team in building the models and submodels for the analysis.

The work was begun with a pilot phase, in which a simplified version of the decision problem was constructed. During this phase, four possible designs were postulated; each design represented increasing levels of sophistication. Figure 5 shows these designs and their potential accomplishments. In the prototype analysis there were 12 possible vehicle designs plus the alternatives of skipping opportunities on cancelling the program at any decision point.

Because of the behavior of the orbits of the Earth and Mars, an opportunity to launch a vehicle toward Mars occurs about once every two years. Consequently, the decision problem was characterized by a sequential decision process, where each decision can be contingent upon the entire project history

Figure 5 Pilot model assumptions for the exploration of Mars

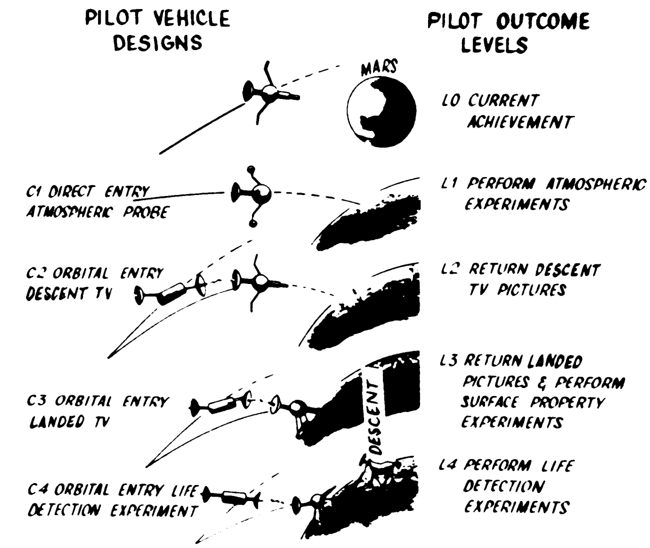

that precedes the decision point. Because of the lead time required in constructing a given vehicle, it was necessary to make each vehicle design decision before the outcome of the previous vehicle's flight was known. A decision tree was constructed to capture the structure of this sequential decision process.

In order to create a decision tree of manageable size, the concept of state variables was introduced. The state variables are a set of variables that are selected during the modeling process and whose value at any point in time summarizes all of the past history of the project relevant to future decision-making. Each node in the decision tree is characterized by a set of values for each of the state variables. The probabilities, cost, and values of subsequent branches are assigned conditionally on the basis of these values. Creativity is required in the selection of state variables. If a good approximation to the total available information is to be obtained, an appropriate set of state variables must be judiciously selected. A major objective in this process is to discover where essentially the same point can be reached via different paths through the program. When such a point is reached, two or more branches in the decision tree coalesce at a single node. The node is assigned the common value of those state variables that are reached at this point along either path. This property, called coalescence, greatly reduces the size of the decision tree characterizing the problem. The sizes of the uncoalesced and coalesced decision trees for both the pilot and prototype decision trees are presented in Table 1.

Table 1 Summary of decision tree scale

	PILOT Uncoalesced	Coalesced	FULL SCALE Uncoalesced	Coalesced
Number of Nodes	3,619	56	476,012,807	3,153
Number of Branches	3,618	126	476,012,806	22,784
Number of Paths	1,592	1,592	354,671,693	354,671,693
Number of Policies	3,005	somewhat less	over 10^{39}	somewhat less

The assignment of the probability, cost, and value parameters to the branches of the decision tree was a task that required the incorporation of information from additional submodels. For the pilot analysis, these models were kept quite simple.

In the prototype analysis, the most complex submodel was the probability model. Essentially, a probability tree was constructed from detailed diagrams that showed the functional steps in any flight to Mars. This tree had on the order of one hundred nodes, and the probabilities assigned to its branches were either obtained directly, from judgment combined with experimental data, or indirectly, from yet another sublevel of probability models. At each chance node in the decision tree, the detailed probability model produced the probability for each possible outcome.

Another unusual model was the value molel, that is, the model that assigned a monetary value to each outcome in the space program. Since the client was reluctant to assign values directly in monetary terms, a cardinal scale of benefits was first employed. This scale was constructed so that the benefit of a perfect project would be one point. A total monetary value assignment to a perfect program then determined the monetary values to be used in the decision tree.

The benefit scale was determined by constructing a value tree. The value tree is simply a convenient method of breaking the total benefit of the project into the incremental benefit of each individual outcome. Figure 6 shows a value tree for the pilot analysis. The value tree was constructed by dividing the benefit of the entire program (one point) into major categories, and then into subcategories identified in increasing detail until no further distinction was desirable. Each tip of this tree is divided into additional categories. Each additional category represents an elemental outcome that may be achieved during the project. For example, in the figure, the number 1.0 beside the node at the extreme left represents the

Figure 6 Space program value tree

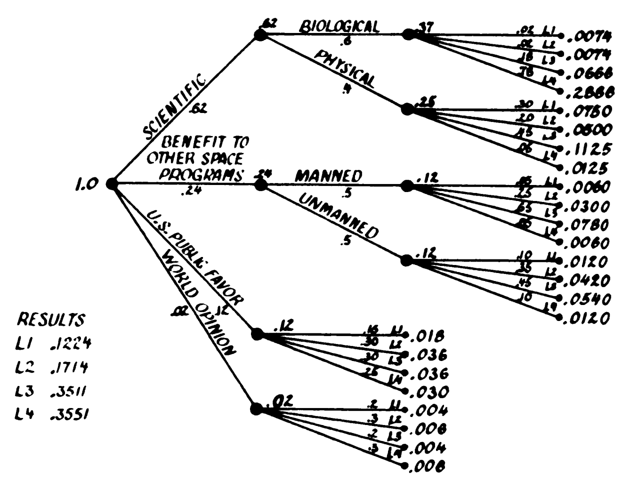

total benefit of all the objectives of the program (achieving outcome L1, L2, L3, and L4 of Figure 5). The upper branch represents all direct scientific benefits of the program and was assigned 62 per cent of the total value. The succeeding biological branch was assigned 60 per cent of the scientific benefit, yielding 37 per cent as the total benefit of the project to biological science. The further subdivision from this node represents the four increments in outcome level that are presented in Figure 5. Finally, the terminal code benefits were added for each level of outcome to give the totals shown in Figure 6. These totals, when multiplied by the total monetary value assigned to the program, determined the assignment of values to each outcome branch in the decision tree. A more detailed value tree was constructed for the prototype analysis.

In the pilot phase, calculations for the decision tree and the three submodels were made on a time-sharing computer system. The programming was carried out primarily by the decision analyst during the formation of the conceptual structure of the problem in the pilot phase.

The pilot model provided a good means of communicating the concepts of the analysis and for making rapid sensitivity calculations. The pilot analysis could be demonstrated during meetings and presentations at which the results of changes in the parameters of the model could be determined almost instantaneously. In many cases, decision-makers would supply their values assignments for purposes of determining how the policy would be changed by them.

Because of the large size of the prototype model, the analysis programs were implemented in a system of programs called SPAN (Space Program ANalysis). The SPAN system is outlined in Figure 7.

Figure 7 Span system operation

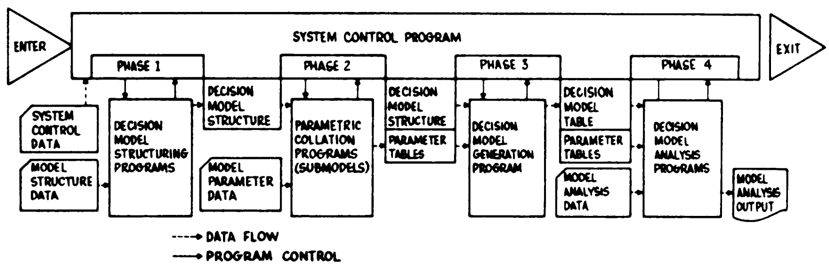

The large size of the decision tree structure made it impractical to draw the complete tree by hand. Thus, the tree was generated by a computer program that utilized structural information describing characteristics of the decision tree to generate a symbolic description of the decision tree. This symbolic description was then compiled into a computer representation more suitable for computation. The generation and compilation were carried out in Phase 1.

In phase 2 the cost, value, and probability model were executed, and from them, the numerical values of these parameters were generated and collated with the symbolic representations produced in Phase 1.

Phase 3 was a computer bookkeeping phase that operates on the decision model structure and the parameter tables for purposes of changing the information into a more efficient format for the analysis programs.

Phase 4 executed analysis programs that performed the roll-back of the decision tree, to determine optimum policies, and the determination of the probabilities of the various events in the tree. It was capable of applying discount factors that represented time preferences and the exponential utility function that represented risk preference.

ELECTRICAL POWER SYSTEM PLANNING

The goal of this application was to create a basis for deciding when and whether to install a nuclear generating plant in Mexico. Because electrical generating plants have very long lifetimes, the desirability of any installation depends on the characteristics of the future system expansion. Consequently, each specific installation decision must be made within the framework of a policy for overall power system expansion.

In order to carry out this analysis, a project team, which included four representatives from Mexico and four decision analysts, was brought together for a period of about one year. The role of the Mexican representatives was to provide technological expertise, to collect necessary data, and to gather judgments regarding the preferences of the country of Mexico.

Figure 8 A decision analysis model of the Mexican electrical system

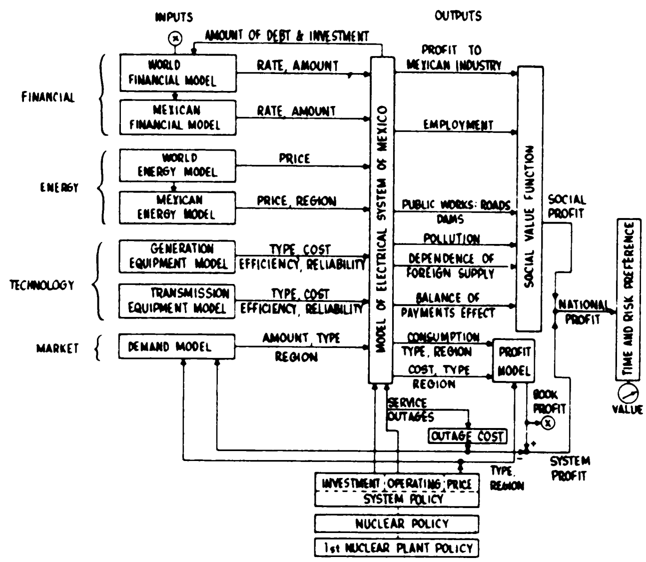

The conceptual framework for this problem is presented in Figure 8. At the left of the figure are the environmental inputs of the power system. These divide into four major categories– finance, energy, technology, and market. The financial model characterizes the terms at which capital is available from both domestic and world financial institutions and markets, as a function of the profitability, debt, and equity of the power utility. The energy model describes the price of all potential fuels–such as oil, natural gas, and uranium–as well as the availability of other energy sources–such as water power–over the time period considered in the analysis. Similarly, the technology model characterizes the availability and prices of various types of generation equipment. Finally the demand model characteristics of electrical demand growth over time, ideally as a function of the price charged for electrical service.

At the bottom of the figure is the policy stating the conditions under which the first nuclear plant should be installed. The figure shows that this policy must be embedded in the general nuclear policy, which in turn is embedded in the system's investment, operating, and pricing policy.

All of the environmental inputs and the policy alternatives feed into a model of the electrical system of Mexico. Application of the model determines the output variables over time. In the lower right corner of the figure, the outputs that indicate financial performance are shown. The amount of electrical consumption, the price of electricity, and the various costs are all combined to produce the usual book profit. Since reliability of service is one of the major considerations in electrical system expansion, the outage cost model is used to determine a monetary deduction from book profit, which yields system profit.

The social value function in the upper right-hand corner of the figure was included so that national goals that are outside the normal purview of the electrical system management could be considered. Its purpose is to assign a monetary value, called social profit, to social benefits of profit to Mexican industry, employment, public works, pollution, dependence on foreign supply, and effect on balance of payments. The sum of the social profit and the system profit is the national profit.

The uncertain time profile of national profit is converted into a single value, which might be called certain present national profit, by means of the time and risk preference model. The best decision policy is the one that maximizes the setting on this 'value meter'.

The development of this conceptual structure into a formal planning tool for system expansion proceeded through the pilot, prototype, and production stages described earlier. It must be pointed out that since an electrical system is so complex, different features of the planning model become important for different installation reasons. Thus, it is crucial that the analyst revalidate the model, through techniques such as sensitivity analysis, to ensure that it adequately captures the essence of each new installation decision.

The analysis was carried out through the development of a system of computer programs that simulate and evaluate the installation and operation of the electrical system over many years. The programs determine the cost of operation, including effects of maintenance, plant mix, system reliability, and possible energy deficits. Within this large simulation of the electrical system, the installation policy routines carried out less detailed simulations and evaluations of the system's future for the purpose of determining the time that each installation should be made and the type of installation it should be. The installation policy was refined so that the resulting installations would maximize the reading on the 'value meter'.

The pilot phase demonstrated the need for elaborate models that were capable of capturing the complexities of the electrical system problem. Thus, during the prototype phase, a modular system of computer programs was constructed. This modular system facilitated the implementation of changes that would naturally occur in the transition to the production phase and permits the appropriate module to be easily updated as the nature of the electrical system changes in the future. The computer model was constructed from a number of independent submodels that communicate through well-defined variables and tables.

One of the most significant submodels developed was the reliability submodel. In the ordinary expansion of an electrical system, each new plant is installed for the purpose of maintaining reliability in the face of demand growth. If plants did not randomly fail, an electrical system could operate with a much smaller capacity. Thus a computational procedure was developed to compute the system reliability from probabilistic demand information and the failure probabilities of each plant in the system. The effect of scheduled plant maintenance on reliability was included in the computation.

An interesting feature of power system expansion is that the system is self-healing. That is, if a 'wrong' plant is installed at any time, or if the environment changes, the effects can be largely compensated for by the choice of new installations. Because an electrical system operates with a mix of plants–some best for steady base load and some best for rapid peaking–the new installations required by the usual rapid system growth can be selected so that the plant mix will be readjusted within a few years.

GAPS IN THE THEORY

Perhaps the widest gaps between theory and practice are in the areas of values and preferences. Methods of solving even the seemingly simple problem of characterizing time preference leave much to be desired. There is a great deal of controversy over the choice of a discount rate, and few guides exist for determining when a discount rate adequately represents time preference characteristics. Suggestions conflict about when the discount rate should be used to represent financing terms, and when it should be used to represent risk aversion.

Utility theory provides an elegant foundation for describing attitudes toward risk. However, seldom, if ever, are all the sources of uncertainty quantified. In addition, since each decision problem is part of a higher-level system, it is often not clear just what risk preference can be normatively deduced from higher-level considerations. In many applications, sensitivity to risk preference can be determined through the use of a family of utility functions, such as the exponential family.

Problems dominated by time or risk preference alone, usually can be adequately treated in spite of the above mentioned problems. However, when time and risk preference must be treated jointly, theoretical foundations are almost nonexistent. Techniques combining discount rates and the exponential family of utility functions were developed for use in the decision trees of the space program planning example [5]. A recent doctoral dissertation considers the joint time–risk preference from fundamental attitudes toward consumption [6].

Some of the most perplexing problems arise, however, in the analysis of public decision problems. In the electrical system planning example, the space program planning example, and in applications to regulatory and natural resource decisions currently in progress, the specification of the value function is a difficult task. The economic literature provides little guidance in the establishment of values for public decisions. In fact, many authors begin their developments with different implicit assumptions about the nature of the values. One example is the literature on marginal cost pricing [7]. I suspect that the resolution of these difficulties will come when the needs for explicit choices of public values are separated from their theoretical consequences.

CONCLUSION

The new discipline of decision analysis has been illustrated in practice with several examples. In my experience, decision analysis has proven to be a useful approach to complex decision problems. It provides not only the principles necessary for analysis, but also a means of bringing the important issues of the problem into focus, so that new alternatives can be created, information gathering possibilities can be evaluated, and the analysis effort itself can be efficiently channeled. Applications have shown the need for new theory and methodology for treatment of values and preferences, especially in public decision problems.

REFERENCES

1. R. A. HOWARD (1966) Decision analysis: applied decision theory. *Proceedings of the Fourth International Conference on Operational Research,* Session 2, International Federation of Operational Research Societies, Boston.
2. R. A. HOWARD (1968) The foundations of decision analysis. *IEEE Transactions on Systems Science and Cybernetics* **SSC-4** (3).

3. J. E. MATHESON (1968) The economic value of computation and analysis. *IEEE Transactions on Systems Science and Cybernetics* **SSC-4** (3)
4. J. E. MATHESON and W. J. ROTHS (1967) Decision analysis of space projects. *Proceedings National Symposium Saturn/Apollo and Beyond.* American Astronautical Society, June 11, 1967.
5. J. E. MATHESON and R. A HOWARD (1967) Decision analysis of multi-furcate trees. Presented at the Institute of Management Science XIV International Meeting, Mexico City, August 24, 1967.
6. A. POLLARD (1969) Ph.D. dissertation (in preparation), Department of Engineering Economic Systems, Stanford University.
7. N. RUGGLES (1949-50) Recent developments in the theory of marginal cost pricing. *The Review of Economic Studies* **17**

10

DECISION ANALYSIS OF A FACILITIES INVESTMENT AND EXPANSION PROBLEM

Carl S. Spetzler
Strategic Decisions Group

Ramon M. Zamora
ChildWare Corporation

CONTENTS

ILLUSTRATIONS

TABLES

I INTRODUCTION

This paper presents the results of a decision analysis of a major facilities expansion decision.* The discussion describes the conclusions, the recommendations, and (in a brief form) the methodology of the analysis.

It was initially decided to focus the analysis on a plant investment decision then facing the decision maker. The investment involved building a complete new facility based on a successful pilot plant that was already in operation. Some of the major areas of uncertainty surrounding the decision were:

- The process involved a chemical reaction that was difficult to control.
- The major product from the facility would be a brightener, but by utilizing new processing techniques a valuable by-product could be produced. The exact yields of both products were uncertain.
- Minute quantities of impurities in the raw materials used in the process could greatly affect the amounts of brightener and by-product produced.
- Market factors, including costs of materials and sales prices of products, were not known exactly.
- The parameters involved in taking the process from the pilot model up to a full-scale production facility were uncertain. Specific parameters involved included the efficiency of the new plant and the additional design costs.
- Some government regulatory effects could be expected in the areas of pollution control of the new plant and pricing controls on the by-product, but neither the extent of the regulations nor the cost imposed could be determined beforehand.

---

* This paper is tutorial in nature but is based upon an actual decision analysis performed by the SRI Decision Analysis Group. The original decision problem has been disguised throughout the ensuing discussion.

The basic decision alternatives under consideration were:

(1) Abandon the project and accept the loss of investment in the pilot plant.

(2) Build the initial production plant.

(3) Build the initial production plant and, if it is successful, consider an expansion investment.

(4) Postpone the initial production plant development to gather additional information regarding:

  (a) The amount of by-product produced.

  (b) The effect of impurities in the raw materials.

  (c) Economic factors.

The investment decision involved a \$25-\$50 million outlay of funds. The lower value is for the initial plant investment; the higher value is for the case where expansion of the facility appears feasible.

The decision analysis involved three parties: the president of the manufacturing division of a large organization (the client and decision maker), the client's technical staff, and SRI's Decision Analysis group. SRI had the responsibility for the application of decision analysis methodology and for the overall system configuration. The client was responsible for the information used in the analysis and provided detailed submodels for the system in his particular areas of expertise.

II CONCLUSIONS AND RECOMMENDATIONS

(1) The venture has an expected Present Value of $12 million, based on present information and a discount rate of 10%.

(2) Adjustment for risk still indicates a positive Present Value.

(3) The value of information, which measures the maximum anyone would be willing to pay to reduce the uncertainty in a variable indicates large uncertainties as to the amount of by-product produced and the expected level of raw-material impurities. Research should be directed toward reducing these uncertainties.

(4) The total value of information is large enough to merit postponing the investment in the production plant.

(5) The decision is highly sensitive to both time and risk attitudes. These implications should be made clear to corporate policy makers and their preferences applied to the results of this study.

III DECISION ANALYSIS--THE DETERMINISTIC PHASE

A decision analysis is carried out in three phases as a normal part of the decision analysis cycle, Figure 1. The first phase is the deterministic phase. The main purpose in the deterministic phase is to structure the decision problem and then to use the structure to identify the crucial variables.

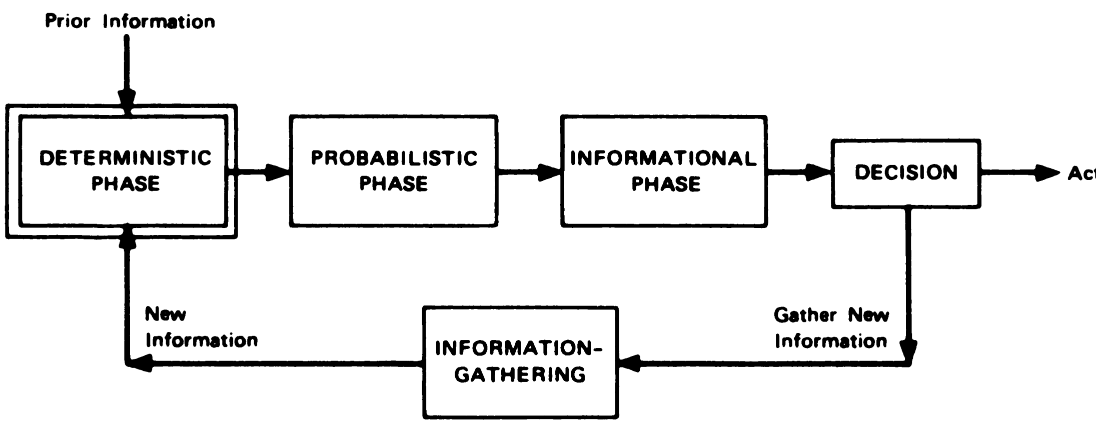

FIGURE 1 THE DECISION ANALYSIS CYCLE—DETERMINISTIC PHASE

Deterministic Model

The structure of a decision problem is initially captured in a model where the uncertainties in the variables are ignored. Figure 2 shows how this deterministic model is organized. The inputs are classified into two kinds: state variables and decision variables. State variables are the factors beyond the control of the decision maker. Decision variables are the factors that can be controlled by the decision maker. The purpose of the whole analysis is to find the best setting for the decision variables while considering the best available information on the state variables. In Figure 2, the state variables are on the left side, the decision variables are on the bottom. Both enter into the systems model. The systems model simulates the building and operation of the plant and generates the resulting cash flow.

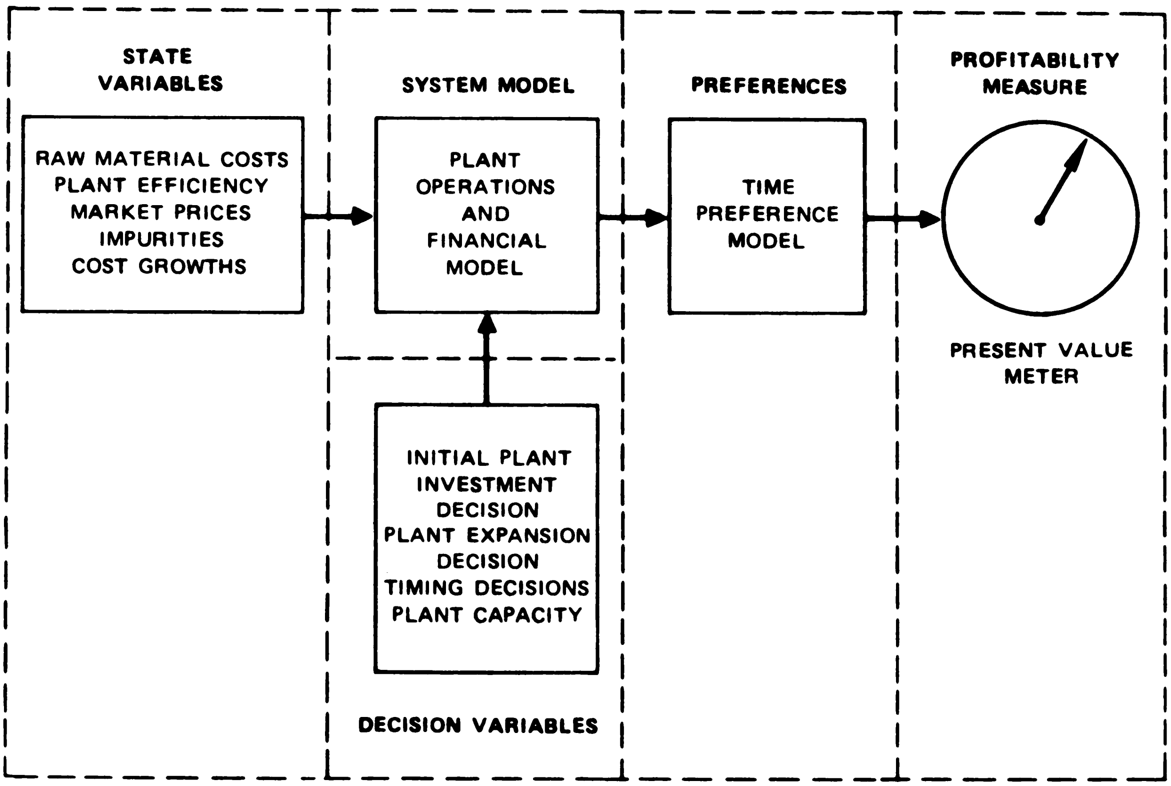

FIGURE 2 ORGANIZATION OF DETERMINISTIC MODEL

The output from the systems model is the cash flows, which show the financial effect over time of specific settings of decision variables. Many combinations of settings of decision variables and state variables are possible and should be examined. However, because it is not feasible to look at the entire cash flow pattern each time, it is important that some single measure of profitability be defined. This measure of profitability should have the following property: For any cash flow that is preferred, the measure of profitability should show a higher value. In this analysis the Present Value (PV) was used as a reasonable approximation of the decision maker's time preference. PV weights the cash flow in the form of discounting with an interest rate. Ideally, the PV of a cash flow should represent the single present sum that would make the decision maker indifferent between choosing that sum or the cash flow.

Since a choice of a discount rate is a statement of time preference, the discount rate should be set by the policy makers of the corporation. This requires the decision makers to face the fundamental question of time preference, i.e., the question of how to choose between cash flows.

Much confusion usually clouds this issue, particularly when rate of return is the measure of profitability used by a company. Often the "minimum acceptable return" is set higher to allow for a margin of safety. But rather than providing safety, the distortion usually confuses the issue. Since risk is included specifically in this decision analysis, the discount rate should represent only the effect of timing. That is, a discount rate of 10% should mean that the policy makers of the company are indifferent as to receiving $1.0 million today or $1.1 million one year from now (after taxes). For the base case in this analysis, a 10% after-tax discount rate was assumed. However, the effects of changes in the rate on the decision are evaluated.

Deterministic Analysis

The deterministic model was first used to calculate the PVs at a 10% discount rate for the major alternatives under the most likely conditions. These PVs are given in Table 1. The results show that under the estimated "most likely" conditions the initial plant should be developed and that no plant expansion plans should be considered.

Table 1

RESULTS OF DETERMINISTIC PHASE

	Present Value (millions of dollars)
Initial plant-investment decision only	$15
Plant expansion decision only	-30
Initial plant followed by plant expansion	5

Next the sensitivity of the initial plant investment decision to change on the discount rate was evaluated. Dropping the discount rate to 7% increased the PV to $25 million. Increasing the discount rate to 13% dropped the PV to $9 million. While the effect of a change in the discount rate is large, it should be noted that it would not change the recommendation.

To identify the crucial variables, a systematic sensitivity analysis was used. In all, the effect of changes in 30 variables was investigated. For each of these 30 variables an extreme range was stated by the decision maker's specialists and the effect of the extreme values on the present value was then calculated. The sensitivities to the 18 state variables that had the greatest effect are listed in Table 2. The order of the listing in Table 2 is by decreasing sensitivity. Changing the value of the first variable--the by-product produced--from 0 to 80 lb/ton changed the Present Value from -$30 to +$35 million.

For the purposes of the probabilistic phase, the first 7 variables were considered crucial state variables. Because of special interest, the plant efficiency (Variable 13) was also included as a crucial variable. Note that the pessimistic end of the range in any of the noncrucial variables would not change the Present Value enough to reverse the initial decision of developing the initial plant.

In summary, in the deterministic phase a model was developed, a likely case was analyzed, and the crucial variables were identified. To this point, the analysis does not differ greatly from the usual approaches. More than one-half of the total professional effort on this research project was utilized on the deterministic phase. This is typical of many decision analysis applications.

Table 2

RESULTS OF DETERMINISTIC PHASE--SENSITIVITIES OF STATE VARIABLES

		Base Case	Tested Range From	Tested Range To	PV Range (millions of dollars) From	PV Range (millions of dollars) To	Change in PV (millions of dollars)
1.	By-product production	36.00 (lb/ton)	0.00 (lb/ton)	80.00 (lb/ton)	$-30	$35	$65
2.	Market price of brightener in 1980	$0.27 ($/lb)	$0.15 ($/lb)	$0.35 ($/lb)	-12	45	57
3.	Raw material cost growth	5.00% (%/yr)	0.00% (%/yr)	8.00% (%/yr)	-7	40	47
1.	Raw material costs	$7.00 ($/ton)	$2.00 ($/ton)	$18.00 ($/ton)	-9	35	44
5.	Impurities in raw material	4.00 (lb/ton)	2.00 (lb/ton)	6.00 (lb/ton)	-10	30	40
6.	Cost multiplier on investment	90.00%	70.00%	125.00%	-12	25	37
7.	Brightener price growth after 1980	4.00% (%/yr)	2.00% (%/yr)	6.00% (%/yr)	-5	28	33
8.	Cost multiplier on operations expenses	110.00%	80.00%	150.00%	3	32	29
9.	Cost multiplier on maintenance expenses	100.00%	70.00%	120.00%	2	30	28
10.	By-product price growth	$0.03 ($/yr)	$0.01 ($/yr)	$0.06 ($/yr)	3	30	27
11.	Water reclamation costs	$0.03 ($/gal.)	$0.02 ($/gal.)	$0.04 ($/gal.)	4	31	27
12.	By-product price, 1970	$0.50 ($/lb)	$0.40 ($/lb)	$0.60 ($/lb)	3	29	26
13.	Plant efficiency	75.00%	50.00%	110.00%	0	26	26
14.	Brightener production	45.00 (lb/ton)	40.00 (lb/ton)	50.00 (lb/ton)	4	30	26
15.	Market price of brightener in 1974	$0.25 ($/lb)	$0.15 ($/lb)	$0.30 ($/lb)	4	28	24
16.	Government regulation costs	$0.02 ($/lb)	$0.01 ($/lb)	$0.03 ($/lb)	3	25	22
17.	Market price of other by-products	$0.10 ($/lb)	$0.07 ($/lb)	$0.12 ($/lb)	4	24	20
18.	Other by-products produced	84.00 (lb/ton)	70.00 (lb/ton)	96.00 (lb/ton)	4	24	20

IV DECISION ANALYSIS--THE PROBABILISTIC PHASE

The probabilistic phase follows the deterministic phase, as indicated in Figure 3. The main purpose of the probabilistic phase is to explicitly bring uncertainty into the analysis. Since this phase applies a number of concepts that may be unfamiliar to the reader, the steps of the analysis are described in some detail. The following steps make up the probabilistic phase.

(1) Encode uncertainty on crucial state variables

(2) Develop profit lottery

(3) Encode risk preference

(4) Determine best action with present level of information

(5) Perform further sensitivity analysis

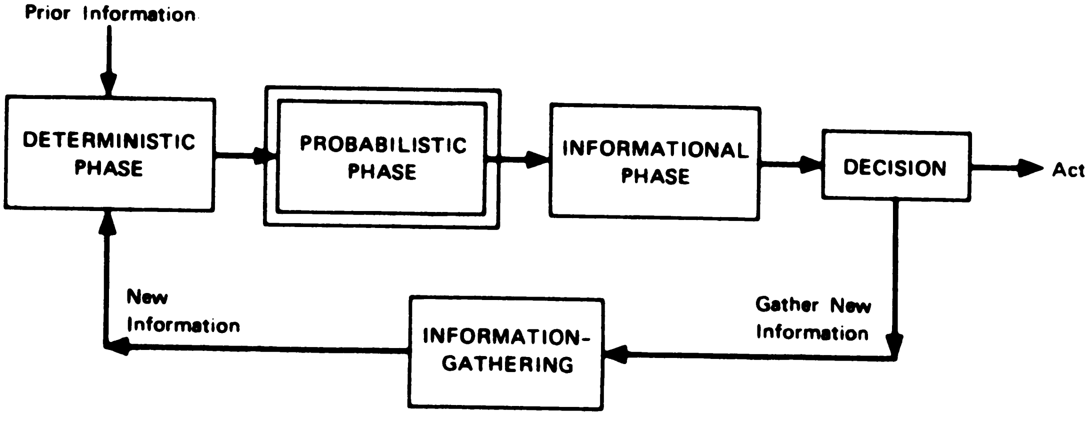

FIGURE 3 THE DECISION ANALYSIS CYCLE—PROBABILISTIC PHASE

Encoding Uncertainty

In the deterministic phase the crucial state variables were identified. Now the degree of uncertainty in these crucial state variables has to be specified.

The two plant investment decisions under evaluation are unique in many respects. Therefore the best available information consists of the judgment of the chemists, chemical engineers, and planning and operations specialists of the client. The "encoding of uncertainty" refers to the process of measuring this judgment. In the analysis interviews were used to encode the judgment of specialists. The interviewee was asked to compare the likelihood of two or more ranges. Typical questions were: "Which is more likely, that the total raw material costs end up being between \$5 and \$10/ton or more than \$12/ton?" From answers to such questions the probability distribution can be inferred. The interviewee need not have any technical knowledge of probability. He simply uses his judgment to answer the questions. Figure 4 is a probability distribution for raw material costs. The graph should be interpreted in the following way: read the x value along the horizontal scale, e.g., x = 6. Then read the corresponding probability that the raw material costs are less than x on the vertical scale, e.g., there is a 5% chance that raw material costs will be less than \$6/ton.

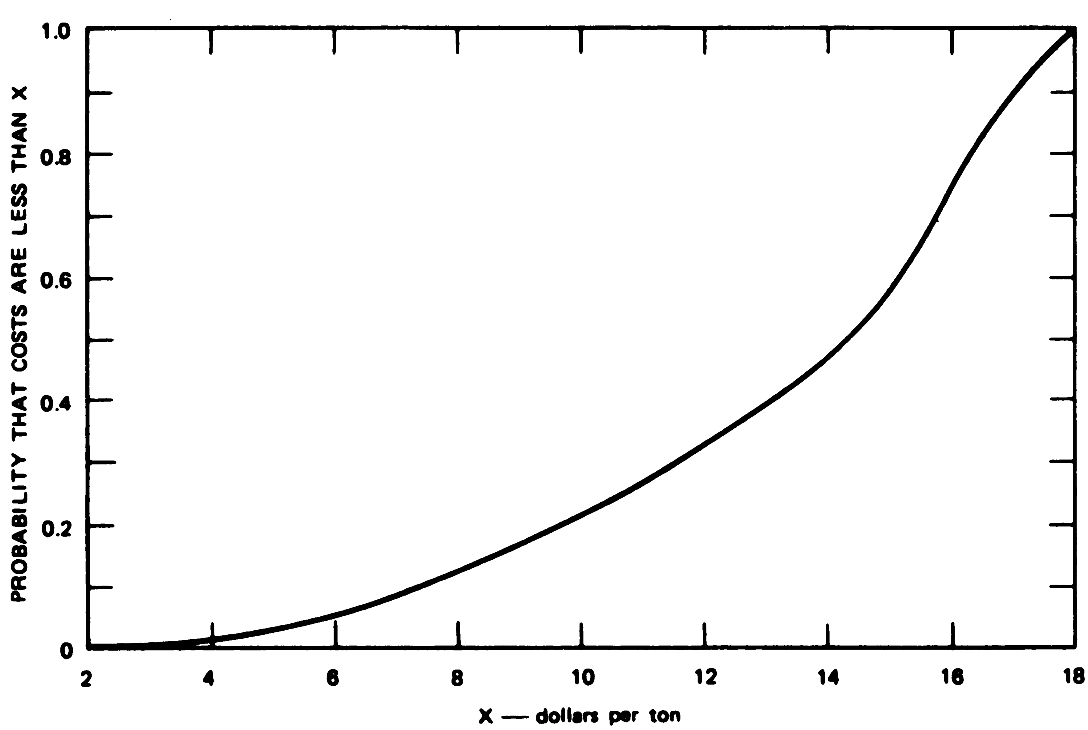

FIGURE 4 PROBABILITY DISTRIBUTION FOR RAW MATERIAL COSTS

Note that the following statements are interpretations of Figure 4.

- The raw material costs will be somewhere between \$2 and \$18/ton.
- It is equally likely that the costs will be above or below \$14/ton.
- There is one chance in ten that the costs will be less than \$7/ton.
- There is one chance in ten that the costs will be more than \$17/ton.

Of course many more such statements can be made. Any graph such as Figure 4 summarizes a specific level of uncertainty in a variable.

Since probability distributions are based on judgment by individuals, different opinions will show up as different distributions. The distributions can then be used as a communications tool for discussing differences in opinion and reaching a consensus.

Figure 4 is based on the consensus of three individuals. One individual was interviewed twice, with six weeks between the interviews. The four distributions before the consensus was reached are given in Figure 5.

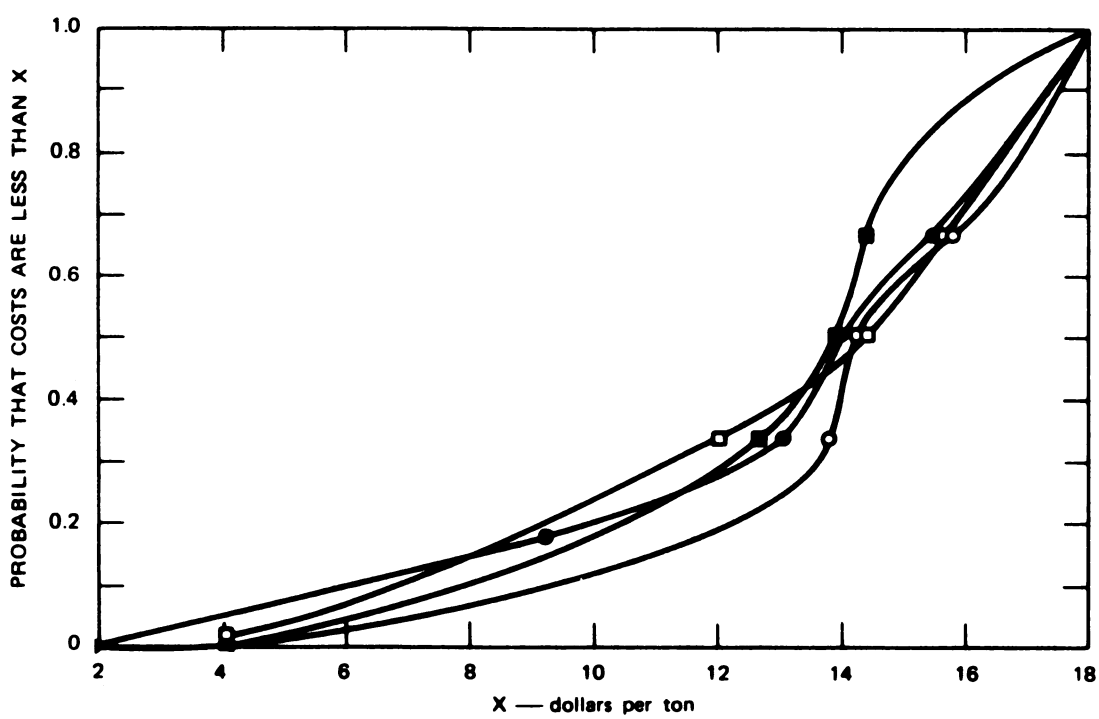

FIGURE 5 DISTRIBUTIONS OF INDIVIDUAL JUDGMENTS

The encoding procedure was repeated for all crucial state variables which resulted in similar probability distribution.

The uncertainty in the brightener price growth after 1980 and in the raw material cost growth is shown in a tree format in Figure 6. Brightener raw material price and cost growth are related. Figure 6 shows this relationship. The circles in Figure 6 are "probability nodes." The first numbers on the branches coming from the probability nodes are probabilities. Thus, the top path represents the following: there is a 30% chance that prices will grow at 6% per year after 1980; if this happens, then there is a 30% chance that costs will also be growing at 8% per year.

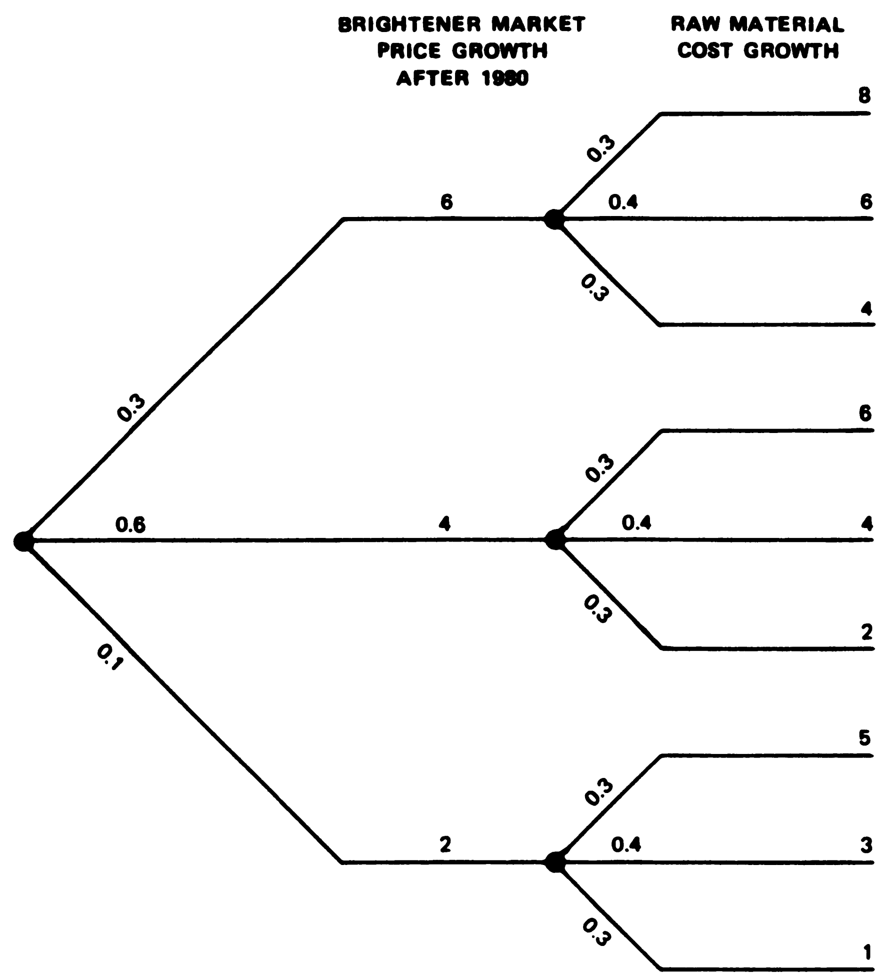

FIGURE 6 COMBINED UNCERTAINTY FOR BRIGHTENER MARKET PRICE AND RAW MATERIAL COST GROWTH

Developing the Profit Lottery

Once the uncertainty of the crucial variables is specified, the next step is to evaluate the combined effect on the profitability. The overall uncertainty in profitability is summarized in a profit lottery. A profit lottery is simply the probability distribution of Present Value. The overall profit lottery for the decision to build the plant is given in Figure 7. The profit lottery summarizes the uncertainty that the

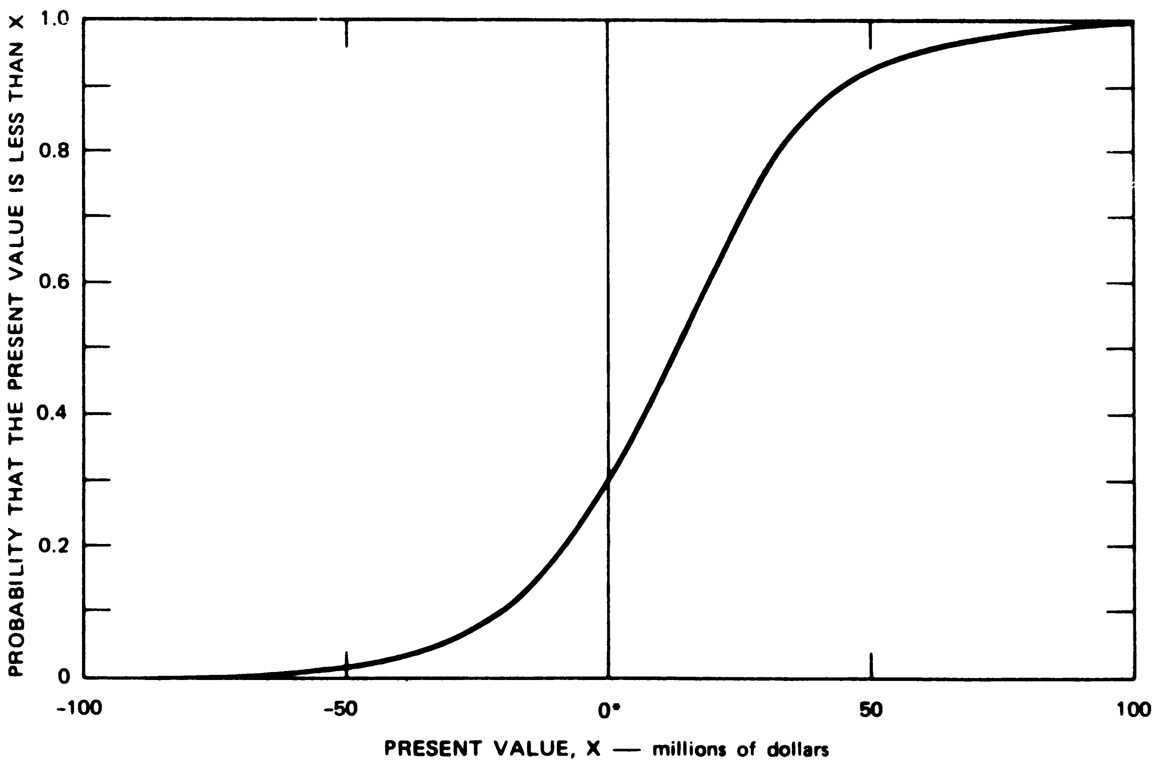

*Zero means getting only the minimum acceptable return (10%) on the investment.

FIGURE 7 PROFIT LOTTERY

decision maker faces without gathering further information. From Figure 7 it can be seen that:

(1) There is a 31% chance of a negative PV.*

(2) There is one chance in twenty of an outcome worse than a $30 million negative PV.

(3) There is a 50/50 chance of being above or below a $12.5 million PV.

(4) There is one chance in twenty of doing better than a gain of a $55 million PV.

Of course, many more such statements are possible. The reader should study the profit lottery in detail to familiarize himself with the graphical representation.

Now let us take a quick look at how the profit lottery is developed. Recall that we have already developed the following:

(1) Distributions on the crucial variables.

(2) A deterministic model that allows the calculation of the Present Value for any set of variables.

There are a number of ways that a probabilistic model can be built. In this analysis a decision tree approach was used. A schematic representation of the tree is given in Figure 8. Decision points are shown as boxes, state variables are indicated by circles. The initial decision is whether or not to build a plant at this time. Not building the plant was used as the base case and is considered as a reference point (PV = zero). If the plant is opened, many situations may result. However, as determined by the sensitivity analysis, only the crucial variables will greatly affect the Present Value of the decision. Therefore to determine the uncertainty in Present Value, combinations of numerical values for the crucial variables should be examined. The schematic decision tree in Figure 8 is a way of organizing this task. If the full tree were drawn out, it would contain 5,833 paths. Each path represents one specific setting of the decision and state variables.

---

* That means not getting the minimum acceptable return (of 10%) on the investment.

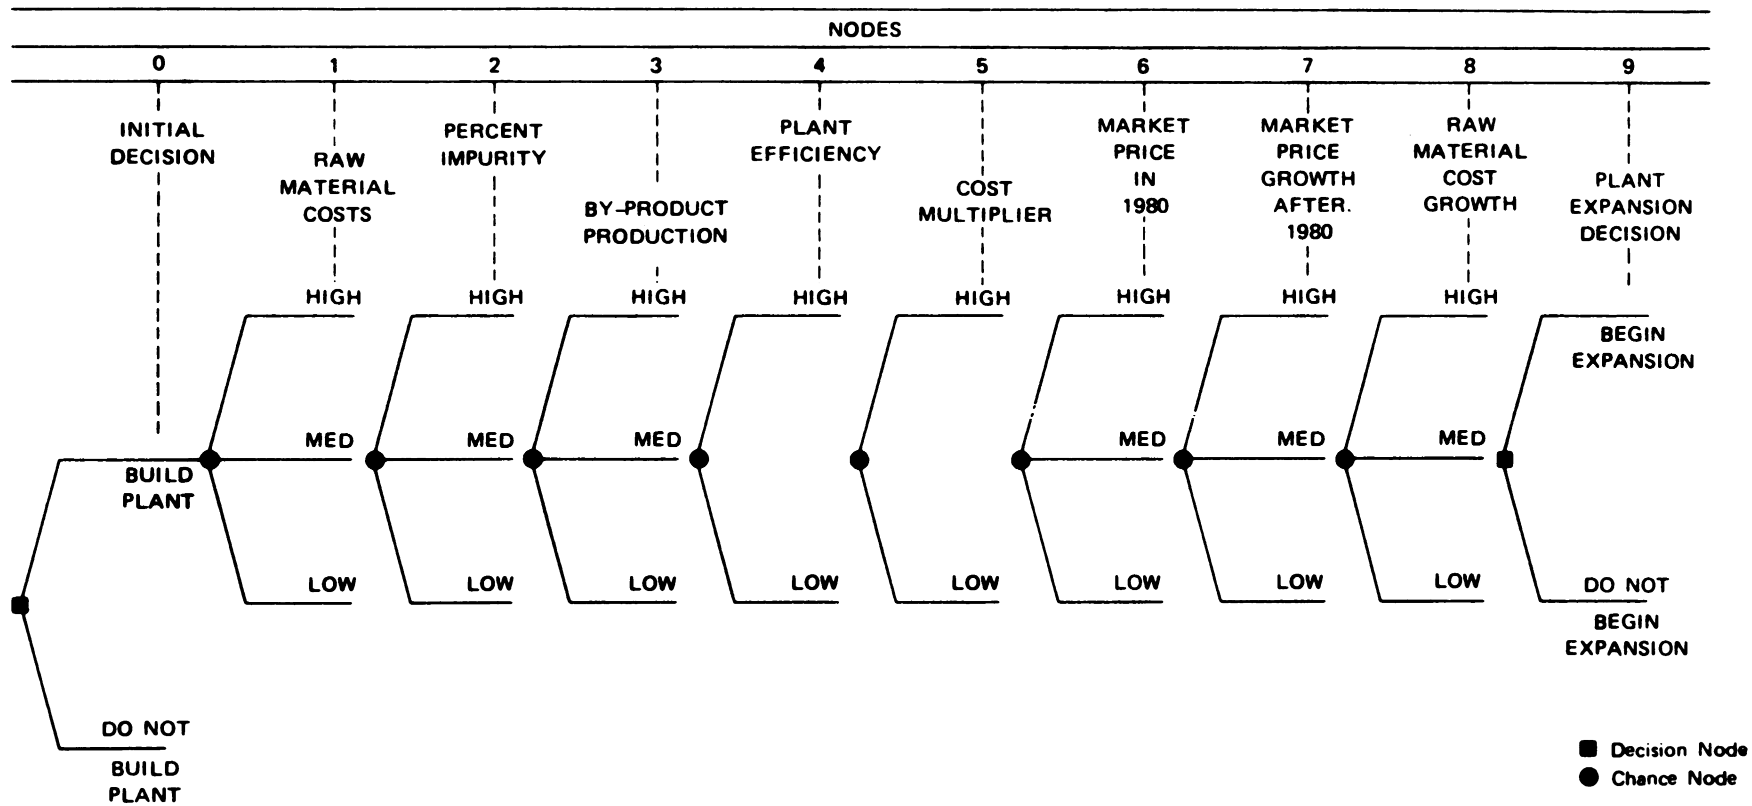

FIGURE 8 DECISION TREE STRUCTURE

Now let us see how the encoded information about the uncertainty can be transferred onto the tree structure. The step is quite simple. Based on the sensitivity to a variable, we decide how many branches on the tree should represent the variable. Since the tree grows very rapidly, the number should be held down. For most of the variables three branches were used. To represent the distribution for raw material costs by only three branches, we need to use the three step approximation for the whole curve. The steps shown in Figure 9 represent a 15% chance of \$6.30/ton, a 55% chance of \$13.25/ton, and a 30% chance of \$17/ton. Table 3 gives a list of the values used for the crucial state variables in the tree.

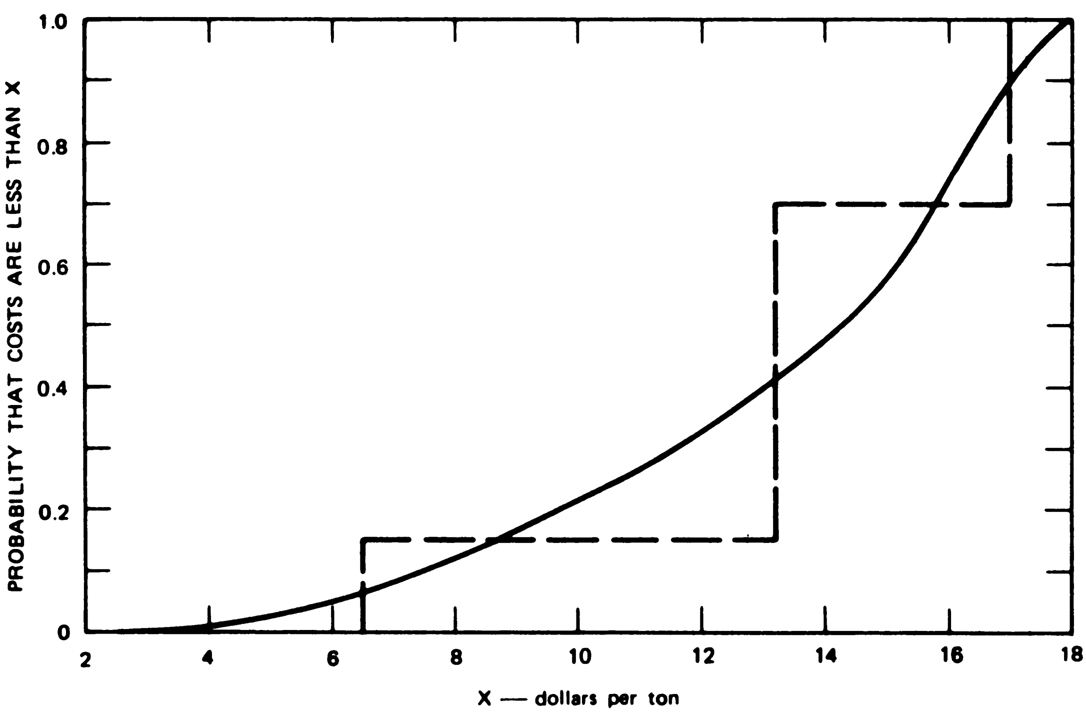

FIGURE 9 THREE-STEP APPROXIMATION OF CONTINUOUS DISTRIBUTION

By means of the tree structure and the information in Table 3 we can now completely specify the 5,833 combinations in the tree. We can also calculate the probability of each specific path by multiplying all the probabilities along it. Thus one path on the tree would represent the following (as shown in the tabulation on the page following Table 3).

Table 3

LIST OF BRANCH VALUES

	Node	Branch	Probability	Branch Value		
Raw material costs	1	1	0.15	$6.30/ton		
		2	.55	$13.25/ton		
		3	.30	$17.00/ton		
Impurities	2	1	.25	2.70 lb/ton		
		2	.55	4.10 lb/ton		
		3	.20	5.20 lb/ton		
By-product production	3	1	.20	18.00 lb/ton		
		2	.30	30.00 lb/ton		
		3	.50	50.00 lb/ton		
Plant efficiency	4	1	.85	82.5%		
		2	.15	102.5%		
Cost multiplier on investment	5	1	.75	90%		
		2	.25	110%		
Market price of brightener in 1980	6	1	.25	$0.18/lb		
		2	.50	$0.25/lb		
		3	.25	$0.31/lb		
Brightener price growth after 1980	7	1	.10	2%/yr		
		2	.60	4%/yr		
		3	.30	6%/yr		
Raw material cost growth		If node 7 branch is		(1)	(2)	(3)
	8	1	.30	1	2	4%/yr
		2	.40	3	4	6%/yr
		3	.30	5	6	8%/yr

	Probability	Branch Value
High raw material costs	0.30	\$17.00/ton
High impurities	0.20	5.20 lb/ton
High by-product production	0.50	50.00 lb/ton
High plant efficiency	0.15	102.5%
High cost multiplier on investment	0.25	110.0%
High market price of brightener in 1980	0.25	\$0.31/lb
High brightener price growth after 1980	0.30	6.0%/year
High raw material cost growth after 1980	0.30	8.0%/year

By multiplying the probabilities, the probability of this specific branch becomes less than one chance in ten thousand.

With the specific values for all the inputs specified, the deterministic model can be used to calculate the Present Value for the specific situation. In the case of the above path the Present Value is \$54.7 million. So the evaluation of the path results in a PV and the chance of occurrence of that PV.

The profit lottery is a summary of the result of evaluating all 5,833 paths on the tree. This was done with a computer model that combined the deterministic model with the decision-tree structure.

Risk Preference

Whenever a decision has to be made whether or not to accept the level of risk in a project, the question of risk preference enters. Risk preference is the decision maker's attitude toward risk. The risk preference of the decision maker can be (and should be) measured independent of any specific project. It requires answering the following type of question: "Would the decision maker be willing to invest in an opportunity that has an 80% chance of gaining \$50 million versus a 20% chance of losing \$20 million?" If the answer is yes, then: "What if a second alternative exists that would result in a certain (completely riskless) \$20 million? Which alternative would be preferred?" For any risky situation there is some riskless value which would make the decision maker indifferent. This value is called the <u>certain equivalent</u>.

The certain equivalent depends on the risk preference of the decision maker. In fact, once a decision maker's risk preference has been measured, the certain equivalent can be calculated. Stating a specific risk preference is a policy issue. The best that could be done in this analysis was to show the effect of different risk preferences.

Table 4 gives four specific lotteries. The first lottery was presented above--an 80% chance of gaining $50 million with a 20% chance of losing $20 million. The column headed "Expected Value" gives the expected values for each lottery. The expected value is the certain equivalent for an individual who acts toward each project as if he could repeat it numerous times. Such an individual is often called risk neutral. In the first lottery the expected value (long-run average) is $36 million. Obviously, since the decision maker cannot count on the long-run averages on such major investments, the decision maker's certain equivalent would be less than $36 million. But, how much less? Each column of certain equivalents in Table 4 represents a specific type of risk aversion and should serve only as a reference. The policy makers need to select a specific level of risk aversion. This can be accomplished by considering such simple reference lotteries as in the table. (Risk aversion increases from risk neutrality to Level 3.)

Table 4

REFERENCE LOTTERIES FOR RISK AVERSION

Lottery Probability	Outcome (millions of dollars)	Expected Value* (millions of dollars)	Certain Equivalents† at Different Levels of Risk Aversion (millions of dollars)		
			1	2	3
0.8	$50	$36	$31.5	$20.4	$ 0
0.2	-20				
0.5	50	25	22.0	16.4	8.4
0.5	0				
0.7	50	35	32.2	26.2	14.5
0.3	0				
0.6	20	8	6.9	4.7	-0.1
0.4	-10				

* Risk neutral.

† Based on an exponential utility curve. For a technical discussion of certain equivalents and risk aversion see References (2) and (3).

The level of risk aversion affects the plant investment decision in the following way:

	Risk Neutral	Level of Risk Aversion 1	2	3
Certain equivalent for plant investment decision (millions of dollars)	$11.7	$8.5	$2.0	$-15.5

Unless the decision maker wishes to act more risk averse than Level 2 in Table 4, the certain equivalent will be positive.

Determining the Best Action Based on the Initial Level of Information

What is the best choice, if it had to be made without further information? The answer to this question depends on the preferences of the decision makers. Since the preferences were not measured in this study, the effect of different preferences was investigated. The effect of risk aversion was discussed in the previous section. Next we will consider the effect of time preference; and then the combined effect of time preference and risk preference.

Figure 10 shows the shift in the profit lottery due to a change in discount rates. An increase in the discount rate from 10% to 13% leads to a steeper curve, since the gains as well as the losses are discounted more. The opposite effect is shown for a reduction in the discount rate. Rather than analyzing these profit lotteries, let us look at the effect on the certain equivalent. Table 5 shows the combined effect of discount rate and risk aversion. Note that for all discount rates of 7% or 10% and risk aversion Levels 1 and 2, the certain equivalent is positive. The boxed numbers in Table 5 represent a reasonable range of attitudes for policy makers of a corporation of the decision maker's size. Depending on the actual preferences of the decision maker, the action of opening the plant should be considered as having an equivalent guaranteed value of from $2 million to $17 million. Thus the plant looks like a good opportunity, even when adjusted for reasonable time and risk preferences.

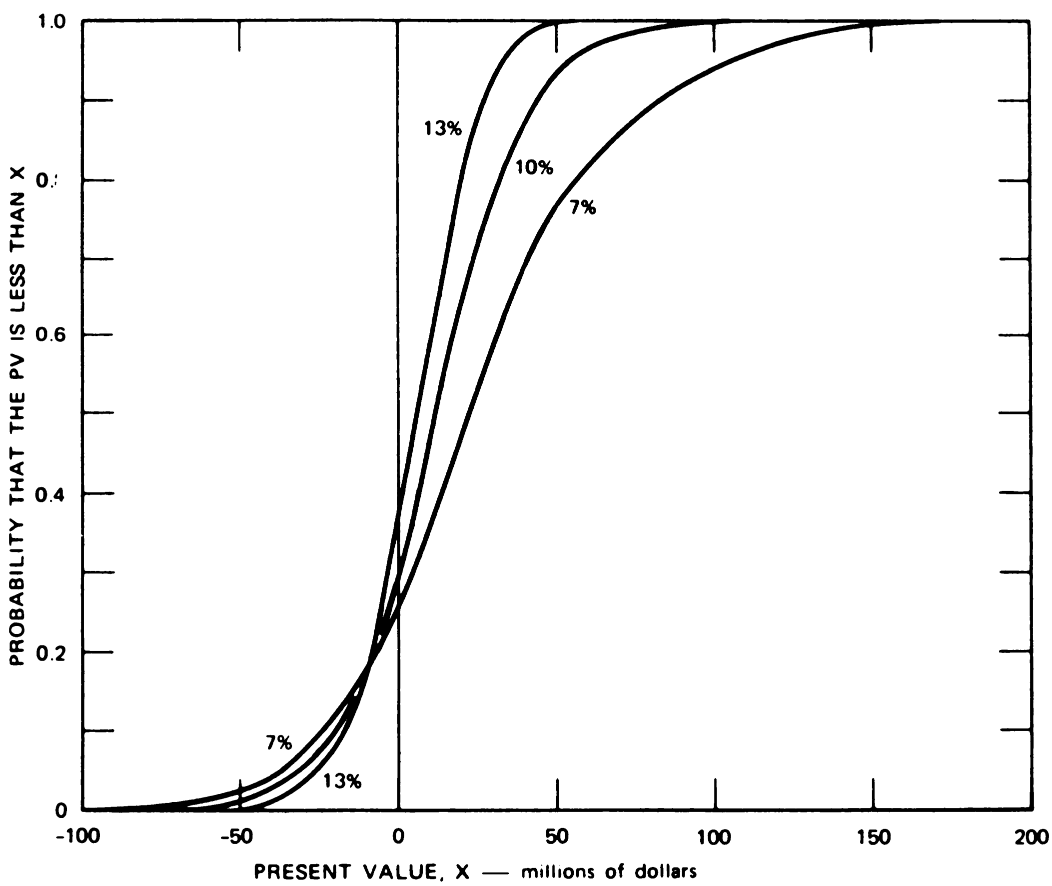

FIGURE 10 PROFIT LOTTERY SENSITIVITY TO DISCOUNT RATES

Table 5

COMBINED EFFECT OF DISCOUNT RATE AND RISK AVERSION ON THE CERTAIN EQUIVALENT

Discount Rate	Risk Neutral (millions of dollars)	Certain Equivalent Risk Aversion Level (millions of dollars) 1	2	3
7%	$25.6	$16.9	$1.7	$-46.6
10	11.7	8.5	2.0	-15.5
13	4.8	3.3	0.1	-8.3

Probabilistic Sensitivity Analysis

It is difficult to comprehend the results of a sensitivity analysis by directly examining profit lotteries. Therefore, let us look at the effect of changes on the certain equivalent. For the basis of comparison, the overall lottery will be evaluated at $2 million certain equivalent. (This corresponds to the 10% discount rate, and risk aversion Level 2.)

Because a tree structure was used to represent the uncertainty in the crucial variables, it is possible to answer the following question: "What effect would it have on the certain equivalent if we only looked at that section of the tree that represents the case of low raw material costs?" (The answer still includes the uncertainty in the other variables.) In Table 6 the effect of looking at only parts of the tree is tabulated.

From Table 6 the following conclusions can be drawn.

- Determining the specific raw material costs before making the decision will not greatly affect the decision.
- Impurities greatly affect the value of the investment.
- The quantity of by-product produced has the most important effect on the certain equivalent.
- The efficiency of the plant would not change the decision by itself.

The above results immediately lead into Section V, the value of additional information.

The Value Added by the Plant Expansion Decision

In the analysis of the decision tree, the plant expansion decision was included as the last node on the tree. Whether the plant expansion was included depended on whether the plant opening appeared profitable. By checking the output of the tree, it was found that a 25% chance exists that the plant expansion will be profitable after the initial plant was successful. The plant expansion added $1.9 million to the expected PV

Table 6

FURTHER SENSITIVITIES
(Overall Certain Equivalent of Lottery = $2.0 Million)

Variable	Probability	Branch Value	Certain Equivalent (millions of dollars)
Raw material costs	0.15	$6.30/ton	$ 3.4
	0.55	$13.25/ton	3.2
	0.30	$17.00/ton	-0.8
Impurities	0.25	2.70 lb/ton	15.7
	0.55	4.10 lb/ton	4.4
	0.20	5.20 lb/ton	-13.9
By-product production	0.20	18 lb/ton	-19.5
	0.30	30 lb/ton	2.7
	0.50	50 lb/ton	16.3
Plant efficiency	0.85	82.5%	0.5
	0.15	102.5%	12.0

and $1.6 million to the certain equivalent. Therefore the expansion decision is a major part of the total certain equivalent of $2 million from the initial plant opening. Favorable prices and costs could result in very large cash flows for the expansion decision. These cash flows are, however, greatly reduced by discounting.

V DECISION ANALYSIS--THE INFORMATIONAL PHASE

The informational phase follows the probabilistic phase as indicated in Figure 11. In the informational phase, the uppermost question is: "Should more information be gathered before acting on the main decision?" The value of gathering additional information can be measured by evaluating a revised decision tree. The essential change in the tree is the postponement of the decision point until after information has been included.

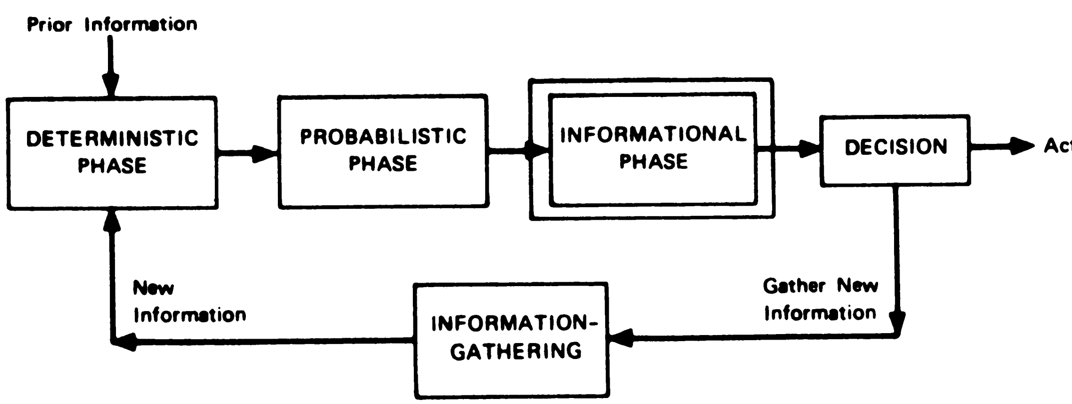

FIGURE 11 DECISION ANALYSIS CYCLE—INFORMATIONAL PHASE

The Value of Perfect Information

A special case is the value of perfect information. Here we investigate the effect of completely removing uncertainty in a variable before having to make the decision. This calculation is straightforward and can be carried out without any additional modeling effort.

If getting better information regarding a variable does not change the decision, that information has no value. Recall that (Table 6) removing uncertainty from plant efficiency would not change the decision. Therefore even perfect information regarding this variable alone would not be of any value. However, if the by-product production results were known, there is a 20% chance that a negative $19.5 million certain equivalent could be avoided, since the investment could be abandoned. Thus perfect information would be valuable regarding this variable.

The value of perfect information for several variables is given in Table 7. Removal of uncertainty on all variables should be valued at \$11.0 million. This includes such factors as prices in 1980. It is highly unlikely that much of this uncertainty can be removed. Note from Table 7 that the greatest value of information is in reducing the uncertainty in the ability to achieve reasonable production quantities of by-product. The next most important variable is the amount of impurities in the raw material. Reducing the uncertainty in raw material costs yields little value. The value of perfect information for plant efficiency is zero, as was expected.

Table 7

VALUE OF PERFECT INFORMATION

	Millions of Dollars
Total perfect information	\$11.0
Single variables only:	
By-product	6.2
Impurities	3.9
Raw material costs	0.3
Plant efficiency	0.0

What Additional Information Should Be Gathered?

Until now we have discussed the value of information without any considerations of the cost of information. Obviously perfect information is not available at any cost.

In the analysis, alternative information gathering programs were not explicitly evaluated. In fact that would be the next step in the cycle. However, the need for such an evaluation is not great, since the value of perfect information has identified areas of possible importance. In particular a high value of information was identified for determining the by-product production limits and for determining the impurities found in an average batch of raw materials. At the time of the analysis there were a number of steps toward getting such information that could be taken without great cost.

One cost that must be considered is the cost of postponing development for production, if indeed the development has to be postponed by the information gathering process. An approximate cost to postponing can be assigned by assuming that the identical cash flow would result after an initial delay. If that were the case a one year delay would simply reduce each PV by the discount factor. The average PV reduction is $1.2 million. Because of the cost of postponing, means for shortening the time to the plant opening should be considered.

Once additional information has been gathered, the analysis can be revised in a short period, since the computer models already exist. The decision analysis cycle can be repeated again, resulting in further recommendation for gathering information, building the plant, or discontinuing this investment effort.

REFERENCES

1. Howard, R. A., "Decision Analysis: Applied Decision Theory", Proceedings of the Fourth International Conference on Operational Research (D. B. Hertz and J. Melese, Eds.). New York: Wiley, 1968

2. Howard, R. A., "The Foundations of Decision Analysis", IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, September 1968.

3. Matheson, J. E., "Decision Analysis Practice: Examples and Insights", Proceedings of the Fifth International Conference on Operational Research (John Lawrence, Editor). London: Tavistock Publications, 1970.

4. Spetzler, C. S., "The Development of a Corporate Risk Policy for Capital Investment Decisions", IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, September 1968 .

5. Spetzler, C. S., "Establishing a Corporate Risk Policy", AIIE 1968 Proceedings.

11

STRATEGIC PLANNING IN AN AGE OF UNCERTAINTY

Michael M. Menke
Strategic Decisions Group

Reprinted with permission from Long Range Planning,
Volume 12, August, 1979, pp. 27-34.

Strategic Planning in an Age of Uncertainty

Michael M. Menke

Strategic planning and decision making in the face of uncertainty have always presented a serious challenge to top management, but the present scale of uncertainty is unprecedented. Decision makers used to be able to rely on the past to predict the trends of the future. Now they are increasingly being forced to make important decisions that depend upon highly uncertain external factors for which the past offers little guidance. In all areas of business and government, there is a vital need to understand and implement procedures that enable decision makers to deal more effectively with uncertainty for planning and allocating their organizations' resources.

Because of differences in their social, political and economic environments, European and Japanese managers are today affected by many acute areas of uncertainty—such as industrial democracy, floating exchange rates, changing social and political values, growing environmental awareness, government regulation, technological change, pollution control regulation, energy cost, and raw material availability—earlier than their counterparts in the U.S.A. These uncertainties affect not only private sector manufacturing industries, but also financial and service industries as well as nationalized industries and government organizations.

This article shows how ineffective methods of dealing with uncertainty can lead to serious mistakes with costly consequences. The cost of overconfidence and people's natural but futile tendency to ignore or to try to eliminate uncertainty is illustrated by the crises seen recently in the world steel and shipbuilding industries. The article then demonstrates how decision analysis procedures that focus directly on the major decision points in the strategic plan enable executives not only to include uncertainty directly in their strategic planning, but also to increase their understanding of the decision process and their ability to communicate the results to others.

The Age of Uncertainty

Earlier in this decade, *Fortune* published a series of articles entitled *The U.S. Economy in an Age of Uncertainty*[1]. The unifying theme of this series was the effect of unprecedented uncertainty on economic growth in the 1970s. Among the fundamental areas of uncertainty discussed were the future course of technology, changing values, attitudes toward pollution and population, future productivity, and the desire for better and more responsive government. The first article in this series showed that high and low projections for the Gross National Product of the United States in 1980 could differ by the staggering figure of $500bn, or even more, based on scenarios that were entirely plausible in 1970. The range of such projections illustrates the magnitude of the uncertainty that now confronts decision makers.

Subsequent events have shown that this series of articles was indeed prophetic. The uncertainties they discussed, plus many others they did not, have indeed made life more difficult for decision makers in organizations all over the world. The recent major crises in the world steel and shipbuilding industries have demonstrated the costs that can be associated with guessing wrong. A retrospective examination shows that in both of these industries, little heed was paid by decision makers in many major organizations to the possibility that actual market developments might turn out to be substantially different from the high growth patterns assumed in the early 1970s. These optimistic assumptions were then used to justify massive investment, which by the late 1970s were in many cases generating equally massive, and well publicized, losses—in some cases over £1m per day.

For important corporate and governmental decisions, the greater the uncertainty in the critical assumptions and value judgments that affect the decision, the more likely it is to be referred to the highest level of authority. For several reasons, however, top management generally has not dealt explicitly with uncertainty in the past. First, successful executives have usually developed an intuitive grasp of the *economic* uncertainties in their businesses during their careers. Second, there was no language for *precise* communication about uncertain events. Third, most approaches for dealing with uncertainty were highly mathematical in appearance but offered minimal insight into the decision.

Nevertheless, these reasons provide only a meagre excuse for ignoring the powerful methods now available to assist executives in making the extremely complex decisions of the future. Moreover, a recent important development offers executives the ability to meet the challenge of uncertainty head-on: the creation of a comprehensive, understandable language and a philosophical framework for treating decision problems involving uncertainty. These methods are practical, well established, and comprehensible to people throughout the spectrum of management responsibility.

The Danger of Using Single Number Estimates

Past attempts at predicting too precisely the outcome of uncertain factors have sometimes led to disastrous consequences. Nonetheless, the results of nearly all current prediction and forecasting processes--whether based on time series extrapolation, econometric modelling, environmental analysis or pure judgment—are single point forecasts for use in subsequent planning and decision making. The Suez Canal provides an interesting example of the perils that can be involved in basing decisions only on a single best guess forecast. In mid-1967, when the Suez Canal was blocked, most major oil companies were forced to reallocate significant resources for transporting oil from its origins to its markets. Typical actions considered included chartering additional tankers, ordering supertankers, and building pipelines, as well as searching for alternative sources of supply. All of the commitments made and contracts signed depended strongly on management's judgment of when, if ever, the Canal would be reopened.

Realizing this, the management of one major oil company sought the best estimate of when the Canal would reopen. After a thorough study, the management decided to use a July 1969 opening as the corporate planning estimate for use in calculating the profitability of all subsequent actions. This directive had significant effects on the alternatives considered, as well as on the length of charter contracts signed, the charter rates accepted, the type of new tankers ordered, and the alternative supplies sought. With all its future planning based on a single best guess as to when the Canal might reopen, imagine the company's situation in 1969 when it found its charters expiring, tankers in short supply, higher rates prevailing, its competitors equipped with supertankers, and no prospect for the Canal to be reopened in the immediate future.

On the other hand, many oil companies and shipowners essentially disregarded the possibility that the Suez Canal might ever reopen and subsequently heavily over-commited themselves in the supertanker area. This has subsequently proved to be a disastrous situation for many of the companies concerned, several of which--although large--were nearly bankrupted. Although many other changes in the operating environment of the shipping industry contributed to its depressed condition during the mid-70s, the inability to cope effectively with this environmental uncertainty has clearly been a major contributing factor to the poor results stemming from many of these decisions.

Probability: The Language of Uncertainty

More comprehensive methods of analysis cannot foretell the future, but by allowing a more complete expression of the uncertainty surrounding future events, they can frequently point out more flexible and thus less risky alternatives. The key to coping with uncertainty in strategy formulation is for managers to discriminate among alternatives whose consequences have widely different degrees of uncertainty, allowing them to search for the most profitable plan when uncertainty is low and to 'pay' for a suitable degree of flexibility where uncertainty is high.

What type of statement can be made to capture the complete judgment of management, including the full range of uncertainty, regarding a specific event, such as the reopening of the Suez Canal? The only logical way to express a complete opinion involving an uncertain event is to use the language of probability.[2] Such an expression does not require extensive knowledge of probability theory on the part of the executive or expert providing the information, but it does subsequently allow the full power of probabilistic methods to be applied to solve the problem. Furthermore, it provides a new way to communicate the judgment—and resolve the differences of judgment—throughout the management hierarchy.

Figure 1 shows how management's view of the future of the Suez Canal could have been expressed in terms of a cumulative probability distribution compiled in July 1967. Each point on the curve represents a judgment of the probability that the Canal will reopen by the date on the horizontal axis.* For example, they believed that there was no chance (0 probability) that the Canal would reopen before mid-1968. They thought there was a 0·50 probability (50 per cent chance) that the Canal would reopen before the beginning of 1970, a 0·75 probability (3 to 1 odds) that it would reopen before mid-1972, and a 0·85 probability that it would reopen before the beginning of 1975. Conversely, they felt that there was a 0·15 probability that it would not reopen until after 1975 and about 0·05 probability (1 chance in 20) that it would never reopen!

This graph summarizes many similar statements. The curve is steepest during the first half of 1969, which means that they felt that the Canal was most likely to reopen during that period. The graph shows that there was a 0·30 probability that the Canal would reopen between 1 January and 30 June of 1969, compared with equally low 0·10 probabilities that it would reopen during the 6 months preceding or following that period.

*Practical methods to encode subjective judgment regarding uncertain events and to express the information through probability distributions such as Figure 1 have been developed and used extensively by SRI (see reference 3).

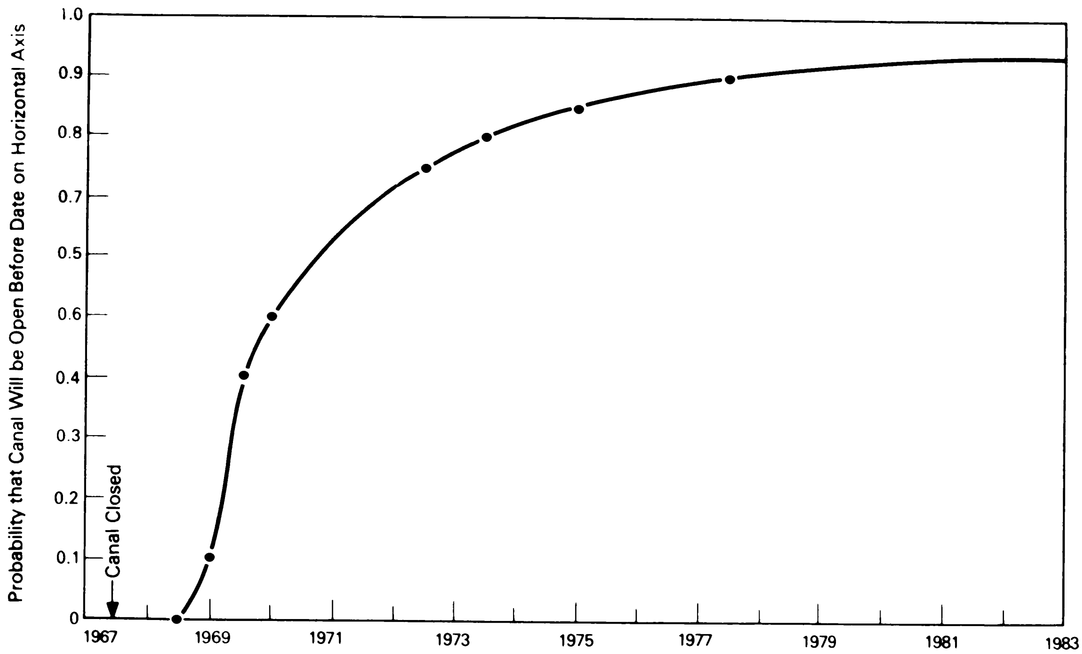

Figure 1. Uncertainty in date by which the Suez Canal will be reopened (compiled in July 1967)

After 1969, the curve flattens out considerably. This could mean that they felt that if the blockage persisted longer than 2 years, alternative transportation methods would become firmly established, and the economic pressure to reopen the Canal would greatly diminish.

Using this curve as a forecast would have resulted in the choice of a very different strategy from the one actually used. This strategy would have taken into account the 0·50 probability that the Canal would not be reopened before 1970 and the 0·20 probability that the Canal would still be closed in mid-1973. It would probably have been a hedging strategy involving a totally different mix of options and commitments than those based on the single number estimate of July 1969.

As we know today, the Suez Canal was actually reopened in 1975. Given this actual result, what conclusion can we draw about the quality of the judgment expressed by the curve in Figure 1? It can be described as reasonable, since the curve in fact indicated a probability of 0·15 that the actual opening date would be even later than 1975. However, one could also argue, since the curve indicated that there was a 50 per cent chance for the Canal to open 5 years or more before it actually did, that the judgment expressed was not very helpful. It must be clearly understood that in unique situations like this that the *a priori* curve can never be proven to have been 'right' or 'wrong' by the outcome.

On the other hand, considerable evidence has been accumulated during the past decade that shows subjective probability has provided a valid and statistically reliable forecasting method for repetitive uncertain events as diverse as interest rates and R + D project success. Morgan Guaranty Trust Company has used the probability method to forecasting short-term interest rates since the early 1970s and has compiled an enviable track record.[4] In fact Ralph Leach, Chairman of the Bank's Executive Committee said several years ago[5] 'The (method) has been a remarkable success so far. . . I can't recall missing a 90-day move since we put it into use.'

In a quite different application, Sandoz AG has used the probability method to forecast the technical success of pharmaceutical R &D projects as early as 6 years before completion and to take appropriate management action in light of these forecasts. An analysis of their experience[6] has shown that the probabilities they have assigned are statistically valid and therefore provide a reliable basis for R &D management decisions. Examples like these indicate that experienced managers can use probability methods to express themselves more clearly and reliably about uncertain events than has been possible before. Moreover, the management action based on these probability assignments appears to have been quite effective in achieving profitable results.

The message of these examples should be clear. Although we cannot decide what will happen in the future, we can nonetheless wisely allocate our resources today in light of the uncertainty that we perceive. We have done a good job when our actions combine with uncertain forces beyond our control to increase our likelihood of achieving desirable outcomes. As the impact of uncertainty increases, more precise ways of specifying and treating uncertainty will be required than those to which management has been accustomed in the past.

Strategic Planning and Decision Analysis

Among the many existing definitions of planning, the definition preferred by SRI is 'the network of *decisions* that directs the intent, guides the preparation for change, and programs action designed to produce specified results'. The crucial importance of decisions in this definition is not accidental, because the key to a more explicit and rewarding treatment of uncertainty in strategic planning is to focus on the major decisions that comprise the strategic plan. The approach that SRI recommends to do this is called decision analysis.[7–9]

Decision analysis combines the ability of decision theory to handle uncertain problems with the ability of systems analysis, operational research and modelling methodology to deal with complex and dynamic problems. Decision analysis, however, is more than a specific technique or method. A recent *Harvard Business Review* article described it as 'a major area of management—probably the most revolutionary advance in management practice in many years'.[10] Decision analysis applies all available rational and logical means to solve difficult management decision problems. Thus, decision analysis shapes methodology to fit the problem, rather than attempting to make the problem fit a specific methodology.

Although the foundations of decision analysis go back several hundred years to the beginnings of probability theory, the means for dealing with the complexity of real problems was not available until the recent advent of reasonably priced computational equipment. Much of the development of decision analysis into a feasible and effective management discipline has resulted from applications performed by SRI's Decision Analysis Group over the last dozen years. The rest of this article explains the basic concepts that allow decision analysis to treat uncertainty and complexity together and then discusses the application of decision analysis to an actual shipping decision.

The Conceptual Framework of Decision Analysis

Any logical treatment of uncertain problems demands a clear distinction between decisions and outcomes. In an uncertain environment, we can prescribe only our decisions; not their outcomes. Ford's Edsel and Du Pont's Corfam are often cited as classic examples of bad decisions. While we can all agree that these decisions produced unfortunate outcomes, only an understanding of the logical context in which these decisions were made would reveal the quality of the decisions themselves.

Suppose that Ford's management had in fact thought the Edsel had only a 0·20 probability of being the last highly profitable intermediate-size car and therefore a 0·80 probability of failure in the marketplace. If they further thought Edsel's failure would signal that the time was ripe for compact cars and had a back-up plan to convert Edsel facilities to produce the Mustang for the compact market, then Ford may have made a very good decision in the context of a broader strategy. It is not possible to audit the quality of unique, long range decisions on the basis of outcomes; decisions are good if they are logically consistent with the information and preferences of management at the time when action was required. A major part of this consistency is how the decisions fit into the overall strategic plan.

An important feature of decision analysis is bounding the problem to a well-defined and manageable scope. The central requirement in the initial stage of analysis is to identify exactly what decision must be made and who is responsible for the resources required to implement this decision. This necessitates a detailed enumeration of the perceived alternatives. Without alternatives, there are no decisions—only worries. Establishing a hierarchy of decisions is one technique developed by SRI to help clarify and demarcate decision problems. Figure 2 shows a decision hierarchy that was instrumental in establishing

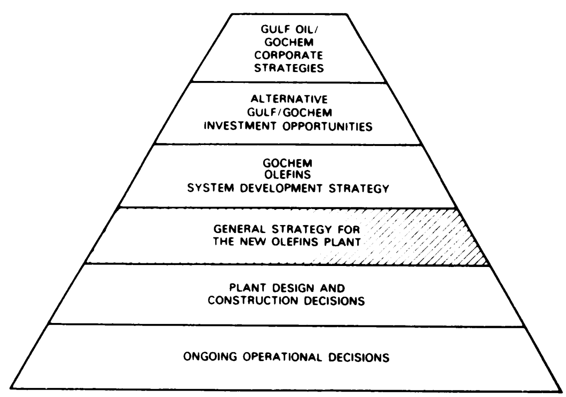

Figure 2. Hierarchy of new olefins plant decisions

the decision focus for the analysis of a new olefins plant strategy for Gulf Oil Chemicals Company (GOCHEM) in the early 1970s.[11] A decision analysis can be performed to allocate resources at any level of the hierarchy, but once a specific decision problem has been identified, decisions at other levels of the hierarchy should be treated only with regard to their impact on the decision at hand. This decision focus clearly differentiates the role of decision analysis from corporate and strategic planning, on the one hand, and from technological and market forecasting on the other.

The application of decision analysis often takes the form of an iterative procedure called the decision analysis cycle. This procedure is not an unviolable method of attacking the problem, but it provides a means of ensuring that essential steps have been considered. The procedure can be divided into three phases, as shown in Figure 3. The deterministic phase defines the variables affecting the decision and relates them to a decision

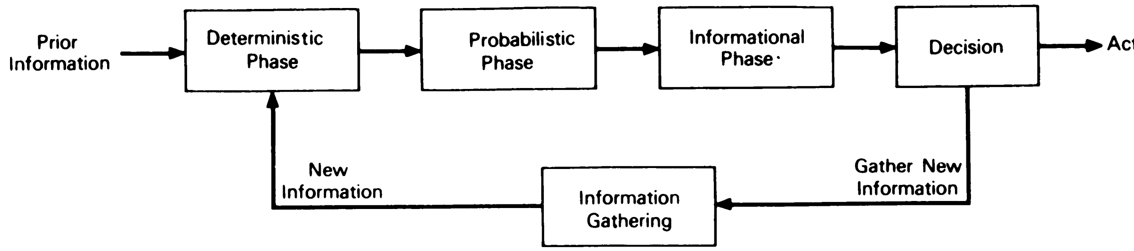

Figure 3. The decision analysis cycle

criterion (usually a measure of profitability like present worth), assigns values to these variables, and measures the importance of the variables through extensive sensitivity analysis. Most problems can be solved at this stage without explicit consideration of uncertainty, although a proper sensitivity analysis must take into account the relative predictability of different decision factors.

The probabilistic phase encodes probability distributions[3], such as that in Figure 1, for the most important variables and then derives for each alternative the associated probability distribution over the decision criterion. This phase also introduces the assignment of risk attitude[12,13] which defines the best solution in the face of uncertainty. Risk attitude measures management's aversion to receiving the worst possible outcomes instead of the best, given the uncertainties present. It is expressed in terms of a certain equivalent value, which is less than the expected value of the probability distribution over the decision criterion by an amount that depends on the organization's ability to tolerate losses.

The informational phase reviews the results of the first two phases to determine the economic value of eliminating uncertainty in each of the important variables in the problem. In some ways, this is the most important phase because it shows just what it costs in monetary terms not to have perfect information. A comparison of the value of information with the cost and reliability of any potential source then indicates whether additional information should be collected or whether it is more profitable to act now.

An Application of Decision Analysis to a Transportation Decision

The following example illustrates how decision analysis deals with uncertainty. A major integrated mining, refining, and marketing company, MRM, had to choose among several alternatives for transporting unprocessed ore from its mines overseas to its centralized processing facilities. Historically, MRM had relied on vessels chartered from World Wide Shipping (WWS) to provide ore carriage. However, most of WWS's ships were old and obsolescent, and MRM expected that WWS would scrap these ships over the next 10 years. Due to structural changes in the shipping industry, it was not obvious to what extent WWS would replace the old ships and whether the replacements would be suitable for bulk transportation of ore. Furthermore, WWS was unwilling to commit itself to any contract containing specific terms beyond 6 years.

In light of these problems and in anticipation of a future shortage of appropriate ships, MRM was considering several alternatives for securing proprietary shipping capacity. Proposals, proposal evaluations, and negotiations had been underway for 2 years when the chairman of the board of MRM attended an SRI executive seminar on decision analysis. SRI's Decision Analysis Group was subsequently commissioned to extend the deterministic analysis performed by company personnel into a probabilistic analysis taking into account the principal uncertainties affecting this decision.

Four specific alternatives were considered for providing ore carriage over a 20 year period, to begin 3 years later:

(1) *Purchase*—purchase two new, large proprietary ore carriers.

(2) *Tug/barge*—postpone the decision for a year and a half and then decide whether to use proprietary carriers, chartered vessels, or two new tug/barge units that operate with the tug pushing at all times, locked into the barge slot by a patented quick release device. The postponement was required so that the tests could be performed to determine whether the interlocking device would operate according to specifications.

(3) *Charter*—utilize chartered carriage for 6 years under existing agreements with WWS and then continue to relay on WWS carriage under terms not yet negotiated.

(4) *Purchase/charter*—Hedge by purchasing only one vessel and chartering all other needed capacity. WWS would provide all additional transportation requested for 10 years, after which a similar amount of carriage would be provided by WWS but under uncertain terms.

For both the purchase and the tug/barge alternatives, circumstances peculiar to MRM dictated that no remunerative use of the ships was to be made when they were not used for ore carriage.

For each of the alternatives, the figure of merit used was the present value of cost. The financial structure of MRM was such that minimizing transportation cost would maximize corporate profitability. The corporate time preference for money, which expresses the corporate valuation of cash flows over time, was reflected in a discount rate. In the analysis, the probability distribution of the present value of cost—called a lottery for short—was calculated for each of the four alternatives. These lotteries amalgamate the uncertainty in the crucial variables in the problem and explicitly display the resulting uncertainty in total cost for each alternative.

Results of the Probabilistic Analysis

The main variables whose uncertainty was explicitly treated were subsidiary ore production, finance rate, WWS future operating costs, the ore carrier and tug/barge contract prices, technological development (whether the tug/barge interlocking device would operate according to specifications), labour stoppages, and total loss of a ship. For each of these factors, expert judgment was condensed into a probability distribution like Figure 1; these distributions were subsequently combined into a large number of scenarios with quantitative probabilities using a decision tree analysis.[14,15] The results of this analysis were the probability distributions on the present value of cost.

The expert judgment incorporated in the probabilistic analysis was provided by the subsidiary management, by mining engineers, and by MRM management personnel. The results reflect their state of information at the time of the analysis. The information encoding was done only to the level of detail that was economically justified. For example, the likelihood that the interlocking device would operate successfully was thought to be about 40 per cent. Because of the economics of the situation, it was not necessary to refine this figure any further. As long as the doubt about success necessitated a one and one-half year delay, the recommended decision would not change. On the other hand, a more detailed encoding was performed for this variable.

The cost lotteries derived from these assumptions are shown in Figure 4. Each lottery explicitly demonstrates the perceived uncertainty in the present value of cost. The vertical axis represents the probability that the present value of cost will be less than the amount given on the horizontal axis. For example, for the purchase alternative, there was a 0·72 probability that the present value of cost would be less than $124m. For each alternative, the expected* (or average) present value of cost is also indicated in Figure 4. The expected costs (in millions of dollars) for the four alternatives were:

Alternative	Expected cost
Purchase	$119·4
Tug/barge	$131·0
Charter	$140·6
Purchase/charter	$130·4

The purchase alternative clearly had the smallest expected cost of the four alternatives. Moreover, even when the uncertainty in all crucial variables was explicitly considered, the purchase proposal dominated the other proposals in such a way that for each level of transportation cost, it had the smallest chance of exceeding that cost level. This type of dominance, called stochastic dominance, does not guarantee that the purchase alternative will *actually* provide the lowest present value of cost; it simply demonstrates that the purchase alternative will always have the greatest chance of providing the lowest cost. The practical implication of stochastic dominance is that the purchase alternative would be superior for any corporate risk attitude. No matter how averse the corporation might be to taking risks, the proprietary fleet would always be preferred. The same conclusion held as the discount rate varied from 6 to 20 per cent.

The $11·6m difference between the purchase alternative and the tug/barge alternative reflected the cost of delaying the decision for one and one-half years while chartering at higher cost during that period. If the tug/barge system had been known to be operational at the time of the analysis, then it would have had roughly the same expected present value of cost as the purchase alternative.

Sensitivity Analysis

The results indicated that the purchase alternative would provide the most favourable transportation cost lottery. To check this conclusion, the results were subjected to many sensitivity tests. In each case, attention was focussed on one or two variables while all other variables were held fixed. In particular, the sensitivity of the results to the following factors was determined: labour stoppages; catastrophic loss of a ship; and variations in ore production, ore carrier price, bank financing rate, refined deliveries, WWS operating costs, and inflation rate. The best decision was insensitive to reasonable variations in any of the above factors, except for ore production.

Possible variations of ore production could alter the decision. For example, if production were to drop to a low enough level, one of the WWS proposals would have provided a lower expected cost. In the opinion of MRM production experts, there was only a 0·22 probability that production would decrease enough that the purchase/charter proposal would be superior to the purchase alternative. The assumptions underlying this judgment favoured the charter proposals, since they were based on the premise that MRM would not adopt

*The expected value of a lottery is computed by first multiplying the value of each outcome of the lottery by its associated probability of occurrence and then summing over all outcomes. The expected value computed in this fashion is identified with the average value that would be achieved if the uncertain situation were to be repeated a great many times. For continuous lotteries, such as Figure 4, the [illegible] becomes an integral.

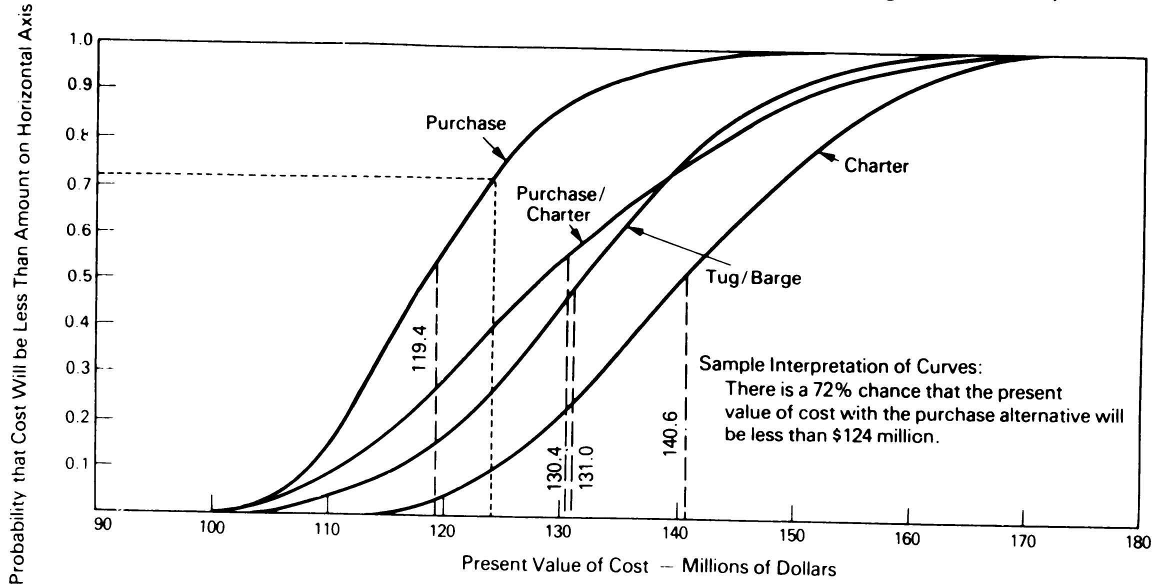

Figure 4. Cost lotteries for the four basic alternatives (Corporate discount rate of 10%).

new operating rules as production decreased. In fact, the reduced production would allow for scheduling changes and might allow handling alternative cargoes.

For other variables examined—ore carrier price, tug/barge price, bank financing rate, refined deliveries, WWS operating costs, and inflation rate—the sensitivity analysis showed that the purchase alternative was dominant for any reasonable variation. This conclusion held even for changes in these variables much larger than those considered possible by MRM management, and for joint variations in variables that might be expected to be correlated, such as WWS operating costs and inflation rate. Thus, although a great deal of uncertainty was present in many important variables affecting the future cost of each alternative, the best decision was clear. Thus the decision was robust.

Two sensitivities are of special interest. First, the catastrophic loss of a ship was examined and found to have no effect on the best decision. If a proprietary ship were lost, it would be insured, and so the cost of the loss could be calculated as the cost of chartering equivalent capacity while the lost ship was being replaced. Even for probabilities of loss ten times greater than the historical figure of one chance in a thousand, or for charter cost premiums that were double the highest conceivable figure, the purchase alternative was superior.

Second, labour stoppages could affect MRM in several ways, but most of these favoured the proprietary fleet alternatives. Only a strike of foreign dock workers, which would suspend all obligations to WWS under a *force majeure* clause, would differentially favour the charter alternatives. However, reasonable estimates, such as a four-week strike every three years, reduced the expected cost of the charter alternative by only \$3·5m. Since the purchase alternative had a \$21m expected cost advantage over the charter alternative, labour stoppages would be relatively insignificant. In short, the only conceivable circumstance that might change the best alternative was a sharp decrease in future ore production.

Informational Phase

In a decision analysis, both the analyst and the client should clearly focus on the decision at hand. In this case, the alternatives were well defined, and the decision was to select the alternative that would provide the necessary carriage at the lower cost, which in turn would mean the highest profitability. As previously mentioned, only the foreknowledge of significantly lower ore production in the future would result in the expectation that a different alternative would provide the lowest total cost. There was thought to be only a 0·22 probability of such a production decrease, and even in this case, the expected cost of the purchase/charter alternative would be only about \$2m less than the purchase alternative. Therefore, the expected value of *perfect* information about ore production was found to be 0·22 × \$2m, or \$440,000. Since any programme to provide further real—and thus imperfect—information concerning the ore deposits available would have cost a minimum of several hundred thousand dollars and would have taken considerable time, it was pointless *in the context of this shipping decision* to gather any additional information on future ore production.

It should be clearly understood that decision analysis can subject all information-gathering procedures—such as geological studies, core samples, pilot facilities, prototypes, and market surveys—to rigorous economic criteria. Information has value only if it could enable one to modify a present course of action to increase profits.

Thus, only information that can potentially change the decision under consideration has value. Using this standard, additional information on any variable other than future ore production had no value to the shipping decision. In spite of the great uncertainty shown in Figure 4 as to what the cost of the proprietary fleet would ultimately be, additional information was not justified because there was only a very small chance that it would lead to the recommendation of a different alternative.

Concluding Remarks

The principles of decision analysis apply to decisions at all levels. In the Figure 2 hierarchy, decision analysis can be used to determine whether to expand production capacity for a particular chemical product, whether to enter an entirely new business area or divest one that is not meeting expectations, or whether to merge the entire organization. At the new plant strategy level, important decisions are plant size, plant location, choice of technology[11] and whether to finance the plant via lease or purchase arrangements. At the highest level of the hierarchy, decision analysis has been used to derive the implications of management's desire for a smooth growth in corporate earnings per share in terms of proper measures to evaluate individual project decisions.

Decision analysis offers promise in any problem area fraught with uncertainty. Decision analysis has been used to evaluate new products, to define marketing strategies, to forecast pharmaceutical sales and determine the optimal production capacities[14], to devise a commodity purchasing strategy, to derive the value of information in mineral exploration[15], and to develop financial portfolio management models[5]. A different but equally pertinent application of subjective probability has been to plan and manage the R &D portfolio of a leading European pharmaceutical company, where the technical success of R &D projects as well as their future market potentials are highly uncertain.[6] Decision analysis has also been applied to many important public sector problems including national energy policy[16], nuclear reactor safety[17], environmental protection[17], weather modification[17] and health care planning.[18]

The example presented in this paper dealt with uncertainty involving a shipping problem of an integrated mining, refining, and marketing company. However, the investment decision presented could equally well have been that of any industrial organization, and the nature of investment could equally well have been in R &D, production capacity, or acquisition. The important point is that all areas of management responsibility need to, and can, develop procedures to assess the uncertainty inherent in the alternatives that confront them, to weigh this uncertainty against corporate risk attitude, and to select those alternatives most likely to contribute to the strategic objectives of the organization.

References

(1) The Editors of *Fortune*, *Challenges for Business in the* 1970*s*, Little, Brown and Co., Boston-Toronto (1971).

(2) M. Tribus, *Rational Descriptions, Decisions, and Designs*, see especially Chapter 1. Pergamon Press, New York (1969).

(3) C. S. Spetzler and C. S. Staël von Holstein, Probability encoding in decision analysis, *Management Science*, **22** (3), November (1975).

(4) I. Kabus, You can bank on uncertainty, *Harvard Business Review*, **54**, 95, May–June (1976).

(5) A. Dee, Probability bar graph adapted by Morgan for rate projection, portfolio strategy, *American Banker*, 21 March (1973).

(6) H. U. Balthasar *et al.*, Calling the shots in R&D, *Harvard Business Review*, **56**, 151–160, May–June (1978).

(7) R. A. Howard, Decision analysis: applied decision theory, *Proceedings of the Fourth International Conference on Operational Research* (1966), D. B. Hertz and J. Melese, eds., John Wiley and Sons, New York (1968).

(8) J. E. Matheson, Decision analysis practice: examples and insights, *OR69*: *Proceedings of the Fifth International Conference on Operational Research*, John Lawrence, ed., Tavistock Publications, London (1970).

(9) H. Raiffa, *Decision Analysis: Introductory Lectures on Choice Under Uncertainty*, Addison-Wesley, Reading, Massachusetts (1970).

(10) J. G. Milliken and E. J. Morrison, Management methods from aerospace, *Harvard Business Review*, p. 6, March–April (1973).

(11) W. C. Roher, Balancing risk and opportunity in commercial development, Speech to the Commercial Development Association, New York, (1977).

(12) R. A. Howard, Risk preference, *Readings in Decision Analysis*, SRI International, Menlo Park CA (1976).

(13) C. S. Spetzler, Establishing a corporate risk policy, *AIIE* 1968 *Proceedings*.

(14) R. Abt *et al.*, The dangerous quest for certainty in market forecasting, *Long Range Planning*, **12**, 52 (1979).

(15) P. D. Newendorp, *Decision Analysis for Petroleum Exploration*, Petroleum Publishing Company, Tulsa (1975).

(16) Steven N. Tani, Decision analysis of the synthetic fuels Commercialization Program, *AFIPS–Conference Proceedings*, AFIPS Press, Montrale, N. J., pp. 23–29, (1978).

(17) R. A. Howard, Social decision analysis, *Proceedings of the IEEE*, **63** (3), March (1975).

(18) F. Offensend and M. W. Merkhofer, Elements of a planning methodology for use by the national caries program, Final Report, SRI Project 4210, prepared for National Institute of Dental Research (NIH), Bethesda, Maryland, August (1976).

12

A DECISION ANALYSIS OF A PETROCHEMICAL EXPANSION STUDY

Terry J. Braunstein
Strategic Decisions Group

INTRODUCTION

In early 1973, Gulf Oil Chemical Corporation (GOCHEM) decided to build a major capacity addition to its olefins system at Cedar Bayou, Texas. This would be the first of many very large, feedstock- and product-flexible ethylene plants in the United States. GOCHEM's decision process was based on decision analysis -- a quantitative weighing of the many foreseeable risks against the promise of profitability -- and as such was a first for the petrochemical industry. The analysis explicitly modeled the actions of competitors long before it was fashionable and captured the complex price relationships between petroleum products and petrochemicals in a probabilistic structure. This paper relates the history of the Cedar Bayou plant decision and the underlying decision analysis process.

THE ETHYLENE INDUSTRY

Ethylene, a basic chemical commodity that forms the raw material base for many chemicals (such as polyethylene, ethylene oxide, vinyl chloride, and styrene) and many products, is one of the largest volume industrial chemical commodities.

The U.S. ethylene production industry, which is centered on the Gulf Coast of Texas and Louisiana, is serviced by an extensive pipeline system connecting producers and users. Ethylene is produced by the major chemical companies and also by the chemical divisions of the major oil companies. It is primarily used by the producer companies in their own downstream derivative businesses. The remainder is sold into the merchant market, and GOCHEM is one of the largest manufacturers of ethylene for this market.

Traditionally, ethylene was produced in the United States through cracking natural gas liquids, mostly ethane and propane, that come from refineries and from natural gas production. However, during the late 1960s, gas liquids production began to peak and ethane and propane were increasingly difficult to obtain. During this same period, the profitability of the U.S. chemical industry generally declined. Poor profit incentives, combined with considerable uncertainty about the source of new feedstocks for ethylene plants, resulted in an industry-wide reluctance to build new plants. However, the demand for ethylene continued to grow, and in the summer of 1972, a major ethylene shortage was predicted for 1975.

A Naphtha-Based Process

During 1972, GOCHEM considered expanding its ethylene facilities. Given the highly favorable projected market conditions, GOCHEM anticipated that it could successfully insulate itself against risk by carefully designed contracts. Since internal and external suppliers of ethane and propane could not commit sufficient quantities for supplying an economic size ethylene plant, GOCHEM considered using the next most common feedstock, naphtha, and process -- naphtha cracking. Naphtha cracking to produce ethylene has been practiced overseas because of both the unavailability of natural gas and because of naphtha's role as a surplus part of the "crude barrel." Gulf already ran plants in Europe, Canada, and the Far East that used the naphtha-cracking process.

Potential Problems with Naphtha

There were, however, a number of potential problems with using naphtha. Since naphtha is the basic constituent of gasoline, potential gasoline shortages in the United States threatened its long-term availability. In addition, it is an excellent feedstock for producing synthetic natural gas (SNG), and a projected natural gas shortage increased interest in SNG plants. Furthermore, government regulations controlling the importation of foreign naphtha were uncertain. Naphtha's value also could be substantially affected by legislation regulating lead and aromatic levels in gasoline.

There were also other potential problems with naphtha. Being considerably heavier than the natural gas liquids, naptha contains a larger fraction of long-chain carbon molecules. Consequently, naphtha cracking to produce ethylene yields more by-products than does cracking natural gas liquids. GOCHEM was not in many of these by-product businesses, and market values for the by-products were quite uncertain. Moreover, if naphtha-based ethylene plants were built by the whole industry, the marketplace could have been flooded for some of these by-products.

GOCHEM'S ETHYLENE PRODUCTION PLANS

Despite these uncertainties, GOCHEM decided in 1972 to proceed with the planning of a naphtha-based ethylene plant. Given the uncertain situation, the company conservatively decided on a 1 billion pound-per-year plant using proven and reliable equipment. Larger plants existed, but engineering studies showed no further significant economies of scale.

While GOCHEM believed it could sell 1 billion pounds of ethylene annually, it was unsure of the marketability of the ethylene by-products. The company also had some doubts about investing in a business with recent low returns. These concerns prompted GOCHEM to obtain outside advice.

In January of 1972, GOCHEM's top management sponsored an in-house seminar on decisions analysis -- a methodology for making complex, risky decisions -- and GOCHEM decided to explore it further as a method for dealing with its proposed ethylene plant. As a result, GOCHEM asked for an analysis of the decisions surrounding the new ethylene plant venture before requesting authorization and funds from the Corporation. The author was a member of the consulting team that carried out this work.

THE DECISION ANALYSIS CYCLE

In a project, the decision analysis generally follows the cycle portrayed in Figure 1. The cycle involves three phases: the deterministic, probabilistic, and informational phases.

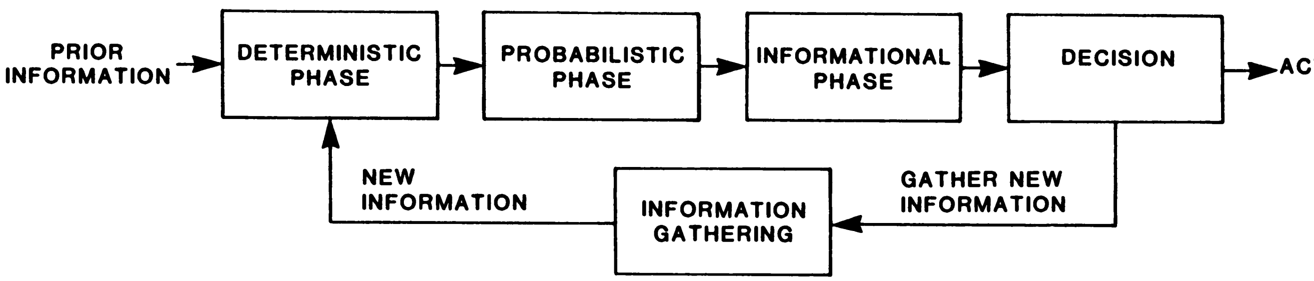

Figure 1: The Decision Analysis Cycle

In the deterministic phase, we isolate the decision that management faces and build a model incorporating the structure of the decision. With this model, we determine the important variables affecting the outcome of each alternative. A preliminary decision, based on a nominal case, is also determined. During the probabilistic phase, these sensitive variables are treated as probability distributions. In the deterministic phase, the model's output is a single number representing an estimate of the outcome, usually of profitability. In the probabilistic phase, the output is a probability distribution of this single outcome. In the informational phase, we calculate the economic worth of decreasing the uncertainty on the key variables. The decision can then be made to proceed with the best alternative or to gather further information. If additional information is obtained, the whole cycle is repeated.

THE DETERMINISTIC PHASE

The first step of the deterministic phase is called bounding the decision, which involves determining the decision to be made and isolating it from other issues. The decision hierarchy is a tool for ordering the relationship between different decisions and for clarifying which decisions are to be analyzed. Generally, decisions at any level of the hierarchy affect those decisions beneath them and are, in turn, affected by the more strategic decisions above them. As a result of extensive interviews at GOCHEM, we created the decision hierarchy shown in Figure 2, which delineated the decision levels involved.

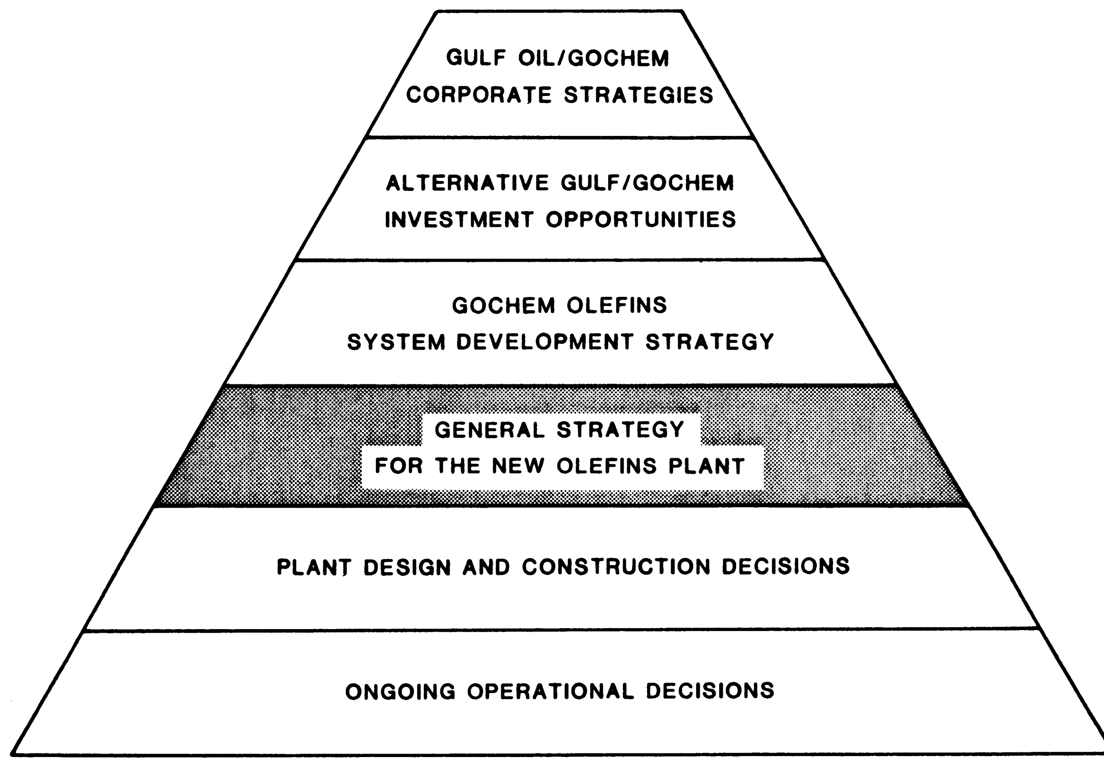

Figure 2: Hierachy of New Olefins Plant Decisions

Although GOCHEM personnel were concerned with many issues, such as location, the use of unionized labor, and plant design, most of those interviewed indicated that the main issues were whether to build a new plant, how big it should be, and whether it should be naphtha-based. Therefore, we focused the analysis on the go/no go decision, on the choice of feedstock, and on the size of the plant.

The next step in the deterministic phase is identifying the alternatives. In the interviews, we found there were alternatives on a number of the main issues. Concerning the feedstock, we discovered that even heavier feedstocks than naphtha could be used, specifically gas oil, the prime constituent of home heating oil. Unfortunately, gas oil has a lower ethylene percentage yield and even more by-products than naphtha. To obtain 2 billion pounds of capacity per year of ethylene out of a gas oil plant requires building a bigger plant, using more feedstock, and producing more by-products. Naturally, the bigger gas oil plant costs more than the proposed naphtha plant, which in turn costs more than a conventional propane and ethane plant. Moreover, though this gas oil feedstock process existed, it was not widely used, and Gulf did not have extensive experience in its operation. In further discussions with the engineers, we found that although feedstock flexibility could be built into plants, it could only be done at a substantial cost.

The Deterministic Model

The deterministic phase uses a structural model that shows the relationship between the decision and state variables and the resultant outcomes. A state variable is an uncertain quantity, not under the control of the decision-maker, that affects the outcome (for example, a by-product market price). A decision variable is a choice under the direct control of the decision-maker.

A fundamental characteristic of decision analysis is that models are developed for specific decisions: in this case, the go/no go, feedstock, and size decisions. The model we built was not intended to be useful for other level decisions, such as operations and engineering design trade-offs, nor was it supposed to be a detailed simulation of reality. Moreover, since it captured GOCHEM's perceptions of the industry at that time, it may no longer be applicable.

The right side of Figure 3 contains a preference model, which captures GOCHEM's attitude toward the time value of money.

GOCHEM management indicated its willingness to have its time preference represented by a simple constant discount rate. It also agreed to use cash flow as the principal measure, although the model also calculated ROI, payback, and profits.

Figure 3 also contains the olefins system model, which consists of two financial submodels: one for the new olefins plant and another for the existing ethane/propane-based olefins system. These financial models output a cash flow and net income to the preference model. Inputs to the models are feedstock and by-product prices and an ethylene price, as well as a specification of the parameters of the new GOCHEM ethylene plant (which come in from the box labeled New Plant Decisions). We included the existing system to capture the greater risk of using the same feedstock in all GOCHEM's plants.

The pricing contract logic is quite involved. Basically, ethylene is sold on long-term contracts that may have escalator clauses based on inflation indexes. Before beginning the construction of a new plant, chemical companies usually obtain letters of intent committing most of the primary products. This feature of the industry is central to the model.

The competitive market model assumes that at least one new plant will be required annually to meet the demand for ethylene (a very conservative assumption at the time). The model simulates the process by which GOCHEM

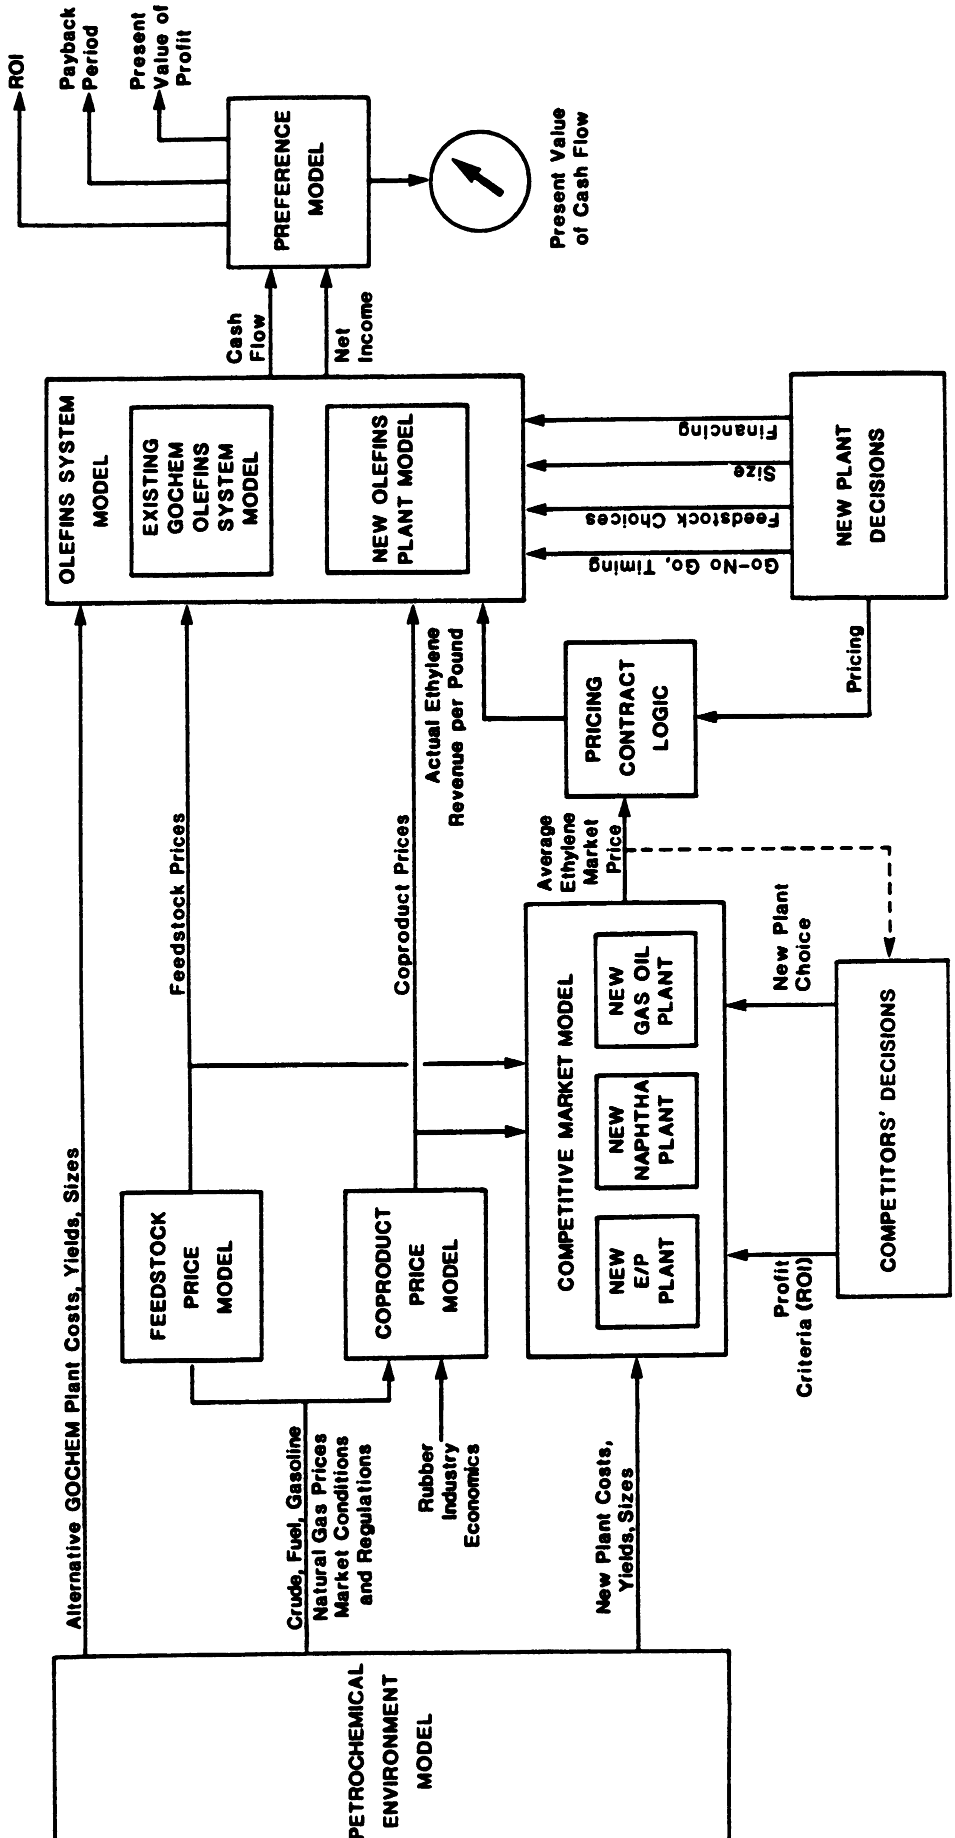

Figure 3: The Deterministic Model

and its competitors built such new plants. By using projected feedstock and by-product prices, they considered the economics of the alternative processes; then, using the most advantageous process and a minimum profit criterion, they computed an ethylene price. Finally, they went into the marketplace and tested the customers' willingness to pay that price. In general, the producer with the lowest profit criteria built the plant. In this way, an ethylene market price was established. The model represents this process for every year in the life of the new GOCHEM plant and uses the derived ethylene price to evaluate its financial outcomes. In effect, the market price of ethylene is determined by new plant economics -- using the most favorable process at the future time. In reality, deviations from these prices result from inaccurate forecasting of future demand and of feedstock and by-product prices. However, although these deviations affected profitability, sensitivity analysis showed that they would not upset the decision conclusions.

The usefulness of such a deterministic model depends on a consistent data set of feedstock and by-product prices. An examination of the basis for the different feedstock and by-product prices showed that they were both uncertain and interdependent. Most prices were strongly influenced by refinery economics. Either the materials were used in the production of gasoline, or they could be converted into gasoline at some process cost. Figure 4 contains a value relationship diagram based on price differences. The uncertainty in price differences turned out to be considerably less than the uncertainty in the absolute prices, which ultimately simplified the choice between alternative feedstocks.

The last step in the deterministic phase is the sensitivity analysis. Each of the estimates for state variables is varied over the full range, one at a time and sometimes in pairs, to isolate the crucial variables, which are those that can switch the decision from one alternative to another.

MID-PROJECT FINDINGS

We had also progressed in time to our mid-project oral report to the executives of GOCHEM. Based on our deterministic model analysis, we arrived at, and presented, the following controversial conclusions:

- The project was not risk-free, despite early contract guarantees. Its profitability was primarily endangered by the future actions of competitors.

- Changes in by-product prices and inflation rate levels, major sources of worry to management, did not greatly affect the profitability of the new naphtha plant. A market mechanism compensated for them.

- The predicted profitability of the naphtha-based plant was quite optimistic.

- Of the feedstocks considered, gas oil seemed better than naphtha.

- A 1-billion-pound plant was too small.

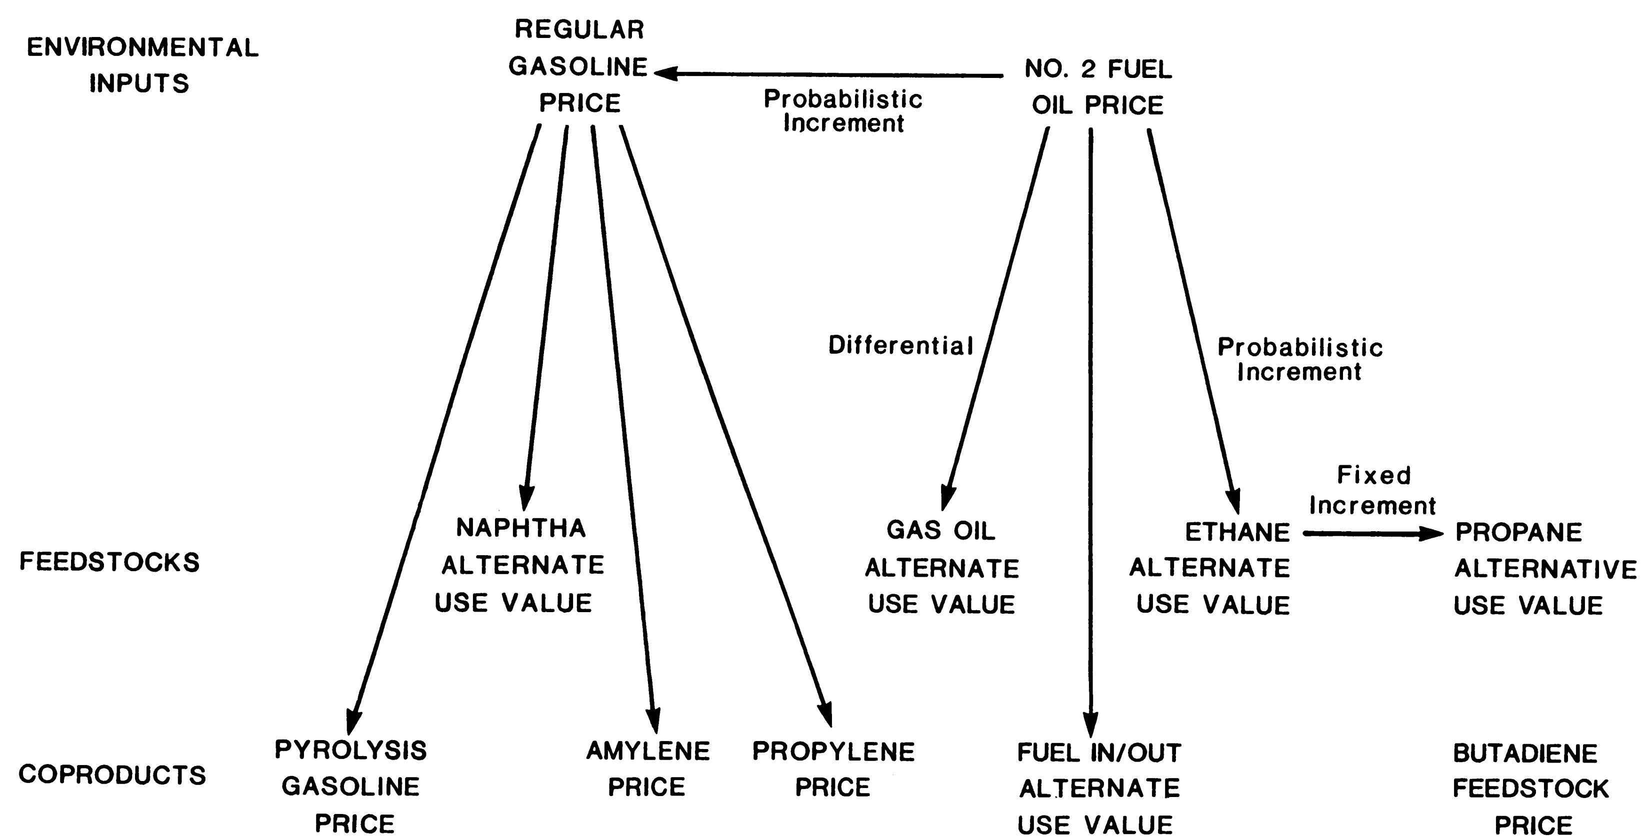

Figure 4: Feedstock and Coproduct Price Structure

The high point of the presentation was the results of the sensitivity analysis, which demonstrated the effect of possible changes in the state variables on the naphtha plant's profitability. It answered "what if" questions about:

- A precipitous fall of the price of prophylene (the major by-product);
- A year's delay in plant startup;
- The construction of a larger plant;
- An increase in construction costs due to inflation;
- The deregulation of natural gas by the federal government;
- Lower or higher than expected ethylene yields.

Although significant changes in some of the state variables produced substantial changes in profitability, only under unlikely conditions were the conclusions changed. A 20 percent larger plant was generally more profitable than the planned size, and a gas oil plant was almost always better than a naphtha plant.

While impressed by the analysis, GOCHEM executives were not convinced. After the first presentation, committees of GOCHEM personnel systematically examined the data base and ranges and checked the model and its assumptions with GOCHEM's financial models. During this period and the remainder of the analysis, the decision analysis team had regular contact with the executives and the rest of the GOCHEM staff; although some changes were made in both the model and the data base, the conclusions remained firm.

An in-house investigation verified that although the economies of scale were not as great as in the past, they were still significant, especially if viewed on a systems basis. The investigation also verified what impact competitors' potential decisions to build larger future plants would have on GOCHEM's profitability. As a result, the executives decided to increase plant size by 20 percent.

The matter of feedstock, however, was not so easily changed. Switching to a gas oil feedstock would require substantially more capital. In addition, GOCHEM had expended considerable time and effort in lining up an adequate supply of naphtha feedstock. GOCHEM executives suggested that a 50/50 flexible plant might be the best alternative. While beginning as a purely naphtha-based, the plant would have the necessary flexibility to change to up to 50 percent gas oil feedstock.

THE PROBABILISTIC AND INFORMATIONAL PHASES

The probabilistic phase uses the same model as the deterministic phase. However, rather than single deterministic estimates for all the state variables, probability distributions are used for those variables found crucial in the sensitivity analysis. This allows us to calculate a probability distribution on the present value of cash flow for each set of alternatives.

In decision analysis methodology, information and values are recognized as subjective, personal judgments. Those used in the analysis must be chosen by the decision-maker, who has ultimate responsibility for the project. Assessing an individual's uncertainty is both an art and a science, because the interviewer must be sensitive to the subtle biases of the subject. The subject's perspective of this interview experience is expressed in this quote from a speech by Mr. William Roher, the president of GOCHEM:

> "For each of the crucial variables, the interrogations begin. For some of those variables, I was one of the experts. I don't quite know how to describe my experience. At times, I thought I was talking to the CIA. . . . After an hour or so of this interrogation, we had, in addition to my Excedrin headache, a curve showing my very best judgment of the probabilities that various prices would prevail. Only time will tell how <u>good</u> my expert judgment was. All I know is that the interviewer brought forth the best I had -- and he did it without once using the word probability."

A disguised version of the cumulative probability distribution of one of the sensitive variables, the ratio of gas oil to naphtha plant investment, is shown in Figure 5.

Another major variable, the 1976 gasoline/fuel oil price differential, is shown in Figure 6. Although GOCHEM's president chose the individual whose probability distribution was to be used, he also suggested obtaining distributions from others who might have different opinions. Gathering probability distributions is a good way to communicate different opinions on a subject and to resolve arguments. As it turned out, the differing distributions did not affect the conclusion and it was not necessary to resolve the differences.

Figure 7 shows a schematic of the decision tree used in this analysis. The actual tree is much more complicated because each node is attached to every open branch of the previous node; as a result, there are numerous different paths through the tree, each one representing a possible scenario. Each of these circular nodes represents an approximation of a probability distribution.

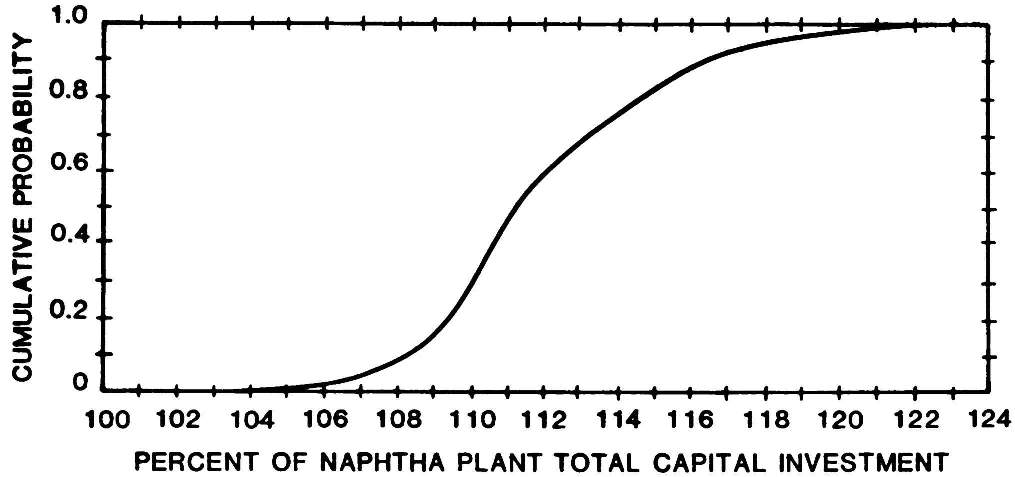

Figure 5: Uncertainty in Gas Oil Plant Total Capital Investment

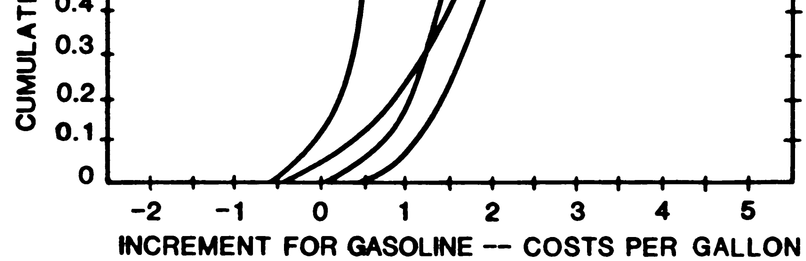

Figure 6: Uncertainty in Difference Between Regular Gasoline and No. 2 Fuel Oil Prices for 1976

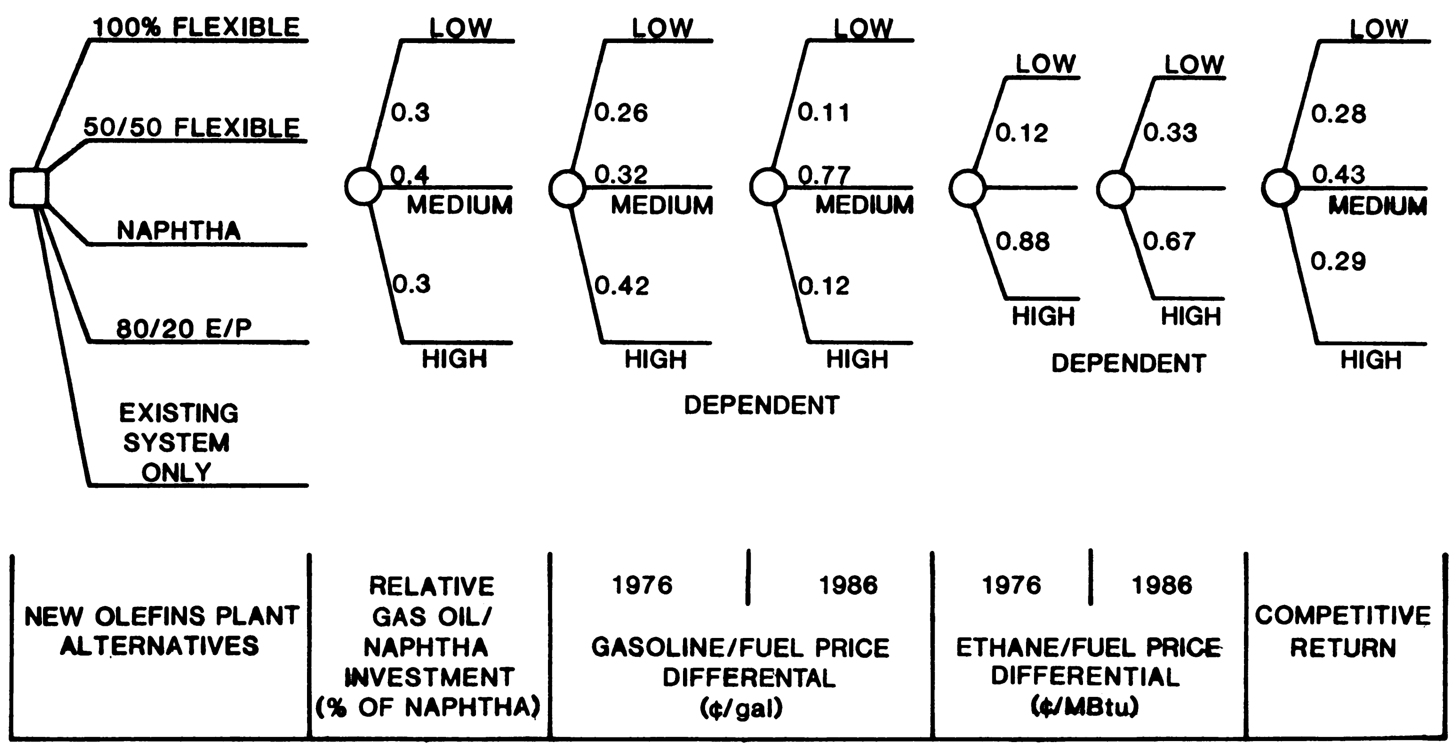

Figure 7: New Olefins Plant Decision Tree Structure

While we were accumulating the financial and technical data to convert a naphtha-based plant into a 50/50 capability plant, we also explored the cost of making a gas oil plant flexible enough to use naphtha. We found that the additional cost of such a plant over a pure gas oil plant was rather small. The investment, however, was still substantially more than what was needed for the naphtha and 50/50 capability plants. With GOCHEM's approval, we added this alternative to the list and called it the 100 percent flexible plant.

The results of the probabilistic phase are captured in the disguised profit lotteries of Figure 8, which are taken from the final presentation to GOCHEM's management. Despite its greater investment, the 100 percent flexible plant is far superior to the other alternatives. Not only does this plant have a higher expected value (represented by the solid bars on the graph), but, more importantly, it has a higher certain equivalent.

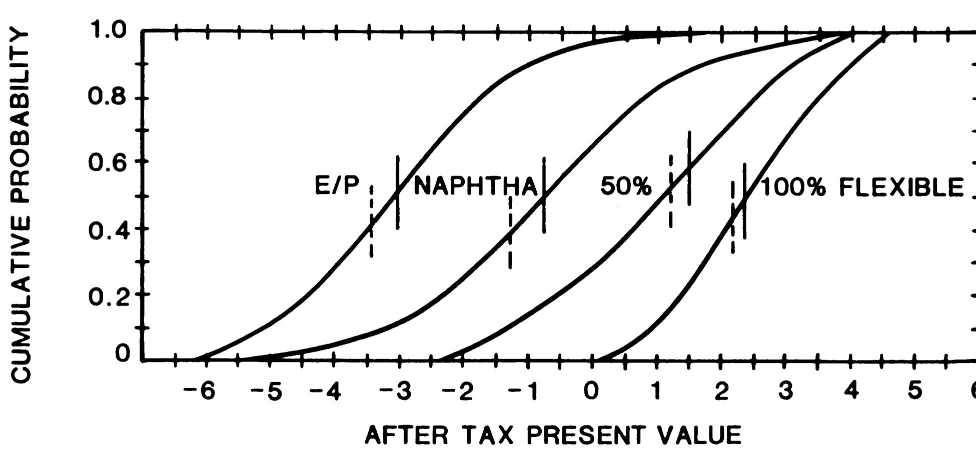

Figure 8: Uncertainty in Profitability of New Olefins Plant Alternatives

The certain equivalent is the amount of money the decision-maker would accept in lieu of the uncertain venture represented by the profit lotteries. It is a function of both the profit lottery and the decision-maker's attitudes towards taking risks. In this case, it was not necessary to measure the decision-maker's risk attitude, because the decision was insensitive to it -- the flexible plant had both a higher expected value and a smaller risk of bad outcomes.

Moreover, informational phase calculations of the expected value of perfect information showed that there was no obtainable information that could justify delaying the decision any further.

The impact of this decision analysis is best summarized by another quote from Mr. Roher's speech:

> "We in Gulf Oil Chemicals saw during our olefins analysis that a business we knew very well had its share of uncertainties and doubt. We also saw that in decision analysis we had a tool for making those uncertainties explicit. With that tool, we could communicate to all concerned -- in easily understood terms of profit -- the magnitude of both the risk and the opportunity associated with those uncertainties.
>
> "The final chapter to the olefins plant case is that decision analysis led us to change our minds. We had had preconceived notions. In the end, we decided:
>
> - To build a larger plant;
> - To choose a different plant site;
> - To build a plant based on different feedstock.
>
> "Even those in our midst who originally thought there was no need for decision analysis agreed we came up with much better decisions using this approach."

13

THE DANGEROUS QUEST FOR CERTAINTY IN MARKET FORECASTING

R. Abt and M. Borja
CIBA-GEIGY

M. M. Menke
Strategic Decisions Group

J. P. Pezier
Investment Intelligence Systems Corporation

Reprinted from Long Range Planning, Volume 12,
April, 1979, pp. 52-62.

The Dangerous Quest for Certainty in Market Forecasting

R. Abt, M. Borja, M. M. Menke and J. P. Pezier

Over the past 4 years, the Pharmaceutical Division of CIBA–GEIGY has developed a series of market forecasts in probabilistic terms, explicitly incorporating the principal uncertainties affecting the future sales of new pharmaceutical products. Such probabilistic forecasting methods have now been applied successfully by many pharmaceutical companies, including CIBA–GEIGY, Hoffmann-La Roche, and Sandoz. At CIBA–GEIGY this forecasting process has assisted the Divisional Management Committee with various business decisions, and in particular with capacity expansion planning. The benefits of this approach have been aptly described by the member of that committee responsible for production worldwide, Dr. Hans M. Götz, 'Since we have used the probability curve for active substance requirements, the most valuable result has not been that we get now any better single number, such as mean, mode or any other figure, but that we have the curve at all. No longer do production people and market forecasters accuse each other because of a "wrong" figure. Both sides are aware of the uncertainties and therefore their dialogue has become more constructive. Decisions are made and the results are reviewed in light of the curve that represents our best judgment when action is required. Thus our principal benefit has been the improvement in the quality and co-ordination of management effort in production planning and resource allocation.' This paper explains in detail the steps needed to prepare a probabilistic market forecast and demonstrates how such a forecast can yield significantly different and better production capacity decisions and improved insight into the related process of market planning.

If a man will begin with certainties he will end with doubts, but if he will be content to begin with doubts he shall end in certainties.

Francis Bacon (1561–1626)

Introduction

The quest for certainty! Which manager has not dreamed of having a crystal ball where he could see exactly what the future holds for him. Unfortunately, no forecasting tool, however sophisticated, comes close to this utopian goal. Even with the most modern market forecasting methods, there always remains a large degree of uncertainty about the future sales of new and even current products. In fact today's rapidly changing economic, competitive and regulatory environments seems to increase the range of uncertainty incessantly. Unfortunately, our intuition is not well prepared to cope with uncertainties. Our educational and cultural background encourages categorical assertions, suggesting confidence while suppressing true feelings regarding uncertainty.

Managers usually prefer to use single number estimates rather than to take into account the full range of possible outcomes, since the latter course would openly acknowledge their inability to predict accurately—suggesting a lack of commitment or even indecisiveness, undesirable attributes for a manager.

The danger is that in many problems single number estimates are inadequate to reach sound conclusions; in some areas (e.g. R & D) no such deterministic estimate may even be possible! This danger is particularly acute when the responsibilities for forecasting and decision making are separate, as they are in most organizations today. Single point estimates are fraught with the difficulties of balancing pessimism against optimism; furthermore, motivational biases often lead to stated predictions that differ from true beliefs. This latter phenomenon can be particularly acute in the case of the salesman who is rewarded for exceeding his target or the product champion who needs to defend the viability of his programme.

A common tendency today to cope with the increasing uncertainty is to search for better forecasting tools promising ever more precise predictions. Nonetheless, the uncertainties persist. A positive alternative is to face the crucial uncertainties squarely. The language of probability offers a simple means for describing and conveying uncertain information explicitly. In the past few years, probabilistic methods have proven their worth in fields as diverse as banking[1] and pharmaceuticals. In the pharmaceutical industry, probabilistic methods have been applied by Sandoz to R & D planning, ,[3] by Hoffmann-La Roche to investment analysis and production decisions,[4] and by CIBA-GEIGY to market forecasting and business strategy.[5] This paper will demonstrate how future sales volumes can be described in probabilistic terms which convey more completely the state of knowledge in the organization; an actual example from CIBA-GEIGYs experience will show how a full probabilistic evaluation of market demand avoided oversizing a new pharmaceutical plant by about 20 per cent, resulting in a substantially reduced capital expenditure requirement. Thus probabilistic forecasts are not a costly luxury, but rather the only sensible and financially conservative way to evaluate demand/capacity problems in a highly uncertain environment.

Forecasting for Decision Making

A Market Requirement/Capacity Planning Problem
CIBA-GEIGYs Pharma Division was considering expanding the capacity of a multi-purpose plant to produce a range of active substances whose production processes were similar. To put the problem simply, the same facility could produce the various substances with minimum down time and set-up costs required to shift production from one of these substances to another. These active substances were the basis for a number of products currently in various stages of clinical development or early marketing. Marketing research was asked to indicate how much active substances would be required in the years to come for each of these products. The purpose of these forecasts—to decide on the size of the plant expansion—was at first left unstated by management.

☆ In many companies the subsequent forecasting activity might have gone as follows. Management had requested a 'best' estimate of what the requirements would be. Marketing researchers, however, would make a 'safe' estimate, i.e. one that has a good chance of being reached or exceeded. Management, well aware of this conservative tendency through long experience, would then apply an appropriate corrective factor. Management would also make additional adjustments to account for its preference to have some idle capacity rather than to lose sales due to production constraints.

☆ On the other hand, it is also well known to management that engineers tend to overdesign due to uncertainties about the final chemical process efficiencies, plant reliability, economy of scale and so on, and that they do not want to promise more than they feel sure to deliver. In the end, management faces an exceedingly intricate problem requiring all the business acumen it can muster. Moreover, marketing research and production people were well aware of these corrective tendencies of management in the first place and may have taken them into account in their original estimates. Once entered, the vicious circle of making compensations to account for other peoples corrections can lead very quickly to total confusion.

Fortunately, marketing researchers and Pharma management of CIBA-GEIGY had already been exposed to probabilistic analysis. Given the large uncertainty surrounding future sales figures, they succeeded in demonstrating to management the need for expressing future active substance requirements in probabilistic terms. This in turn enhanced management's basic action goal, namely to determine the production capacity that they should install now to meet future requirements.

What a Decision Analysis Approach has to Offer
Information is both subjective and objective; it is a state of knowledge and a basis for action. Marketing researchers are confronted with these two aspects of information. First, the information they collect generally comes from different sources with various degrees of objectivity and subjectivity and they must collate it into their own, necessarily subjective view of the future. Second, forecasts are made for a purpose and by defining clearly this purpose one can improve and simplify the forecasting problem.

A decision analysis approach[6,7] has much to offer the marketing researcher in both respects. On the one hand, decision analysis has many methods and techniques for acquiring and combining marketing information. Two helpful techniques here are probability encoding and structural modeling. On the other hand, decision analysis techniques are very helpful in simplifying and guiding the forecasting effort. Two relevant techniques here are sensitivity analysis and value of information. These four techniques are briefly described below.

First, through structural modeling a complex forecasting problem can be decomposed into simpler constituents and formulated to take maximum advantage of the available expertise. Structural information, i.e. information about the relationship among various problem constituents, can be separated from numerical information about the value of specific variables. These relationships can then be examined and criticized by many experts and managers. Experts can concentrate on the relationships within their fields of expertise; management can concentrate on the integration of functional areas.

Second, probability encoding techniques have been developed by decision analysts[8] and psychologists[9] to extract experts' opinion about the value of an uncertain quantity in the form of a probability distribution. It may seem surprising that this has been one of the greatest challenges to decision analysis application. Experiments, however, have repeatedly shown that the assessment of an uncertain quantity in probabilistic terms is usually distorted. One such widely documented distortion is the so-called central bias which leads people to believe that they know more than they really do. As a consequence, bigger risks are assumed than were foreseen and people are often surprised when the unexpected happens.

Third, sensitivity analysis can be used to focus management attention and staff effort on the most crucial factors, i.e. those factors whose uncertainty has a dominant effect on the forecasts. Thus the forecasting and modeling efforts can be well balanced, refinements being applied first in those places where they can be shown to be the most useful.

Finally, if the forecasting task is in the context of a well defined decision, the overall level of resources to be committed to this task can be evaluated. Only information that can improve the decision making process is valuable. It is worth acquiring if its cost of acquisition does not exceed the benefits expected from improved decision making. This capability of decision analysis to evaluate and direct an information gathering programme is quite unique. Furthermore, it offers the marketing researcher a new and powerful tool to measure the economic value of his efforts.

To avoid paralysis by analysis the methods described above are implemented in an iterative procedure starting with simple, 'back-of-the-envelope', calculations and introducing each new refinement when the need arises. This cycle, described in Figure 1, starts with structural modeling and single number estimates. The outcome to be forecast is clearly defined in relation to the various decisions and environmental factors that affect it. Assumptions are made about which decisions are considered as given and which decisions should be internalized, i.e. should follow from a prescribed strategy. Also the purpose of the forecasting exercise should be stated in decision terms. Sensitivity analyses are conducted to identify the most crucial variables on which the efforts of the probabilistic analysis should be concentrated.

The probabilistic phase extends the previous deterministic analysis, using probability encoding to represent explicitly the uncertainty in the most crucial variables and then computing the resulting degree of uncertainty in the forecast. A formal model, preferably computerized, proves very useful here due to our poor aptitude to combine uncertainties intuitively and our reluctance to engage in the tedious and lengthy task of doing it by hand.

The informational phase is only possible in the context of a decision problem. One first computes the expected value of reducing or eliminating the uncertainty in the most crucial factors. Then possible information gathering programmes and modeling refinements are examined to see if they could effectively reduce uncertainty at a cost—including the cost of delay—less than the expected value to be derived. If such programmes can be found they should be undertaken and the forecasting problem should be re-examined in the light of the newly available information.

Structuring the Pharmaceutical Requirements Forecast

Figure 2 depicts the problem of forecasting pharmaceutical active substance requirements within a comprehensive decision model framework. The market model at the upper left is influenced by external factors belonging to the market and competitive environment, as well as internal factors emanating from the research model and the production model. The basic outputs of the market model—sales revenue, promotional expenses and active substances requirements—become inputs to other models in the overall decision model structure.

Based on a given direction and level of research effort, products currently under development can be expected to be registered and introduced on the market at some later year. The characteristics of these products can be described by estimates of effectiveness, tolerability, dosage requirements, etc. Research and development will also determine chemical process efficiencies and other production cost factors. Subsequent research and development efforts can lead to improvements in the production process and to new dosage forms.

The production model develops several categories of cash outflows: fixed cost, variable cost and investments. On the basis of plant configuration and capacity decisions, and knowledge of the production process, it also indicates what quantities of active substance can be produced in the years to come.

This comprehensive view of the decision model framework points out two complications of the marketing researcher's task. First, the forecast of active substances requirements is just one of several kinds of information to be obtained through a careful market analysis. Sales revenues and promotion expenses are also derived from the same market analysis. Therefore, the level of detail of the market analysis should be dictated by these multiple needs, and not only by capacity planning considerations.

Second, there is a hierarchy among the research and development, production and marketing decisions shown in Figure 2. R & D decisions are very long term; the market forecast for products in the late stages of clinical trials or early stages of market introduction depend on the current outcomes of R & D decisions taken years earlier. At the other extreme, most marketing

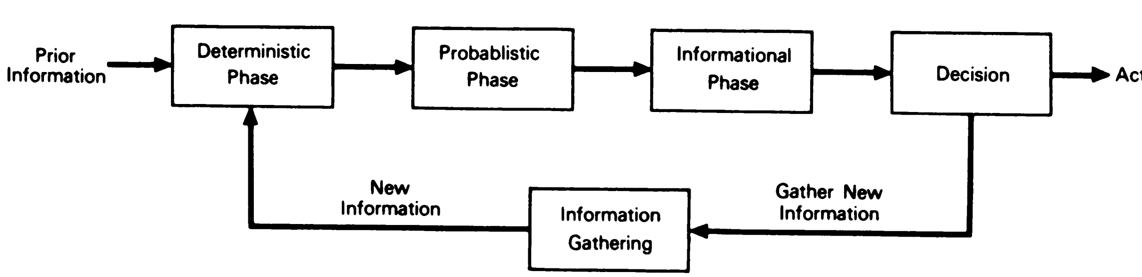

Figure 1. The decision analysis cycle

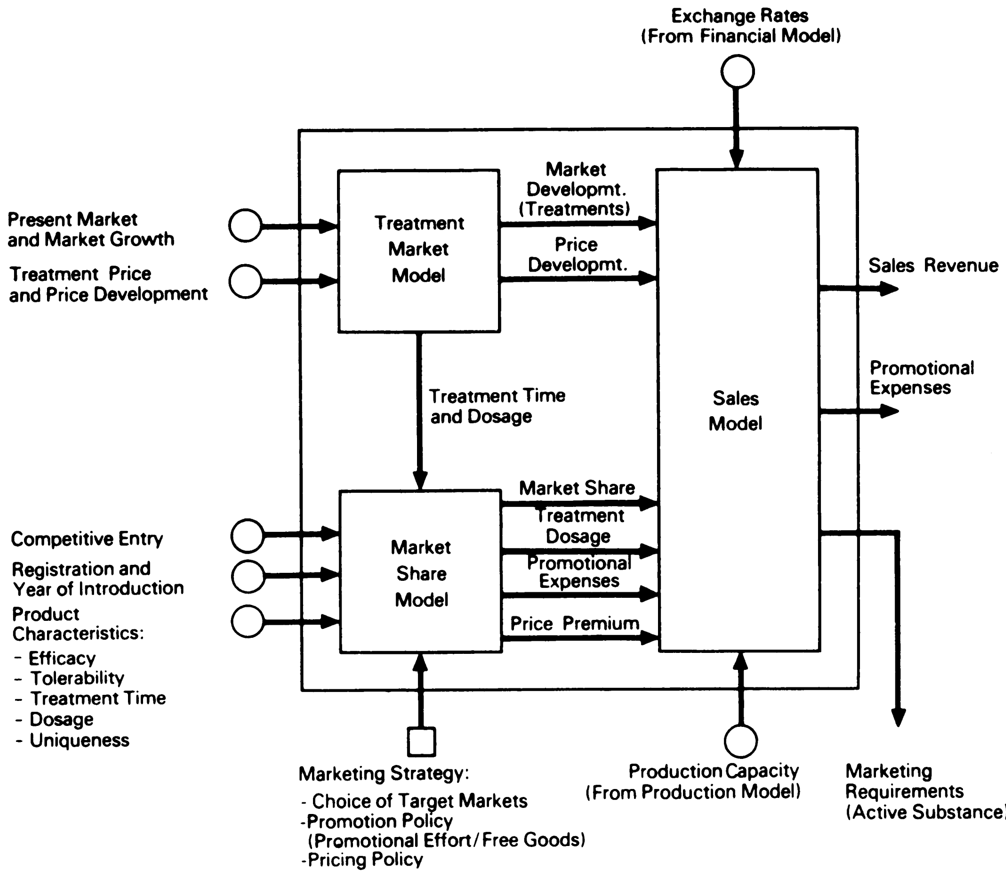

Figure 3. Structure of the market model

it may be sufficient to limit the calculations to peak sales levels, whereas for products that will still be in the early stages of marketing attention must be given to market penetration factors.

Probabilistic Forecasts

Pinning Down the Crucial Unknowns

In this and the two following sections the discussion will be illustrated by the case of a product used for the treatment of acute ailments, sold in a homogeneous market and having almost reached peak sales by 1982.

The first calculations with the forecasting model are carried out with single number estimates. Specialists in research, development, medicine, production and marketing are asked to provide realistic estimates for the variables in their respective fields of expertise. Management stipulates policies and guidelines that it wishes to apply although these decisions may be revised at a later stage in view of the forecasts. The combination of these assessments and hypotheses form a reference case.

For the product under scrutiny, the reference case led to an average requirement of 72 tons per year of active substances from 1982 to 1987. It is not easy to say whether this figure is the most likely value, or the expected value, or even has some minimum probability of being achieved or exceeded. Even if each input variable were consistently assessed according to one of these criteria, the output forecast would generally not correspond to the same criteria. The purpose of the reference case is to test the forecasting model and to provide a well documented reference figure for further analyses.

The next step is to conduct sensitivity analyses to test the importance of uncertainty and identify the most crucial factors. In the sensitivity analysis input variables are varied over a 'reasonable' range of possible future values and the corresponding variations in requirement forecast are computed.* An effective sensitivity analysis requires therefore that the forecasting model be sufficiently reliable to capture the logical relationships among the factors influencing sales. The model should also be

*In the simplest case, each variable is swept over its range of likely values, the other variables being held at their base case representative values. However, if several variables are strongly related they should be varied simultaneously. This could be a complex and lengthy task. Often a useful alternative is to capture the dependencies in a more explicit model structure containing more but independent variables.

easy to implement so that the effect of each variation can be calculated quickly.

The most crucial variables are those whose initial range of uncertainty contributes the most to the uncertainty of the forecast. How crucial a variable is thus depends both on the sensitivity of the forecast to changes of this variable and on the range of uncertainty that has been attributed to it. The last point can be criticized as a logical circularity. Indeed, a variable whose uncertainty has been underestimated initially might not be included in the list of crucial factors as it should.

Two lessons can be learned from this observation. First, it is a reminder that a forecasting model cannot be expected to give a 'correct' picture of the future. All that can be required and achieved is that it should be a truthful representation of the marketing researcher's judgment. Second, it emphasizes the importance of assessing comparable ranges for all input variables, e.g. 90 per cent confidence intervals. Which variables turn out to be the most crucial and therefore merit further attention depends wholly on this kind of subjective judgment.

In practice, the situation is rarely as confusing as it might appear because in most forecasting problems a few crucial variables clearly account for most of the uncertainty. One can therefore conservatively consider as crucial variables that, *a priori*, would contribute as little as 1 per cent of the uncertainty in the overall forecast and still retain only a manageable (usually five to ten) number of variables.

Figure 4 lists the main variables in the case at hand. Those variables capable of causing the largest variations of the requirements (see the rightmost column of Figure 4) have been listed first. As a simple rule of thumb, the uncertainty in the forecast of requirements can be considered as proportional to the sum of the squares of the variations indicated in the last column of Figure 4.* This rules shows, for example, that the first variable on the list, product quality contributes about 400 times more to the overall uncertainty than the eighth and last variable on the list, time from entry to maturity.

For all practical purposes it is usually sufficient to concentrate on the first five or so variables on the list in order to estimate the uncertainty of the requirement forecast. If one applies the conservative rule of the 1 per cent cut off, the sixth variable, competitive entry, should also be included. We shall see next how uncertainty about these six variables was quantified and a probabilistic requirement forecast was derived.

Encoding the Uncertainty of Crucial Variables

The most comprehensive description of an uncertain quantity is in the form of a probability distribution over all possible values. For a yes–no event, e.g. FDA approval, the probability distribution is simply the probability that the event occurs and the complementary probability that it does not. For a continuous quantity, e.g. North America market share, a probability distribution indicates the probability that the uncertain quantity will take a value within any given interval.

A continuous probability distribution can be conveniently depicted by a cumulative probability curve. The curve represents the probability (vertical axis) that an uncertain quantity will not exceed a given value (horizontal axis). Figure 5 shows a cumulative probability curve and indicates how some of the commonly used measures—expected (or average) value, median, and

*It can be shown that under general conditions (independence and linearity) this measure of uncertainty is consistent with the statistical concept of variance.

State Variable	Range of State Variable		Range of Requirements (Tons)		Variation of Requirements (Tons)
Product Quality	Me-Too	Hit	35	140	105
Treatment Market Growth	25/5/0	30/25/15	49	115	66
Daily Dosage[(a)]	40%	70%	48	84	36
Introduction Year	-1	+2	82	55	27
Treatment Time[(b)]	75%	100%	64	85	21
Competitive Entry	YES	NO	60	72	12
Free Goods	10%	15%	72	78	6
Time From Entry to Maturity	4	8	74	69	5

(a) Per cent of a Nominal Daily Dosage

(b) The Average Time Required to Cure an Acute Disease as a Per cent of Time Required with Current Leading Drug

Figure 4. Sensitivity analysis of active substance requirements

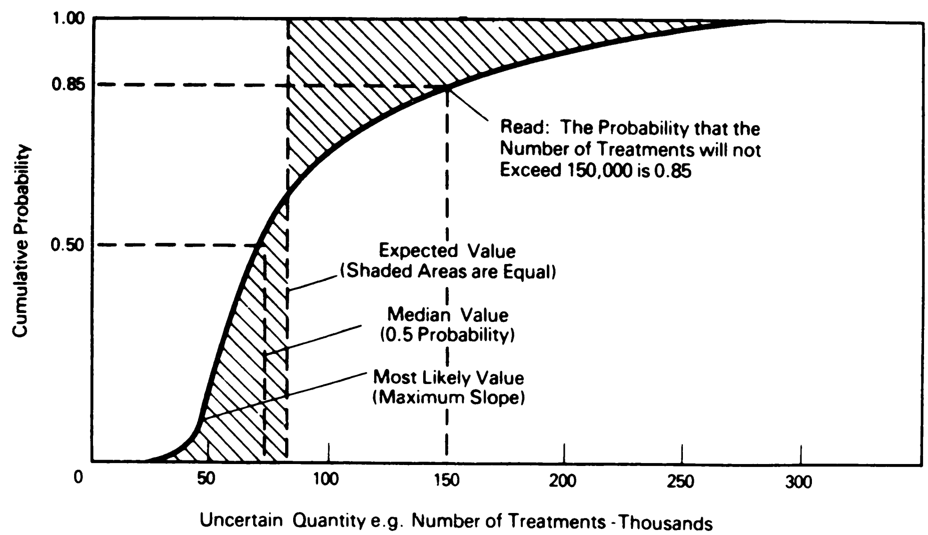

Figure 5. Description of an uncertain quantity by a cumulative probability distribution

most likely value—can be found on the curve. The flatter the curve, the more uncertain is the quantity being described.

How to transform personal knowledge into a probability distribution has been and still is one of the most challenging tasks for the decision analyst. From 10 years of experience encoding experts' judgment in the course of solving client decision problems, SRI is convinced that in many instances the most effective way to elicit these opinions is through an interview conducted by an experienced analyst. Although this is time consuming, and therefore costly, we feel that it can greatly reduce both motivational and cognitive biases—resulting in much more comprehensive and reliable inputs being provided to management.

Steps in a Probability Encoding Interview. A certain number of conditions should be met before attempting to encode an expert's knowledge about an uncertain quantity. The quantity considered must be crucial as defined above so that the expert is convinced of the usefulness of the exercise. However, the expert should not be concerned by the value of the uncertain quantity to the point where his judgment would be prejudiced. The quantity should be clearly outside of the control of the expert. The problem context must be carefully described and the quantity clearly defined (as a test, consider whether a clairvoyant, a hypothetical person with an infallible crystal ball, could tell the exact value). The experts should feel familiar with the units of measurement and the terms used to describe the quantity so that he may concentrate his attention on his beliefs rather than being entangled in mental gymnastics.

When these conditions are met, the analyst can start preparing the expert for the encoding session. The purpose of the preparation is to minimize the effect of some all too frequent cognitive biases. Among these are the differences in availability of various forms of relevant information (for example, too much weight may be given to the most recent information), the tendency to focus on one representative value (e.g. the value in the business plan) and to ignore other possibilities and the tendency to assume that we know more than we really know, the so-called central bias. Availability problems may be alleviated by listing hypotheses, information sources and pertinent factors taken into account. The focus can be shifted away from a representative value by first discussing scenarios for extreme cases. Extreme scenarios can in turn minimize the effect of central bias. Central bias can be further reduced by a discussion of whether the quantity can be predicted easily or with difficulty and whether it is expected to have an ordinary or an exceptional value.

Only then should the actual encoding session take place. Experience shows that this session should take the form of an interview between the analyst and the subject. After all, quantifying personal opinions without outside help is about as difficult as psychoanalysing oneself. General rules that interviewers have found useful are to:

(1) Avoid influencing the subject, either by providing him information or asking leading questions.

(2) Be sensitive to shifts in the subject's judgment (new information or new rationale) which may invalidate previous answers.

(3) Ask simple questions, e.g. ask for relative rather than absolute judgments, analogue rather than numerical. Use simple comparisons with easily understood reference events.

(4) Discourage subject concerns about consistency; vary the questions so that there is no simple logical progression from one to the next; record the answers outside of the subject's view.

(5) Verify the answers at the end of the interview by showing implications. Test the credibility of answers by constructing several fair bets from the

distribution and verifying whether the subject is indifferent about taking either side of the bets.

This approach has been discussed and presented in greater detail elsewhere.[6] It has been outlined here due to the critical importance of carrying out the encoding task carefully in order to develop a valid forecast.

Combining the Judgment of Multiple Experts. Resolving the opinion of a group of experts into a consensus opinion which then represents the organization's judgment is a difficult task since there is no perfect procedure. Having experts express their views, however, in the form of a probability distribution rather than a single number estimate is already a substantial step toward reaching a consensus. It is much easier to have experts agree on a range of uncertainty than on a single number. Various methods similar to Delphi techniques[10] can be used to bring experts toward a common view and should be preferred to blind averaging methods.

Ultimately, however, since experts rarely agree, additional bridging mechanisms may be required. If a decision-maker is not present or does not want to impose his view, some kind of weighting of the experts may be feasible. The rationale is that even when forecasting or decision making is done by groups, a consensus view about each of the elements of the problem should be reached. This 'group state' of knowledge, which in the end is to be adopted by the decision maker, may well be achieved by calibrating disagreeing experts, i.e. by granting to each expert a certain probability of being right, and then computing a weighted average of the experts' opinions. Subjective judgment and past performances can be used to assess the weights.

Advantages of Information Expressed by Probability Distributions. There are many reasons why a probability distribution is an ideal way to express one's judgment. It is the most complete expression of the expert's opinion, displaying what he does not know as well as what he does. Many experts are more honest when providing probability distributions, since it allows them to express their full range of opinion. The risk conveyed by this graphical representation of uncertainty may act as a stimulus for management to prepare contingency plans and probabilistic analysis can allow hedging strategies to be evaluated on a sound economic basis. Finally, when the uncertainty in our current state of knowledge has been clearly expressed, it is possible to attribute an economic value to reducing that uncertainty through marketing research.

Probabilistic Analysis with a Probability Tree

Probability trees provide a very clear representation of a series of uncertain quantities affecting a forecast. To facilitate graphical representation and communication, continuous distributions must first be summarized by a few discrete values with corresponding probabilities. This simple technical step has been illustrated on Figure 6. The continuous variable, daily dosage as a per cent of dosage for the current market leader, is represented by a

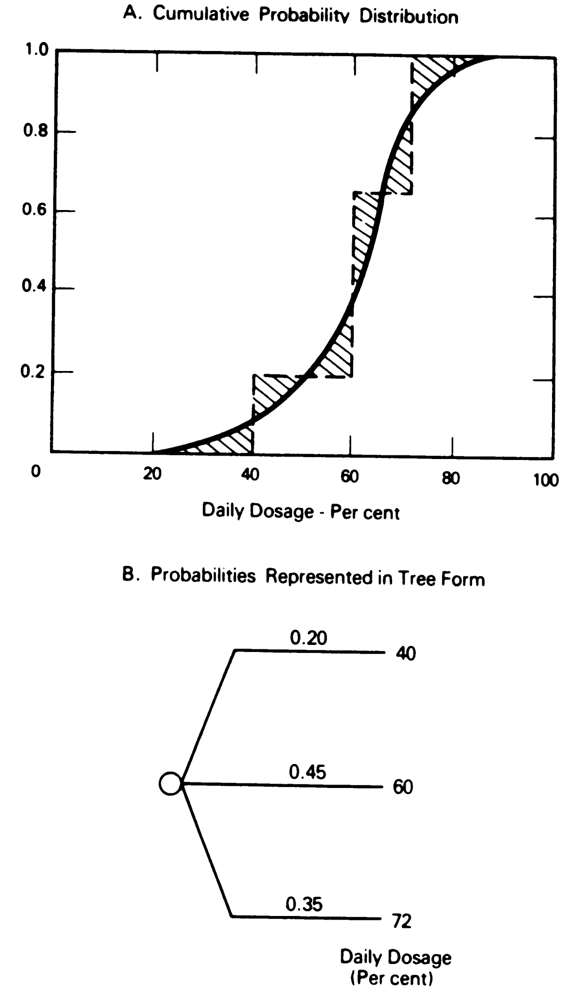

Figure 6. Probability distribution for a crucial state variable: the daily dosage

continuous probability distribution in Figure 6A. Figure 6B represents the same variable in a discrete tree-like form. The three discrete values 40, 60 and 72 per cent with their corresponding probabilities were obtained by drawing a three step function that 'closely' resembles the continuous cumulative probability distribution of Figure 6A.

There is no universal definition of how close the discrete and the continuous representations should be. In practice, if the area of the curvi-linear triangles above and below the continuous curve and the step function (cross-hatched area in Figure 6A) are equal two by two the expected values of the two distributions will be exactly the same. If, in addition, there are at least three discrete values and all the triangles are approximately of the same size the discrete representation will exhibit a degree of uncertainty close to that of the continuous distribution.*

*For example, the standard deviation of a three step discrete distribution will generally be within 5 per cent of the standard deviation of the continuous representation.

Figure 7 is a schematic representation of the probability tree containing the first six variables in the sensitivity analysis list (Figure 4). The most important variable of the list, product quality, has been decomposed in two stages: first a categorization of the product as a Me-too, a product with marginal advantage (Mad) or a Hit, and second a specification of market share uncertainty for each quality category assuming no competitive entry. The unlikely event of failing to register the product, which was not shown in the sensitivity analysis since its consequences are obvious, has been represented at the base of the tree.

Dependent variables have been linked and shown explicitly. Except for registration which clearly conditions the subsequent marketing of the product, the only dependent variables are product quality and treatment share. All the other variables (not explicitly linked) are assumed to be independent. A complete description of the tree would necessitate the repetition of the tree fork for each variable at the tips of the branches of the variable at its left. Thus the complete tree has 1 + 4 × 3 × 9 × 2 × 3 × 2 = 1297 terminal branches. To each terminal branch corresponds an active substance requirement and a probability. The probability tree is simply a means to span all the possible values the uncertain variables can assume and to calculate the probability of each combination of values.

The results of the probability tree analysis can easily be organized in the form of a cumulative probability distribution for the active substance requirements. The 1297 possible requirements are ranked from the smallest to the largest and for each requirement, the probabilities of smaller or equal requirements are summed up. When plotted, these probabilities give the curve of Figure 8 where, because of the large number of small increments the cumulative distribution appears to be a smooth curve.

There are, of course, alternative means of deriving complete probabilistic forecasts, or more limited information, from the same model and inputs. Two other techniques that have been in vogue are Monte Carlo simulations and scenarios. Both in our view suffer from deficiencies

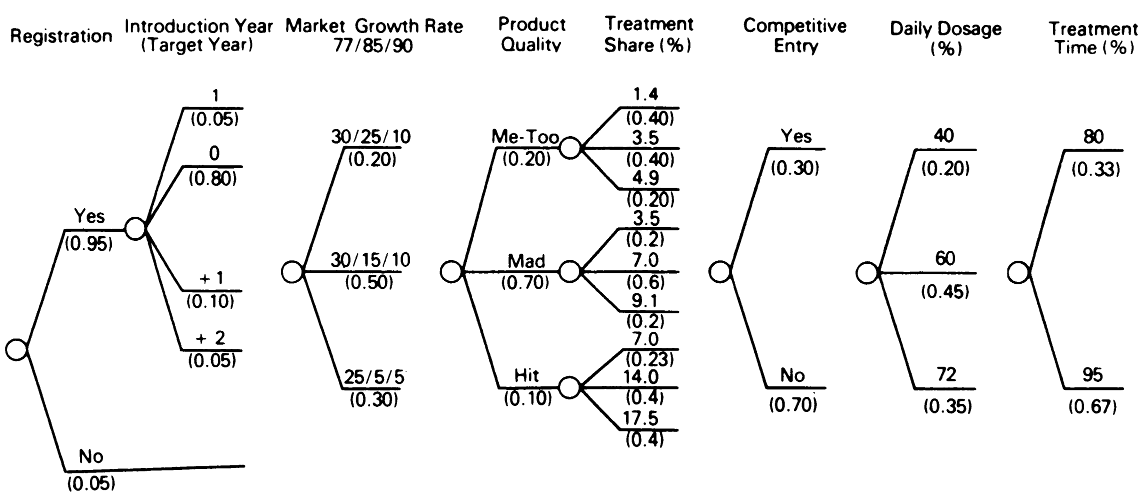

Figure 7. Probability tree for active substance requirements

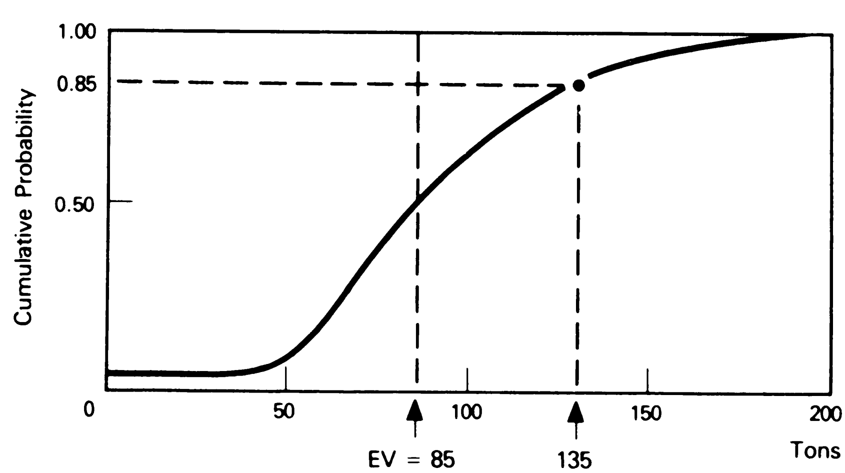

Figure 8. Active substance requirements for product A

and do not offer any real advantage over the probability tree approach. With scenarios, it is very difficult to define a small number that truly span the range of interest. With a probability tree, on the other hand, each path corresponds to a specific scenario and the tree structure organizes them in a well balanced way. A Monte Carlo simulation ,can cover all scenarios, but lacks the communication potential of the tree and is sometimes more costly to calculate.

Capacity Planning Based on Probabilistic Forecasting Requirements

What is the advantage of expressing the active substance requirements in terms of probability distributions over demand? If demand were known exactly the capacity expansion decision would be trivial. However, demand is usually uncertain and one must therefore balance the risk of oversizing the plant against the possible opportunity loss of not satisfying demand. The proper balance of risk and opportunity cannot be achieved without having a complete probability distribution over the possible levels of demand as represented in Figure 8.

Consider first the case of a single product plant. For example, what plant capacity should be recommended if the active substance requirement can be anywhere between 50 and 200 tons per year? A 60 ton capacity plant is almost certain to create supply shortages and therefore a great deal of lost sales. On the other hand a 150 ton plant will be much more costly and is very likely to be underutilized. If we consider first a very small capacity plant, sales will almost certainly be limited by production constraints so that a unit increase in capacity can be expected to yield nearly a unit increase in sales volume. Such an increase would presumably be profitable; otherwise there would be no point in commercializing the product at all. As larger capacities are considered sales are less likely to be curtailed by production constraints and the expected increase in sales becomes a smaller fraction of any new increase in capacity.

At some point the expected increase in sales will have become a sufficiently small fraction of the increase in capacity that any further increase will not be profitable; the desired capacity will have been reached. This break even ratio of incremental sales over incremental capacity can be derived from a standard financial evaluation of project cash flows. Call this ratio p; what then is the plant capacity such that incremental sales would be a fraction p of incremental capacity? The answer is quite simple; it is the capacity for which the probability of supply shortage is p and therefore the probability of satisfying demand is $1 - p$. Such capacity can be read directly off the probability distribution curve for demand. For example, returning to Figure 8, if the break even ratio p is 0·15, then the corresponding probability of satisfying demand should be $1 - p = 0{\cdot}85$ and the optimal plant capacity should therefore be 135 tons per year.

Thus even for the single product plant, a complete probability distribution for demand is essential to make an intelligent production capacity decision. However, the business risk has not been eliminated or even reduced; the method described above simply allows one to balance risk and opportunity effectively. The methods of References 4 and 6 also allow a useful graphical portrayal of the risk/opportunity balance. No single number estimate—however obtained—can reliably and repeatedly find this balance.

For the planning of multi-product facilities, it is even more important to convey information on product demand in terms of probability distributions rather than single number estimates. For such facilities, even the capacity/demand balance as determined above for the individual products will usually be insufficient, due to the portfolio effects that apply to situations involving several products with independent markets and uncertain demands. The portfolio effects reduce the uncertainty in total demand for several independent products and the ideal capacity for a multi-purpose plant is not simply the sum of the ideal capacities for the individual products.

To illustrate this point, consider again the active substance requirement for product A shown in Figure 8. A reasonable capacity, considering the economics of

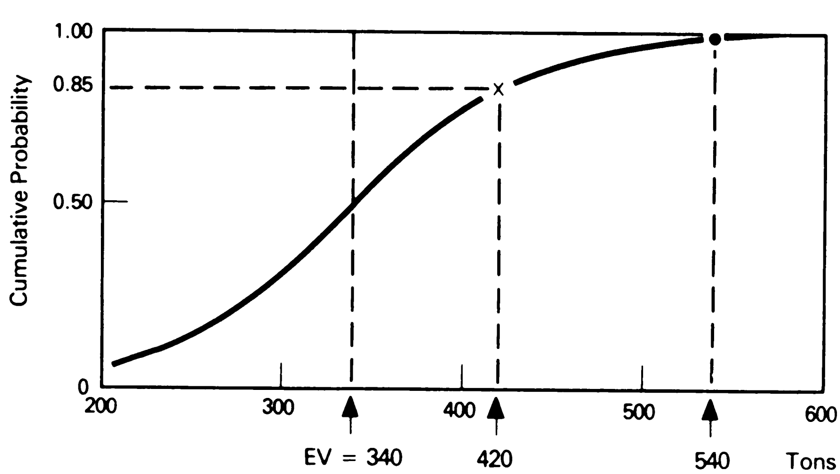

Figure 9. Active substance requirements for four products identical to A but independent of one another.

the product, would be 135 tons. Assume now for the sake of simplicity that the plant can produce four active substances for products with the same uncertainty on demand and the same cost factors as product A, but whose sales performances are independent. Summing up the individual requirements would imply a required capacity of 4 × 135 = 540 tons. However, if the total requirements for the 4 substances are computed using independent probabilities for the sales of the four products, one obtains the total requirements distribution shown in Figure 9. If all four products have similar costs and margins, then the desired total capacity again should correspond to a probability of 0·85 of satisfying total demand, namely 420 tons. This is of course much less than the 540 tons capacity which would result from examining only the individual capacity/demand balances, although the latter procedure was more representative of the way management have made such decisions in the past.

The reason for this discrepancy is that adding uncertain quantities is a more complex operation than adding known quantities. When uncertain quantities are independent the relative uncertainty of the sum is less important than the relative uncertainty of each constituent; this is the well known portfolio effect. Thus the sum of 85 per cent requirements for each of the four products is not an 85 per cent total requirement but has in fact about 99 per cent chance of not being exceeded. This line of reasoning directed Pharma management towards a more modest overall demand requirement, which then had to be compared with technically feasible and economically effective plant sizes to determine the final capacity recommendation. It is clear, however, that a reduction in the total active substance requirements of over 20 per cent should yield a very substantial capital expenditure saving, paying many times over the modest extra effort required to perform the probabilistic analysis.

Conclusions

In spite of the many obvious deficiencies inherent to deterministic forecasts of uncertain quantities, this practice is deep-seated and widespread among managers. There are many understandable reasons for this, including our educational and cultural background, the psychology of business organizations, the reluctance of successful managers to adopt new methods and the difficulty of controlling in an admittedly uncertain environment. An additional deterrent is the time (and cost) required to gain experience with the probabilistic approach.

Nonetheless, organizations which persevere are finding that the probabilistic approach to forecasting is not only quite feasible, but also offers many advantages that can not be gained in other ways, such as:

☆ the ability to communicate is improved
  —among marketing experts who find it easier to reconcile their opinions after major uncertainties have been admitted and explicitly stated;
  —between marketing and development staff who often cannot express more than a probability of technical success;
  —between marketing and production where capacity decisions may hinge upon requirement uncertainties for several products.

☆ the process of structuring the relationships among the various factors impinging upon sales gives better insights for marketing planning.

☆ sensitivity analysis guides marketing research toward the most important factors and value of information calculations can justify the overall level of effort.

In the end, the marketing forecast, the planning process and the resulting recommendations can be more effectively explained to management. Once these benefits have been realized and fully appreciated, there can be no turning back.

References

(1) I. Kabus, You can bank on uncertainty, *Harvard Business Review*, **54** (3), 95–105, May–June (1976).

(2) H. U. Balthasar and St. Gutzwiller, Steady state and portfolio concept in R & D management, *R & D Management*, **5** (3), 201–207, June (1975).

(3) H. U. Balthasar, R. A. A. Boschi and M. M. Menke, Calling shots in R & D, *Harvard Business Review*, **56** (3), 151–160, May–June (1978).

(4) J. W. Dangel, Investment analysis at Hoffmann-La Roche, prepared for OECD/IBRD Seminar on Industrial Planning, March (1976), available from Hoffmann-La Roche or SRI-Europe.

(5) For additional information on how decision analysis has been applied at CIBA–GEIGY, contact Dr. Rene Abt, Vice-Director of Marketing Research and Development.

(6) R. A. Howard, Decision analysis: applied decision theory, in: D. B. Hertz and J. Melese (eds.), *Proceedings of the Fourth International Conference on Operational Research*, pp. 55–71. Wiley, New York (1968).

(7) H. Raiffa, *Decision Analysis*. Addison-Wesley, Reading, Massachusetts (1968).

(8) C. S. Spetzler and C.-A. Stael von Holstein, Probability encoding in decision analysis, *Management Science*, **22** (3), 340–358, November (1975).

(9) A. Tversky and D. Kahneman, Judgment under uncertainty: heuristics and biases, *Science*, **185**, 1124–1131 (1974).

(10) B. Brown and O. Helmer, *Improving the Reliability of Estimates Obtained from a Consensus of Experts*, Rand, California (1962).

14

AN INSIDE VIEW: ANALYZING INVESTMENT STRATEGIES

Robert F. Egger
F. Hoffmann-La Roche

Michael M. Menke
Strategic Decisions Group

Reprinted from Planning Review, May, 1981, pp. 32-39,
published by Robert J. Allio & Associates, Inc. for
the North American Society for Corporate Planning, Inc.

AN INSIDE VIEW: ANALYZING INVESTMENT STRATEGIES

By Robert F. Egger and Michael M. Menke

OVER THE PAST TWENTY YEARS, the management of F. Hoffman-La Roche, a Swiss multinational pharmaceutical and chemical producer, has developed a comprehensive process for the analysis of important investment decisions.* This process has evolved gradually and naturally, according to management needs and readiness. Today it embodies a full range of decision and risk analysis procedures, which allow rapid but thorough quantitative evaluations of important investment, business development and strategy decisions. The results of this analytical process are routinely presented to top management. In addition, the quantitative evaluation is actively sought out by line managers and functional experts throughout the company.

While careful systematic analysis may appear costly, it is far less costly than pursuing an inferior business strategy. In fact, Roche has found the cost of providing comprehensive decision analysis support to management to be modest in relation to the benefits. Specifically, decision analysis at Roche has resulted in greater management productivity and effectiveness, and enhanced transparency of the management decision process. To confirm these conclusions about the benefits of decision analysis at Roche, we obtained an hour-long interview with Dr. A. Hartmann, Vice Chairman of the Board of the Roche group. His comments and insights substantiate our findings and are integrated throughout the article.

Although the decision analysis methods Roche uses are not new, their success at Roche is in striking contrast to some recent reports of other management researchers. In an article entitled "Why Risk Analysis Isn't Working"[1], William Hall reported that a survey of 12 large U.S. companies points to five perceived problems associated with the use of modern, quantitative methods such as decision and risk analysis. First, says the article, it is impossible to demonstrate in practice that quantitative analysis leads to better decisions. Second, the conclusions of a careful quantitative analysis are often entirely invalidated by unforeseen developments, such as the OPEC oil embargo of 1973. Third, too much quantification may suppress creativity or even cause the premature abandonment of a lucrative breakthrough. Fourth, management may be lulled into complacency by the impressive appearance of objectivity given by the results of a detailed quantitative approach. And finally, quantitative approaches may fail to account for organizational realities and prejudices that limit the alternatives that management sees as well as the success of those they choose.

These observations are typical of those made by skeptics of the quantitative approach. Although the first two points have some merit, they apply as well to less analytical decision-making approaches. Moreover, in organizations with sustained experience, such as Roche, there is strong evidence that quantitative decision analysis methods are outperforming more traditional methods.[2] The next two points must be weighed against the benefits derived from quantification. However, the opinion of the Roche management is that these "disadvantages" have been insignificant. The final criticism is more a critique of the Western approach to decision-making than of the quantitative approach. Drucker's discussion of how the Japanese make decisions shows that they focus far more on bringing about a consensus of the essential people than their Western counterparts.[3] This helps them align organizational and individual interests, thereby

1. William Hall, "Why Risk Analysis Isn't Working," *Long Range Planning*, Vol. 8, No. 6, pp. 25–29, December 1975.
2. H. U. Balthasar, et al, "Calling the Shots in R&D," *Harvard Business Review*, pp. 151–160, May/June 1978.
   Rene Abt, et al, "The Dangerous Quest for Certainty in Market Forecasting," *Long Range Planning*, Vol. 12, No. 2, p. 52, April 1979.
3. Peter F. Drucker, *Management: Tasks, Responsibilities, Practices*, Harper & Row, New York, 1973.

*This article reports on the use of decision analysis methods at F. Hoffmann-La Roche's corporate headquarters in Basle, Switzerland. It does not imply that similar methods are used by each operating unit of the Roche group. The article discusses actual operational experience of Roche during the period 1968–1980.

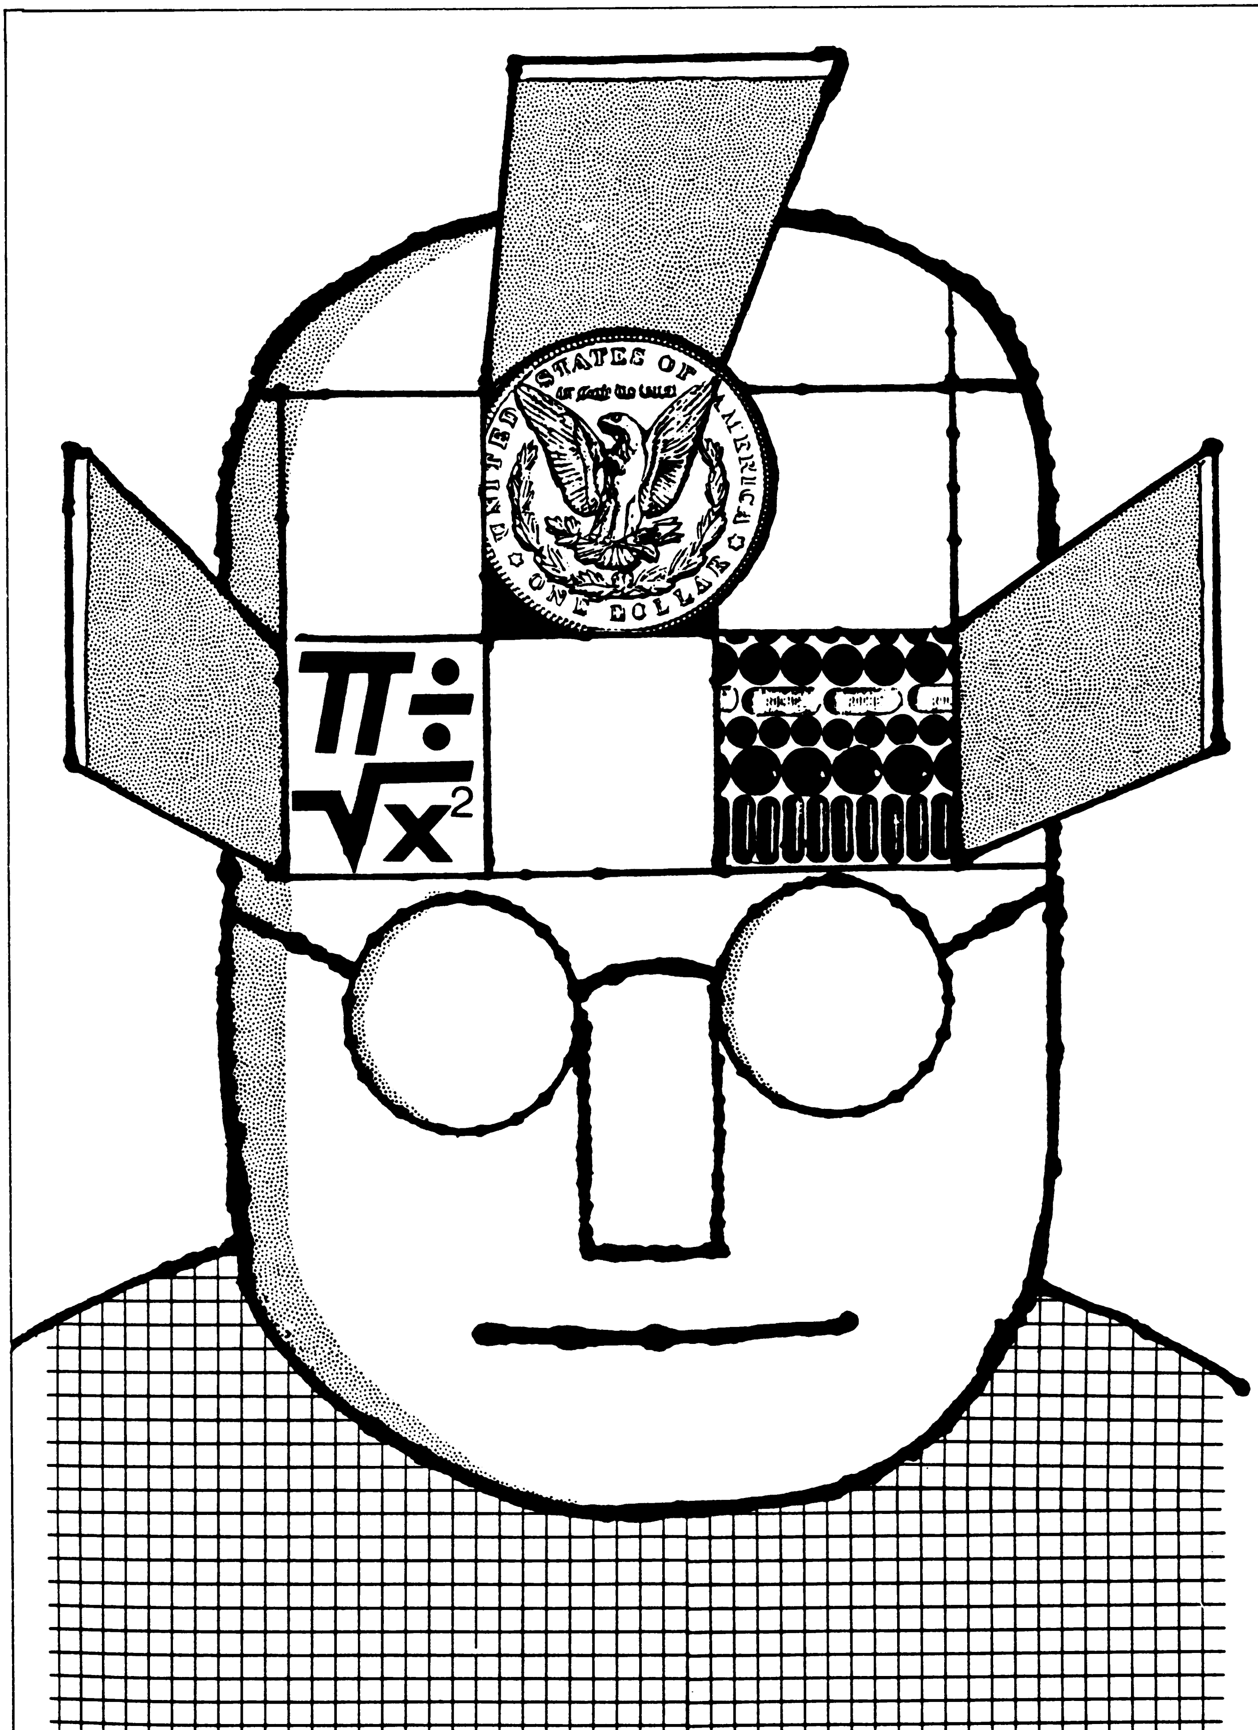
UNITED STATES OF AMERICA
ONE DOLLAR

enhancing the chance for success. The criticism is actually a pertinent observation of why the quantitative approaches that are working so well at Roche are not as successful in some other organizations.

In this article, we present eight key reasons why the analytical process at Roche has been so successful and so well accepted by managers and experts alike. We then describe the evolution of decision analysis at Roche. Finally, we explain the steps involved in the analytical process, as it exists today.

Key Reasons For The Success Of Decision Analysis At Roche

Roche believes that the use of decision analysis contributes greatly to the management decision process. The objective measure of this is that no major investment decision is taken by the General Management without a thorough quantitative decision analysis. In fact, the Roche decision analysis team is requested to perform at least one major analysis every week. The decision orginators (often daughter companies' presidents or business divisional directors) voluntarily ask for this service before pressing for a board decision.

Below, we outline eight key reasons for the success of decision analysis at Roche. We believe that many of these reasons can serve as recommendations to other organizations who wish to derive more benefit from quantitative methods.

- *Focus on Difficult Decisions and Critical Decision Factors* — At Roche, the full power of the decision analysis approach is reserved for the most difficult decisions. These may be, but are not necessarily, the biggest decisions. Decision analysis helps pinpoint the sensitive areas and critical points in a decision situation. Management is thus able to concentrate its efforts on the most difficult aspects of the most difficult decisions.
- *Emphasize Practical Results, not Techniques* — The decision analysis team meets regularly with the Roche General Management to discuss its evaluations and present its recommendations. This helps the team avoid the overly mathematical, backroom image that plagues many analytical departments today.
- *Provide Fast Service at Reasonable Cost* — The extensive automation of the analytical tools at Roche keeps service quick and analytical staff lean. In addition, the focus on difficult decisions reduces the manpower requirements and, thus, the cost.
- *Heighten Productivity of Management* — Roche management has experienced and acknowledged that proper use of decision analysis saves them time. This can be explained by the increased problem focus of sensitivity analysis, the increased problem transparency of the risk and opportunity analysis, the greater uniformity of decision preparation, and the associated improvements in communications. These benefits have become increasingly apparent over time.
- *Don't Confuse a Decision with a Forecast* — Roche management has learned to appreciate the essential difference between planning, forecasting and decision-making. Rather than trying to achieve "accurate" forecasts of highly uncertain factors and then deduce "optimal" decisions, they are searching for robust decision alternatives given the wide range of possible futures.
- *Find a "Product Champion"* — The Technical Department, the Vitamin and Fine Chemicals Division and several daughter company presidents at Roche liked the decision analysis process and therefore used it. Their interest helped to overcome the skepticism of other managers.
- *Emphasize Training at all Levels* — Over the years, there has been a strong emphasis on training Roche management in decision analysis methods. Today, more than 400 managers, including top management, have a solid grounding in these methods. This widespread dissemination of a common way of structuring and analyzing problems has resulted in greater consistency in the preparation of decisions throughout the company.
- *Involve All Experts and Managers* — At Roche, all project experts and concerned managers are involved in the decision analysis process. This allows participants to clarify their ideas and voice their opinions during the process, promoting subsequent implementation success.

It would be misleading to imply that Roche management fully accepts or completely comprehends the decision analysis results in all cases. Sometimes one or more members of the General Management are uncomfortable with the analytical recommendation. In such cases, the discrepancy between that manager's intuition and the analytical results is resolved in a private discussion. Occasionally, experienced senior managers have identified shortcomings in the quantitative assumptions or the structure of the business models. However, once identified, these shortcomings can usually be quickly corrected.

Quantitative analytical methods have thus established themselves as a valuable supplement to intuition and experience at Roche. In particular, sensitivity analysis has often identified strong and weak points that operational experience has subsequently confirmed. Moreover, risk and opportunity analysis has made it possible to quantify all known critical factors. The careful quantification of factors relating to economic viability and profitability in turn permits more extensive consideration of the nonquantifiable aspects of decision-making. At Roche, even experienced managers have found that quantitative methods can usefully support intuition in complex situations.

The Evolution Of Decision Analysis At Roche

The primary businesses of the Roche Group are pharmaceuticals, vitamins and fine chemicals. Together they account for about 80 percent of total revenue, which in 1979 was over 5 billion Swiss Francs (SFr.). In addition, the company also competes in the flavors and aromatics, cosmetics, medical diagnostics, agro-chemicals and electronic instruments markets. The Roche Group employs 42,000 people in more than 40 countries. There are 11 research and nearly 60 production centers. During the last few years, capital investments in fixed assets ran at approximately 600 million Swiss Francs per annum.

The complexity of the management can be illustrated by the existence of 250 relatively independent "daughter companies," of which 45 are important and 8 are quite large. While these subsidiaries are allowed great independence, certain centralized controls are necessary for consistent resource allocation and coordination of corporate policy. For example, all major investments are reviewed by General Management, a committee of top managers who meet weekly and have executive responsibility to the board of directors on behalf of the shareholders. One of the many responsibilities of the General Management is to review and approve over 500 credit applications annually. The heavy burden of this investment decision review created the need for a highly effective and efficient decision analysis team, reporting directly to the General Management.

Initially the decision evaluation process was developed for the analysis of major capital investment projects at Roche. This was natural since decision and risk analysis techniques were first promoted primarily for evaluating capital projects. As the process has evolved and as management has become more comfortable with the results of quantitative analysis, the scope of their requests for assistance in evaluating decisions has broadened considerably. Roche now has extensive experience with a wide range of business development decisions that go far beyond the typical capital investment analysis.

Today, the main areas of application are capital investment, product and process development, acquisitions and licensing and business/market strategy. For investments, decision analysis has been used for various capacity additions, evaluating alternative distribution systems and process efficiency improvements. For development, decision analysis has been applied to new agrochemical products, contract research for automated diagnostic systems and R&D project termination decisions. Many acquisitions and pharmaceutical license agreements have also been evaluated using decision analysis. Market and business strategy analyses involve many alternative choices for target markets, promotional policies and pricing levels, together with the associated capacity for make/buy decisions. The benefit from a systematic and comprehensive decision analysis process is especially pronounced in these complex strategy decisions.

In 1959, having realized the importance of the time value of money in evaluating long-range, strategic decisions, Roche introduced discounted cash flow (DCF) analysis. Initially the DCF return (or internal rate of return) was the preferred decision criterion because of its apparent simplicity for ranking alternatives and its intuitive appeal to Roche's financially oriented top management. As time has gone on, the DCF Net Present Value (NPV) has been used more and more, especially since it has many advantages when dealing with irregular cash flows or multiple scenarios. In the beginning the calculations had to be performed manually and consumed a large portion of the total analysis effort, even though only one "best guess" set of assumptions was evaluated.

In 1969, with the increasing need to measure the impact of deviations of the key variables on the profitability of various alternatives, Roche introduced sensitivity ayalysis. A sensitivity analysis answers a wide variety of "what if" questions (e.g., how will results be affected by a production cost increase of 15 percent or a price decline of 10 percent?). To perform sensitivity analysis on a large number of variables, with each variable having a wide range of possible deviations, requires a large computational effort which would not have been practical without computers. Automation not only served to facilitate sensitivity analysis, but also freed more time to gather and refine the fundamental input data on which any quantitative procedure relies. In 1975, the evaluation procedure at Roche was enriched by another important dimension: analysis under uncertainty. Since the success of management decision depends upon many highly uncertain uncontrollable factors, Roche added recently developed methods to quantify uncertainty and incorporate risk in the decision evaluation process. This requires the use of subjective probability judgments and decision tree analysis, both of which significantly increase the information and insight available to management. The uncertainty dimension also creates greater computational requirements, which have been met due to recent improvements in data processing speeds and storage capacities.

Today, because of decision analysis, Roche management is acutely aware of where the risks lie in a decision. Risks such as the regular supply of raw materials, prices of raw materials and products and environmental hazards are recognized in advance and therefore monitored. In any situation involving five or more factors, intuition is not considered enough, even for excellent managers. Especially in highly technical areas, such as new process development and energy savings, decision analysis has permitted the rational solution of complex problems.

The decision analysis process used today by Roche is an adaptation of that developed by Resource Planning Associates, Inc. (RPA), and others over the past 15 years.[4] Figure 1, which illustrates the building blocks of the present Roche decision analysis process, shows that the explicit treatment of uncertainty is merely the logical extension of the discounted cash flow analysis and sensitivity analysis.

The cash flow model used by Roche follows standard financial analysis methods. However, a flexible version of this model is entered on a computer with a graphic display that allows rapid deterministic evaluation of most capital investment projects as well as many more complex decisions, based on a set of numerical assumptions for critical factors. According to the nature of the specific decision, one or more of these factors may itself be calculated by more detailed sub-models.

Once expert judgment has assessed comparable ranges for the uncontrollable factors, the computer model can quickly determine the corresponding impact on the decision criteria and, in particular, whether the "best" decision changes. Many "what if" questions known to be of interest to General

4. R. A. Howard, "Decision Analysis: Applied Decision Theory," in D. B. Hertz and J. Melese (eds.), *Proceedings of the Fourth International Conference on Operational Research — 1966*, pp. 55–71, Wiley, New York, 1968.
Michael M. Menke, "Strategic Planning in an Age of Uncertainty," *Long Range Planning*, Vol. 12, No. 4, pp. 27–34, August 1979.

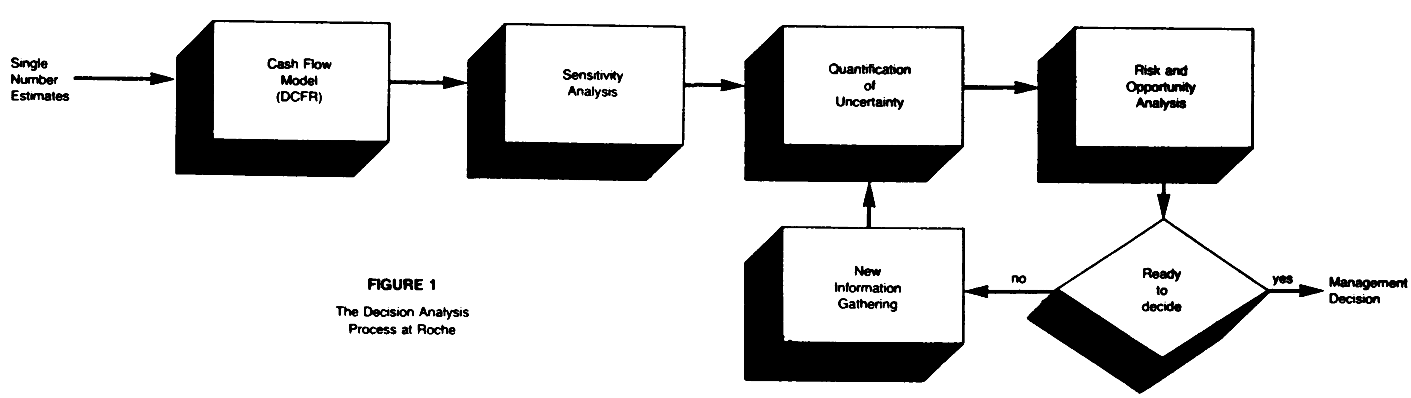

FIGURE 1

The Decision Analysis Process at Roche

Management are examined as well. In addition, breakeven analysis shows the relative sensitivity of the uncontrollables. As a result, management understands the most sensitive aspects of the decision, as well as whether one decision alternative remains superior, despite many changes in the assumptions. For the vast majority of decisions, this level of analysis is sufficient to provide a clear and robust recommendation.

However, when four or more factors are simultaneously unpredictable, the number of different conceivable combinations of factors is so large that even the most experienced executive may have difficulty appreciating the true impact of risk and uncertainty. Attempts to quantify uncertainty have led invariably to the mathematical theory of probability whose foundations were laid over 200 years ago by Bernoulli, Pascal and Fermat. However, only with the more widespread use of decision and risk analysis since the 1960s have specific techniques been developed for assigning probabilities to possible outcomes of an incertain event. The techniques that Roche now uses to encode probabilities are based on interviews with experts in a particular field. This kind of probability assessment may be a subjective judgment, but it nevertheless represents the best available knowledge. By using probability to describe uncertainty, management is able to incorporate its knowledge of what it does not know, as well as what it does know, into the decision process.

Once possible scenarios and the associated probabilities for the most sensitive variables are determined, the decision and its environment can then be structured as a decision tree. Each branch of the decision tree describes one particular scenario about the future evolution of the decision environment for one specific production alternative. Also associated with every branch is a probability, which is the product of the individual probabilities of the corresponding levels of the key variables: investment cost, price level, market growth rate and manufacturing cost. Even a relatively straight-forward decision analysis, with a small number of alternatives and only a few uncertain factors considered explicitly, can generate numerous scenarios.

The quantitative evaluation of a decision tree is what Roche calls a Risk and Opportunity Analysis. This evaluation uses the cash flow model to determine the Net Present Value (NPV) corresponding to the scenario associated with each individual branch of the tree. The profitability of each

FIGURE 2
Risk/Opportunity Profiles
for Four Manufacturing
Plant Alternatives

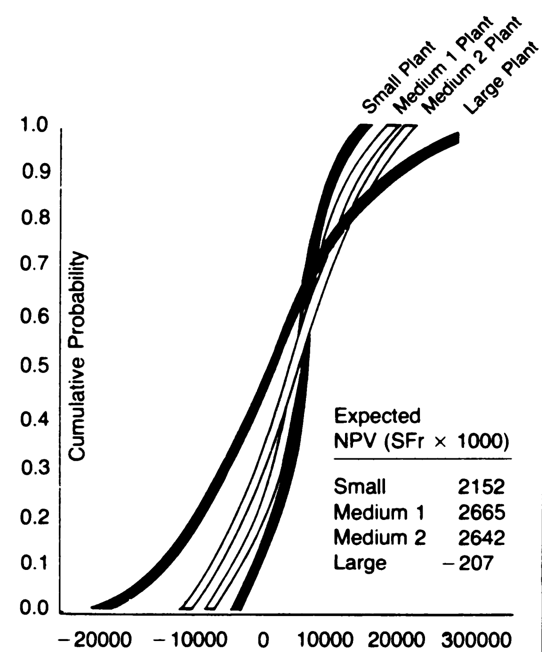

scenario can then be combined with its corresponding probability to give a risk and opportunity profile for each decision alternative. Figure 2 illustrates four risk/opportunity profiles corresponding to four investment alternatives of the decision tree. The results can be interpreted as follows:

- The smallest plant has the smallest downside risk.
- There is a 15 percent chance that future circumstances will be such that the large plant will yield the highest NPV, compared to a 50 percent chance for the smallest plant to be most profitable.
- A decision to invest in the large plant instead of the medium one would result in an average loss of 2.9 million (2.665 — (−0.207)) Swiss Francs in present value terms.

This graphic presentation of the risk and opportunity analysis in the form of a cumulative probability distribution on NPV shows the decision makers at a glance whether a given project offers an acceptable reward in relation to the associated risks. Figure 3 emphasizes the risk and opportunity balance of Medium I decision alternative. This al-

FIGURE 3
Risk/Opportunity Profile
Medium 1 Plant Alternative

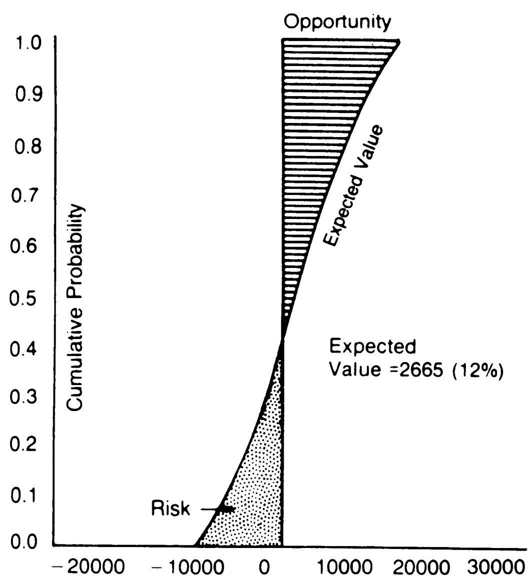

ternative is singled out for further consideration because it has the highest expected NPV and a degree of risk that can be tolerated by a company the size of Roche.

The final decision as to whether a given alternative, with its particular risk and opportunity profile, should be implemented depends upon the company's attitude towards risk. Roche has not yet established a quantitative corporate risk policy. However, over time, the Roche management committee is gaining a clear understanding of the future implications of different balances of risk and opportunity.

Before recommending a final decision, it is helpful to determine the value of additional information. By analyzing the decision tree under the assumption that the outcomes of one or more uncertain variables could be known before the decision, it is possible to measure the value of such "perfect" information. Comparing the values of perfect information on the various uncertain variables suggests which, if any, additional information should be collected before arriving at a final decision.

Conclusions

The evaluation of management decisions is an iterative process requiring specialized information from various experts throughout the organization. The information must be tested for relevance and reliability and refined accordingly. The systematic analysis of specialized decentralized knowledge is the key to a rational assessment of complex decisions. At Hoffmann-La Roche, the sensitive use of analytical procedures involves and integrates the many parties concerned with decisions and turns major top management decisions into group decisions.

Decision analysis appears to work well at Roche for all groups concerned: the General Management who has overall responsibility and who reviews most major decisions; line managers who first initiate and then implement decisions and functional specialists who provide most of the detailed judgments required to anticipate the results. There is every reason to believe that a similar style and philosophy of analytical support could also succeed in many other organizations.

In the opinion of Roche's Vice Chairman, Dr. Hartmann, "Roche is not willing to take risks simply for prestige purposes. On the average and over the long run, business ventures have to pay and it helps to support our goals by figures. Money is a scarce resource and decision analysis serves as a form of insurance which can reduce the risks of unnecessarily wrong decisions. Decision analysis allows Roche to minimize the risk of losses and, where risks are unavoidable, decision analysis helps the company to accept and control those risks from the beginning." □

Mr. Robert F. Egger is an officer of F. Hoffmann-La Roche, Basle, where he has been responsible for investment and project appraisal, decision analysis, and systems development. Prior to joining Roche in 1969, he was project manager at Alusuisse for its aluminum reduction plant in Iceland; he also has previous experience in American and Swiss manufacturing industry.

Michael M. Menke is a principal with Strategic Decisions Group, a management consulting firm specializing in strategy development and decision analysis for capital-, risk-, and research-intensive organizations. His areas of special interest include strategic management and planning, decision and risk analysis, and R&D management. Prior to joining SDG, he was manager of the decision and risk analysis program for the European office of SRI International in London.

15

MANAGING THE CORPORATE BUSINESS PORTFOLIO

James E. Matheson
Strategic Decisions Group

MANAGING THE CORPORATE BUSINESS PORTFOLIO

An important aspect of managing the modern corporation requires selecting and giving direction to a "portfolio" of business activities that will provide a balance of risk and return. Although quantitative methods have been developed during the last two decades to manage portfolios of common stocks, these methods have provided little help to the business executive, because a business portfolio differs from a stock portfolio in several important ways. The composition of the portfolio is perhaps the most important difference. Whereas a stockholder can compose his portfolio by combining small shares of many stocks, a business executive cannot do this in composing his business portfolio. To achieve economies of large-scale production or dominant market share, he must select a fairly small set of economically-sized business activities. In addition, the business executive does not have the stockholders' flexibility to easily adjust his portfolio: the executive must live with his decisions for a much longer time.

In this paper, we shall develop some new methods for balancing risk and return in the corporate business portfolio. We will define risk and return measures and show how executives can use them to build and manage a sound corporate portfolio. Since these measures are based on some key concepts of decision analysis, we will review them here. We will use simplified methods to develop graphic tools that give insight into the management of business risk. However, should they be needed in a particular application, more accurate methods are available to treat the same issues.

INDIVIDUAL BUSINESS DECISION ANALYSIS

Decision analysis of an individual business area incorporates three major factors: complexity, time, and uncertainty [1]. Complexity is treated by carefully constructing a business model that reflects all key issues. Time is usually treated by constructing a dynamic model that can project future cash flow or earnings, which is then discounted to produce a net present value of cash flow or earnings. This net present value is the primary measure of a business's value. The analyst then tests the sensitivities of the model to determine the crucial uncertainties the business faces. These crucial uncertainties are described by assessing probability distributions for them from experts on each subject.

These assessments of uncertainties on the crucial factors are combined to produce a probability distribution on the net present value, which we call the "profit lottery." A typical profit lottery is illustrated in Figure 1. The cumulative form of the probability distribution shows the probability that the net present value will not exceed (i.e., be less than or equal to) any given amount. For the example of Figure 1, there is a 0.1 probability (10 percent chance) that the net present value will be less than zero and a 0.9 probability (90 percent chance) that the net present value will be less than \$60 million. This means that there is an 80 percent chance that the net present value will be within the zero to \$60 million range.

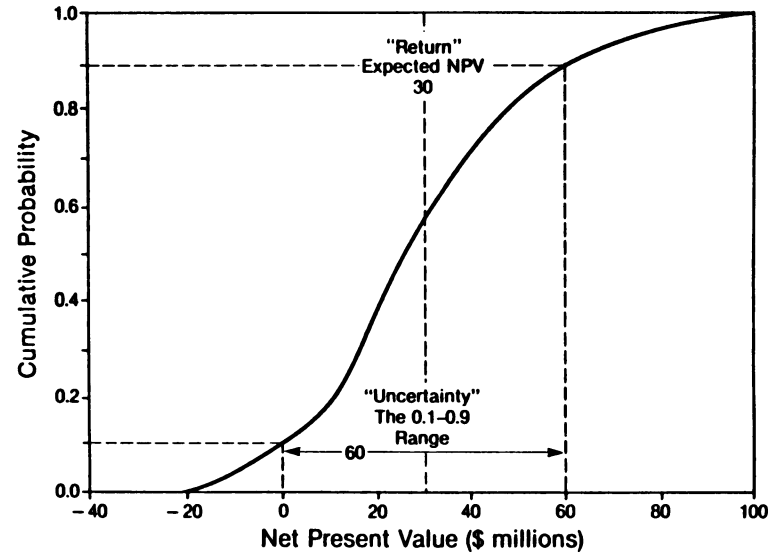

Figure 1: Probability distribution on net present value of profit, or "profit lottery," for Business A

For the purposes of this paper, we shall define the return measure (or more simply "return") as the expected (or mean) net present value. The return would, therefore, be the "fair bet" to place on the profit lottery if risk were not of concern. Unless a business provides a hedge against other corporate uncertainties, this measure of return should be the corporation's upper limit on the value of this business. A corporation wishing to avoid risk would value the business at a lesser amount, as we shall see later.

To deal with risk, we also define an uncertainty measure (or "uncertainty") as the difference between the 90th percentile and the 10th percentile of the profit lottery. The profit lottery illustrated in Figure 1 has a return of $30 million and an uncertainty of $60 million. A fundamental result of decision theory is that the risk associated with a profit lottery is approximately proportional to the square of its uncertainty [2]. If we use this risk measure (or "risk" for short), it is appropriate to trade off risk and return in proportions that represent the business's aversion to risk. Note that this also means it is not inappropriate to trade off proportions, uncertainty, and return in a linear manner. Since proving this would require a mathematical digression, we will not present it here [3].

If we plot risk and return as a point, such as that of business A in Figure 2, we can express a corporation's trade-off of risk and return by a straight line with a slope defined as the corporation's "risk tolerance." Tracing this line down to the return axis yields a point called the "certain equivalent." This point represents a business having zero risk that the corporation would see as equivalent to business A (in the sense of an exchange or minimum selling price). In fact, any business along the risk tolerance line passing through the point for business A would be equivalent to business A and would have the identical certain equivalent. It is easy to see that a lower risk tolerance (slope) would weigh the risk more heavily and result in a lower certain equivalent, while a higher risk tolerance would do the opposite. An infinite risk tolerance would be represented by a vertical slope resulting in a certain equivalent equal to the return only, an attitude we call "risk neutral."

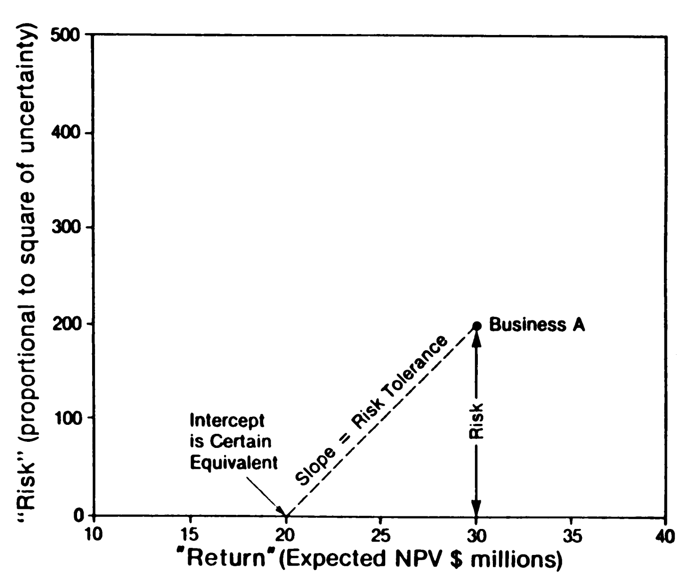

Figure 2: Risk, return, and certain equivalent of Business A

In this example, business A has a certain equivalent of $20 million. This can be viewed as a deduction of a $10 million risk penalty from the $30 million expected value. In fact, the geometry of the straight-line trade-off does exactly this. In word equations,

$$\text{certain equivalent} = \text{expected value} - \text{risk penalty}$$

where

$$\text{risk penalty} = \frac{\text{risk}}{\text{risk tolerance}}$$

or

$$\text{certain equivalent} = \text{expected value} - \frac{\text{risk}}{\text{risk tolerance}}.$$

We can give further meaning to the risk tolerance by applying the graphic analysis to several independent ventures [4], as illustrated in Figure 3. If we construct a line from the origin of the graph with a slope equal to the risk tolerance, it divides the graph into "accept" and "reject" regions. Any business to the right of this dividing line will have a positive certain equivalent, while any business to the left will have a negative one. Also, the ratio of risk to return for any business to the right of the dividing line will be less than the risk tolerance, and that of any business to the left of the dividing line will be greater than the risk tolerance. Thus, the risk tolerance can be interpreted as <u>the maximum tolerable ratio of risk to return</u>. In practice, one way to establish the risk tolerance is by questioning corporate executives about the acceptability of a set of ventures. Their answers, which they may have to adjust for consistency, form the risk tolerance dividing line.

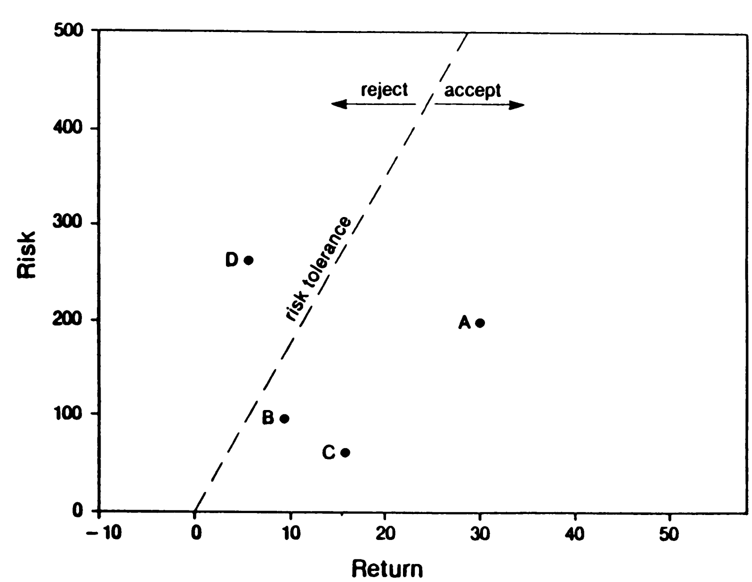

Figure 3: Accept and reject regions defined by the risk tolerance

BUSINESS PORTFOLIOS

Decision theory also shows us that the risk and return of independent ventures, or business areas, may be added together to form the risk and return for the portfolio they compose. Figure 4 shows an example of three independent business areas combined into a portfolio. This independent portfolio serves as a reference point for other situations.

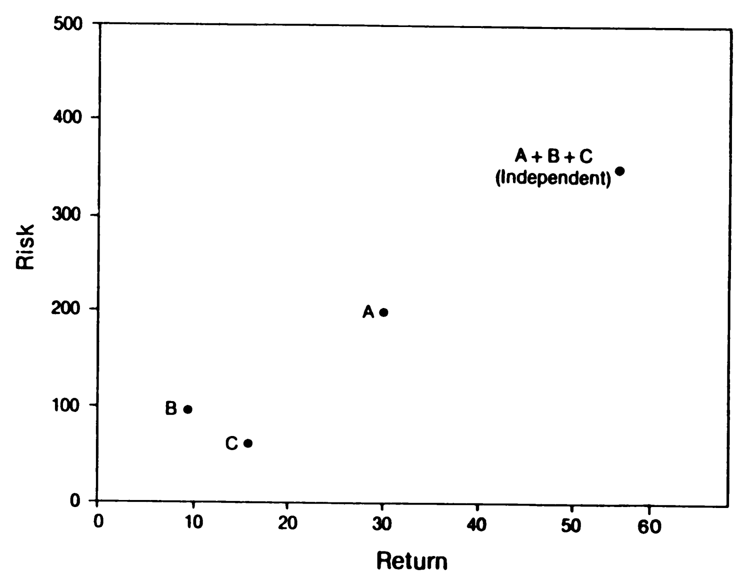

Figure 4: Construction of the independent business portfolio

The business areas could have return synergies or dissynergies, for example, resulting from the sharing of common facilities or the competing for common facilities, that cause the sum of their returns to be greater or less than that of the independent point shown in Figure 4 (without changing the sum of their risks). Figure 5 illustrates these possibilities.

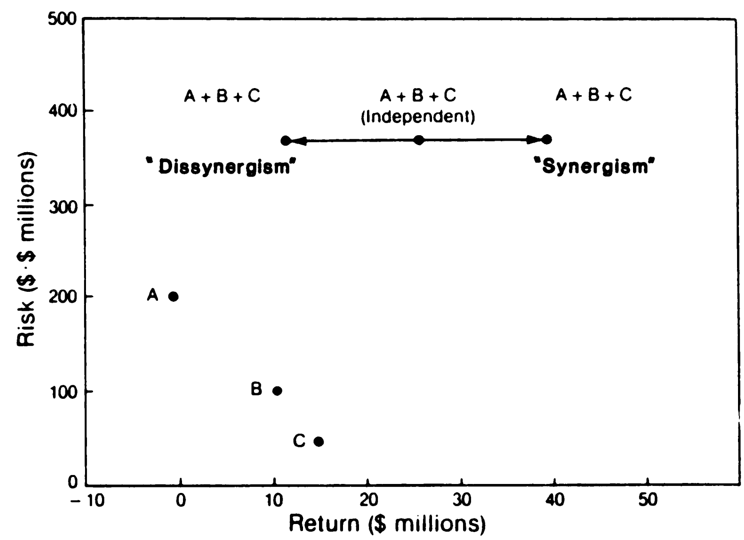

Figure 5: Synergism and dissynergism

Also, the business areas could be influenced by the same or related uncertainties, making their profits either positively or negatively dependent (without changing the sum of their returns). If the business area profits tend to rise and fall together, this increases the portfolio risk. We call this effect "risk concentration." If the business area profits tend to go in opposite directions, this decreases the project risk. We call this effect "risk compensation." Figure 6 illustrates these two possibilities.

Dealing with return synergies and dissynergies primarily requires examining common resources and other common factors. Dealing with risk concentration and compensation requires thoroughly understanding the influence of uncertainties on all the business areas. The remainder of this paper will pursue the treatment of risk.

ANALYZING PORTFOLIO RISK

We will illustrate some methods for treating risk with an example based on a real but disguised case. The organization chart for a successful diversified energy company called "Lamarco" [5] is shown in Figure 7. Historically, Lamarco was a gas transportation utility business. To reduce its dependence on one business, it recently undertook a major diversification program. Unfortunately, it could not tell if it had diversified enough.

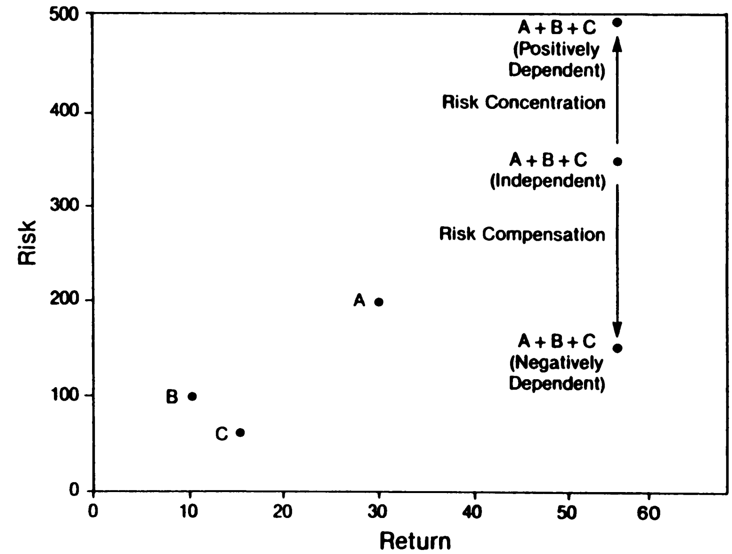

Figure 6: Risk concentration and compensation

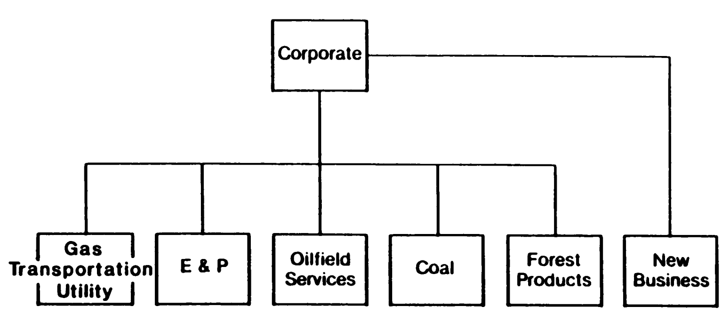

Figure 7: Lamarco's organization chart

Each of Lamarco's businesses was analyzed to determine its risk and return. These are plotted in Figure 8. As expected, the businesses having higher returns also have higher risks.

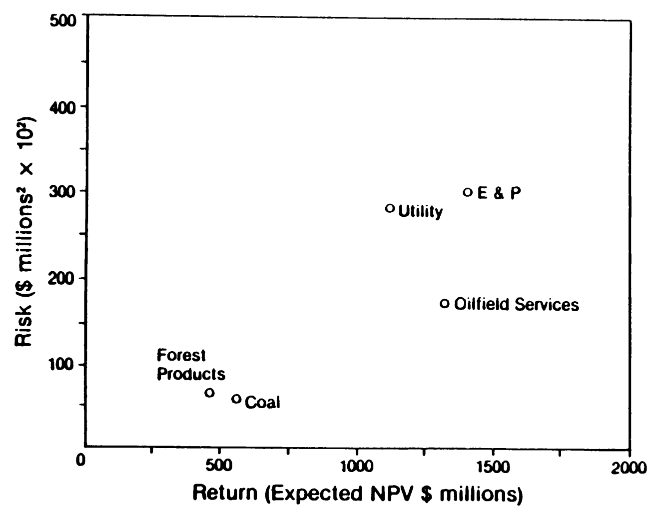

Figure 8: Risk and return points for Lamarco's businesses

Analyzing the sources of uncertainty for each business revealed several common uncertainties. The most critical ones were oil price, inflation, and regulation. In addition, each business had independent uncertainties that were unrelated to the other businesses. The critical uncertainties are portrayed in a simplified probability tree in Figure 9. You can see that each uncertainty can take on two values in each situation. More values could be used if additional details were required. Probabilities of each value are written near each branch point of the tree, and overall scenario probabilities are obtained by multiplying the three probabilities along the paths defining each scenario.

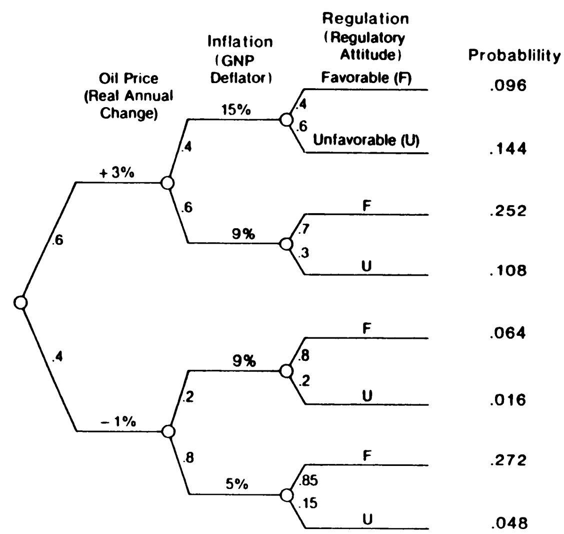

Figure 9: Scenarios for the common uncertainties in Lamarco's businesses

The analysis of each business was carried out for each of the common scenarios. The returns of each business and of the total portfolio are shown for each common scenario in Figure 10. We can easily see risk concentration among the first four businesses; for example, all their returns are very high for the first scenario. We can also see a little risk compensation produced by the forest products business, which has its lowest return under the first scenario and its highest under the seventh scenario. By doing further analysis, we can use the common uncertainties to determine the probability distribution on total profits, which is shown in Figure 11. When the risk and return of the total corporate profit is plotted, as seen in Figure 12, we see a case of high risk concentration.

Lamarco faces much more risk than a corporation made up of similar independent businesses. In fact, the corporate ratio of risk to return is higher than that of any of the individual businesses. A few computations produce the comparisons seen in Figure 13. Although Lamarco's diversification had successfully lowered its exposure to regulatory treatment, the diversification had also increased its exposure to inflation and oil price uncertainties. These observations create new strategic challenges for reshaping Lamarco's corporate portfolio.

Returns

Expected Net Present Value ($ millions)

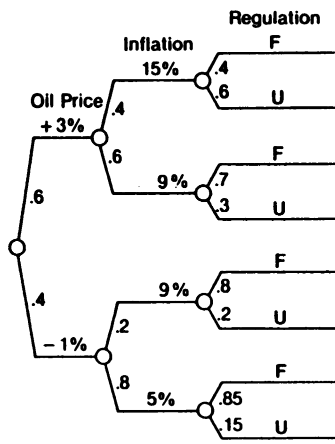

Utility	E & P	Oilfield Services	Coal	Forest Products	Total
1512	1713	1561	636	334	5756
1042	1712	1563	589	336	5242
1398	1549	1458	628	469	5502
928	1543	1450	579	469	4969
1189	1260	1214	598	498	4759
839	1269	1215	543	498	4364
1005	1098	1097	487	557	4244
598	1098	1096	443	554	3789

Figure 10: Effect of the common uncertainties on each business

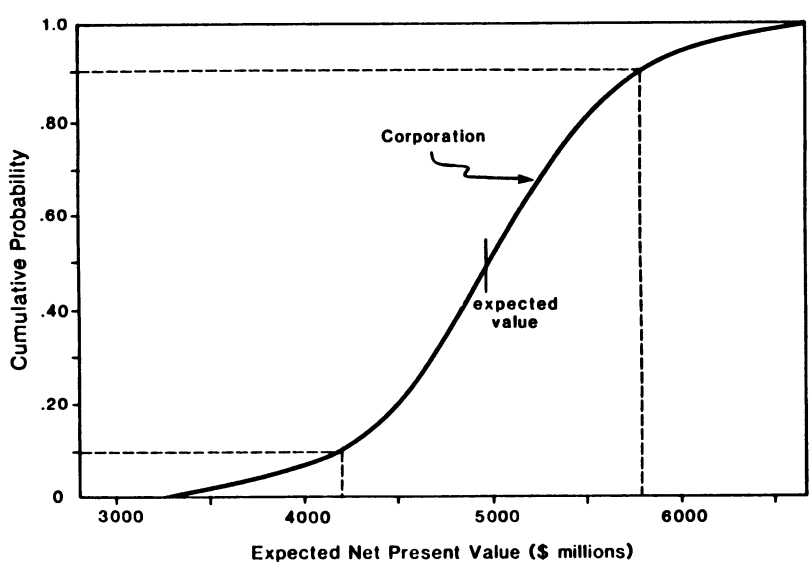

Figure 11: Probability distribution on Lamarco's total net present value

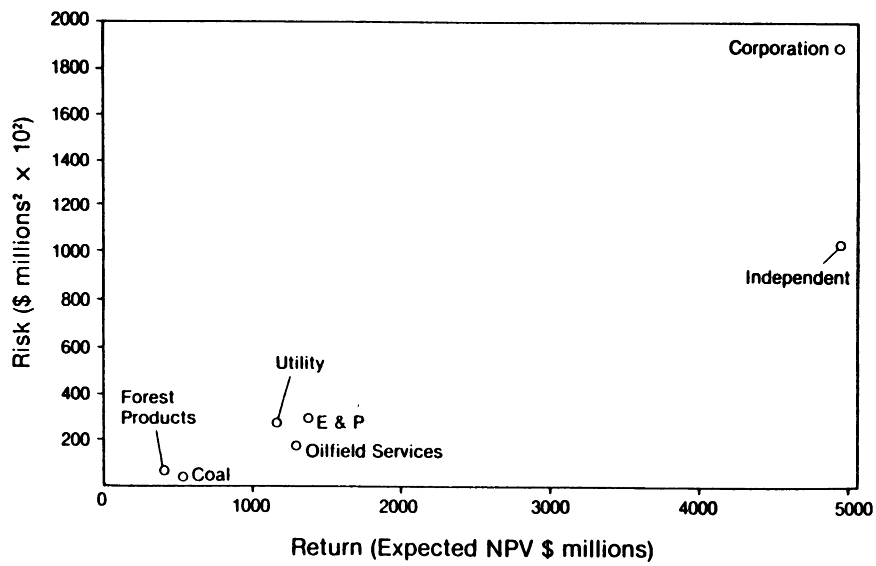

Figure 12: Lamarco's business combined into a corporate portfolio

Regulation Unfavorable 4884 — Regulation Favorable 4968

$\underbrace{\qquad\qquad\qquad}_{\$84\text{ million}}$

Oil Price Rate

−1% 4276 — 3% 5384

$\underbrace{\qquad\qquad\qquad}_{\$1108\text{ million}}$

Inflation Rate

5% 4176 — 15% 5448

$\underbrace{\qquad\qquad\qquad}_{\$1272\text{ million}}$

- The company has successfully diversified away from exposure to unfavorable regulatory treatment.
  - —The conditions that are likely to lead to unfavorable regulation have compensating effects on other businesses.
  - —Further diversification from regulation is not necessary.
- But, the company is exposed to slower rates of inflation and oil prices.

Figure 13: Sensitivity of Lamarco's performance to individual uncertainties

We can gain additional insight into portfolio development strategy by understanding the quantitative impact of the underlying uncertainties producing the risk (and risk concentration). Figure 14 shows the impact of the crucial uncertainties on each business. The width of the bars is proportional to the impact on returns as each uncertainty is varied from its 10th percentile to its 90th percentile. If returns rise with this variation, the bar is plotted to the right, and if returns fall, the bar is plotted to the left. The risk of any business area, which is listed (along with the return) below the bar, is approximately proportional to the sum of the squares of each uncertainty bar in the column [6].

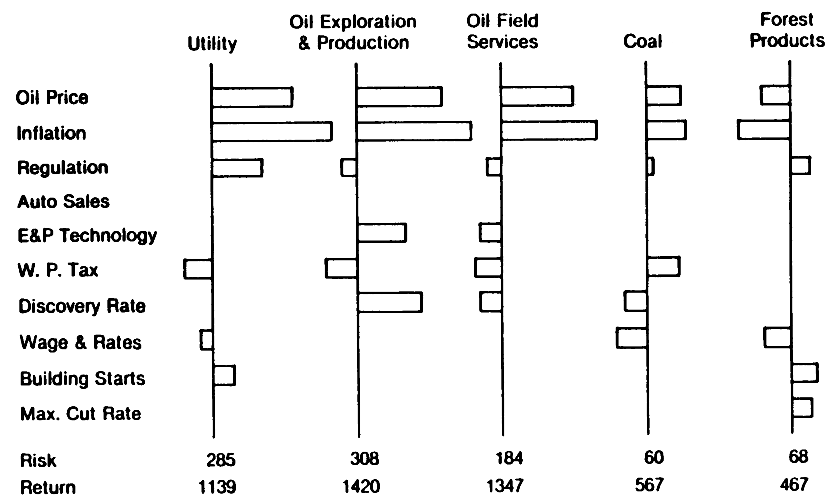

The width of the bars represents the uncertainty on the variable.

Figure 14: Impact of each uncertainty on business returns

Since we are dealing with a case of risk concentration (non-independence), we cannot simply add the risks across business areas. However, we can add the uncertainty bars to form an uncertainty profile of the total business portfolio, which is shown in Figure 15. In adding these bars, some uncertainties tend to concentrate (e.g., oil price and inflation) while others tend to compensate (e.g., regulation). This addition and cancellation is the source of risk concentration and compensation, which can be dramatically seen in Figure 15. Most importantly, this process lets us graphically see the risk problems of the portfolio. In Lamarco's case, oil price and inflation uncertainties are clearly creating most of the corporate risk.

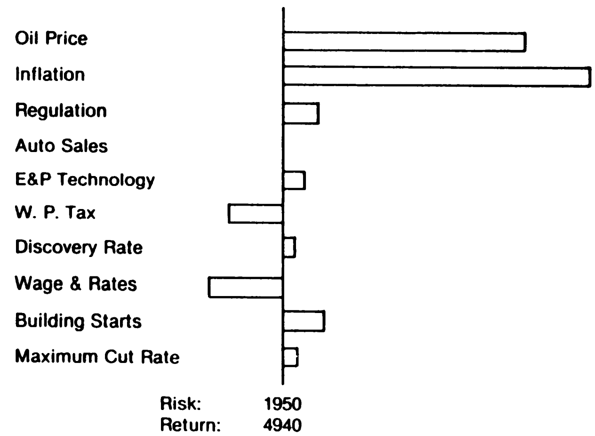

The portfolio risk is found by adding the business area uncertainties for each variable and summing the squares.

Figure 15: Impact of each uncertainty on business portfolio returns

Once we understand the sources of business risk, we are in a good position to create ways of reducing it. In the Lamarco case, we would challenge Lamarco executives and their analysts to come up with ways to reduce this risk while maintaining the returns. One possibility would be to manage each business in ways that reduce oil price and inflation risk in return for increased risk exposure to the less dominant uncertainties (or even a calculable amount of reduced return). Another possibility would be to create or acquire a new business having risk-compensating characteristics. In Figure 16, for example, we see the effect of adding a hypothetical aluminum business that has negative sensitivity to oil price and inflation. If we assume that Lamarco buys the aluminum business for its expected present value, it would have zero net return. However, it would still be valuable to Lamarco in reducing risk. Figure 17 shows the potential acquisition on a risk-return plot. Adding the aluminum business reduces the risk to about that of the independent case without giving up any return. As a result, it increases the certain equivalent by about $400 million.

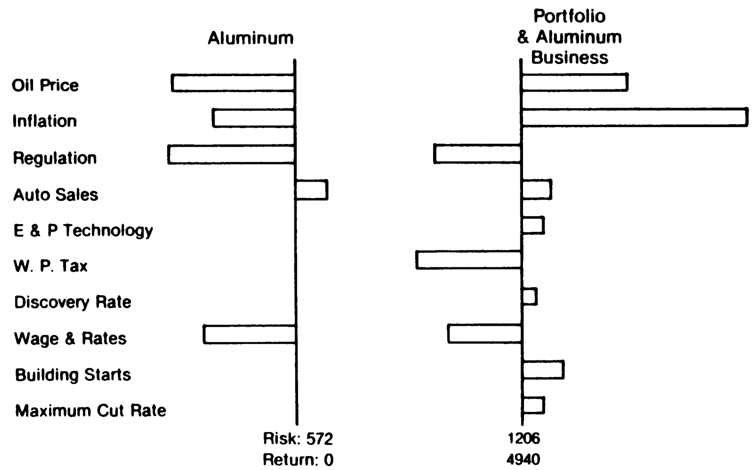

The returns are expected to be negligible. Does such an acquisition make sense?

Figure 16: Effect of adding an aluminum business to Lamarco's portfolio

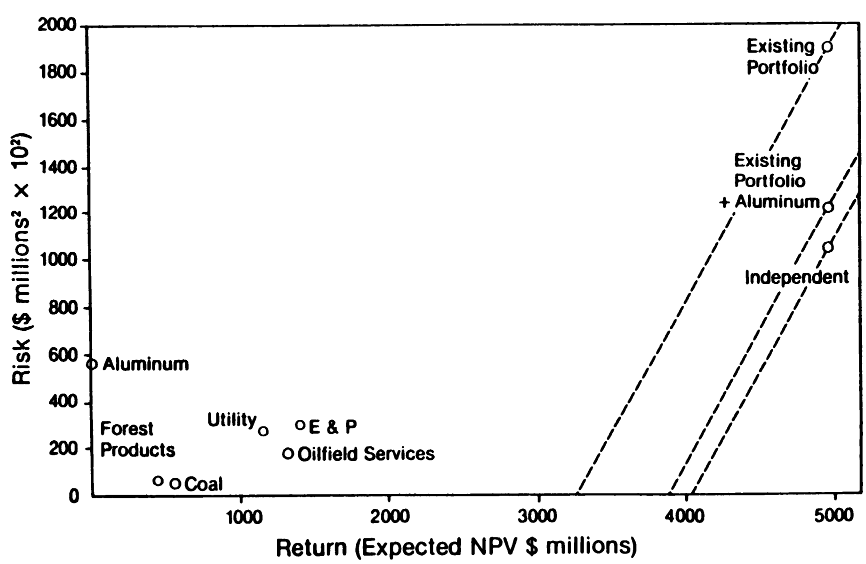

Figure 17: Reduction in risk due to a potential aluminum business acquisition

Lamarco could also consider divestitures. Figure 18 shows how Lamarco's corporate portfolio would change if each business were divested by giving it away. If it were sold at a specific price, this amount would have to be added to the return, thus shifting its point to the right. Depending upon buying and selling prices, divesting coal, acquiring an aluminum business, or possibly doing both could considerably improve Lamarco's situation. These graphic methods stimulate creativity and allow quick screening of ideas for developing better corporate portfolios.

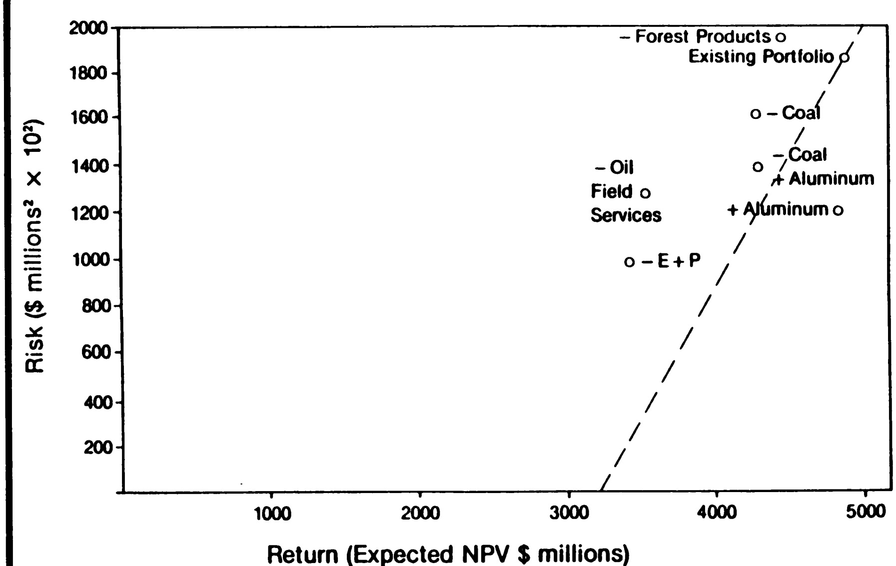

Figure 18: Effect of divestiture on Lamarco's risk and return

CONCLUSIONS

We have shown a new method for managing the trade-off of risk and return in the corporate business portfolio. This method incorporates graphic displays that identify the sources of business risk and shows how business units combine to either concentrate or compensate risk. The insight developed from a business portfolio analysis allows the executive to give better direction to each business unit and to plan acquisitions and divestitures that enhance the risk-return ratio for the whole corporation.

FOOTNOTES

1. Ronald A. Howard, "Decision Analysis: Applied Decision Theory," Proceedings of the Fourth International Conference on Operational Research, Wiley-Interscience, New York, 1966, pp. 55-71.

2. The risk measure is technically defined as one-half of the variance of the profit lottery. For most "ordinary" profit lotteries, the risk measure will be approximately one-thirteenth of the square of the uncertainty measure.

3. An example furnishes some intuitive justification. Suppose we could invest in some multiple, M, of the profit lottery of Figure 1. Then, we would have

$$\text{return} = 30 \times M,$$

$$\text{uncertainty} = 60 \times M,$$

$$\text{risk} = 3600 \times M^2,$$

$$\frac{\text{uncertainty}}{\text{return}} \quad \frac{60 \times M}{30 \times M} = 2,$$

$$\frac{\text{risk}}{\text{return}} \quad \frac{3600 \times M^2}{30 \times M} = 120 \times M$$

   Since the ratio of uncertainty to return is constant, a corporation using this ratio that liked the original profit lottery would be led to invest in any multiple of it, an absurd result. However, the ratio of risk to return increases as M increases. For a sufficiently high value of M, the corporation using this ratio would limit its investment because it would see that the risk to return ratio would exceed its ability to tolerate risk for a sufficiently high value of M. Intuitively, we would prefer the second ratio.

4. By independent ventures, we mean that the corporation could independently accept or reject each of the ventures, and that knowledge of specific results from any of the ventures would not influence our opinion (profit lottery) for any of the other ventures or other holdings of the corporation: that is, they are probabilistically independent.

5. The author would like to acknowledge Terry Braunstein for carrying out the "Lamarco" case that provided the stimulus for this work.

6. If the uncertainties are correlated, products of the uncertainty bars multiplied by the appropriate correlation coefficient should be added into the sum of squares. Also, this approximation assumes an approximately linear profit model.

16

OVERVIEW OF R&D DECISION ANALYSIS

James E. Matheson
Strategic Decisions Group

OVERVIEW OF R&D DECISION ANALYSIS

By James E. Matheson

Uncertainty, complexity, and long lead time are fundamental characteristics of research and development. Breakthroughs cannot be forced, and unexpected results are often the most valuable. Molding new results into new and better products is a lengthy and complex task. And the process of capitalizing upon new ideas is still subject to uncertainties in development, production, and market acceptance. Although these features make R&D an exciting and rewarding business, they simultaneously create logically difficult management problems.

The direct results of R&D are rarely useful by themselves; the real reward comes later as the results are embodied in products and services. Therefore, logical R&D management must comprise two complementary activities: the *economic evaluation of potential research results* and the *selection of R&D activities likely to produce valuable results.*

Since R&D activities necessarily precede the use of the research results, management attention is often too narrowly focused on the research activities themselves. It is all too easy to fall into the trap of pursuing "good research," as perceived by the R&D community, and of assuming that successful implementation will follow. However, the R&D process can be effectively managed and motivated only if management properly devotes its attention to the value of the potential results. Figure 1 shows the logic for R&D management.

The R&D activities transform R&D possibilities into R&D results, but for ultimate success, these results must be adapted and adopted by users to transform them into realized value. Importantly, the "flow of value" is in the opposite direction. Given the user's ultimate values and range of decision alternatives, an economic value can be placed on potential R&D results. This step is called the *economic evaluation of potential R&D results.* The next step, which involves combining these values for R&D results with the possible ways of carrying out R&D activities to achieve these results, is called the *development and selection of R&D programs.*

All managed R&D activity, whether carried out publicly or privately, can be regarded from this vantage point. The details of the methodological approach, however, depend significantly on whether the sponsor of the R&D activity is also the user of its results. Most governmental R&D is carried out to achieve public research results, which are then capitalized upon by various user groups. Thus, the government has little control over user decisions and receives little of the user value directly. Therefore, the government R&D manager logically should carry out an economic evaluation of all the public benefits. On the other hand, when a private company carries out R&D for incorporation into its own products and production facilities, it is both the sponsor and the user. In this case, the R&D manager can carry out a somewhat simplified economic evaluation because he can better anticipate his own future user values and decisions.

R&D programs directed toward a specific implementation target, such as a new or improved production method or product, are easily approached by decision analysis. The decision analysis should begin by establishing the range of R&D results that might be achieved, as shown in Figure 2.

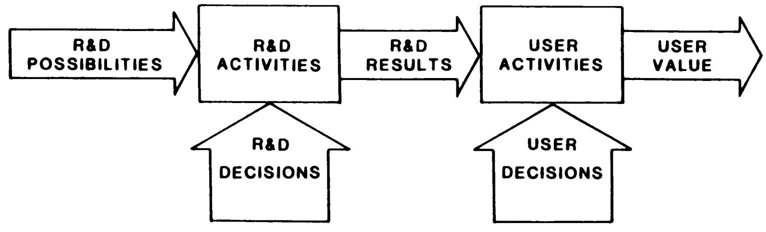

Figure 1: R&D Management Logic

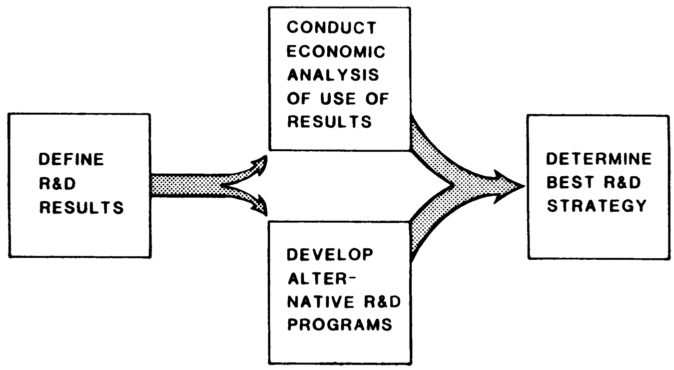

Figure 2: R&D Decision Analysis

Establishing these results allows temporary separation of the analysis into two tasks: the *economic evaluation of potential R&D results* and the *development of alternative R&D programs* that might create these results. A preliminary economic analysis will provide guidance for developing R&D programs by establishing both the economic value of varying degrees of research success and the value of time in attaining results. The latter may often be indicated by competition for, or changes in, user needs. Both of these values help show where parallel R&D efforts are required to increase the chance of achieving the results in a timely manner. Developing alternative R&D programs must include identifying major uncertainties in achieving results, determining the more elemental tasks and results that compose the overall result, and allocating effort to program elements. Developing step-by-step R&D strategies, where crucial activities are undertaken first and successive activities are undertaken only after initial success, can lower the cost and risk of R&D. However, more expensive and "risky" alternatives must also be developed to be weighted against the need for timely R&D results.

In the last task shown in Figure 2, the alternative R&D programs are combined with the economic evaluation to obtain the overall "bottom line." Because of uncertainty in both the R&D results and in using the results (e.g., competing developments, user perception), the alternative R&D programs will have varying degrees of expected values and uncertainty. Large organizations may have the ability to pursue the plan with highest expected value. When R&D resources are more limited, however, an organization may desire to limit risk by choosing an alternative having a lower expected value but more assurance of return.

The decision analysis should not, however, stop with this choice, because often there are sources of information that may reduce some of the crucial uncertainties and allow an organization to make a better choice. For example, a more careful study or survey of potential users may reduce the uncertainty in successful implementation and provide better direction for detailed research objectives. Also, the research process itself, as well as external events, will continually produce new information that should be used to update the analysis and reoptimize the R&D program. In a lively R&D area, even the options to cancel must be reviewed regularly.

R&D programs having less definite paths to the user can be (and indeed have been) analyzed on a similar basis. For example, general research to support advancing technology in a line of business can be treated by capturing the generic features of existing and new products. The economic value of improvements in these features on the whole product line is estimated in the economic evaluation. As guided by the economic evaluation, an R&D program is then developed to provide overall support to the business.

Even basic research can, and should, be treated using this approach. In this case, the approach would rely more on human judgment because the research results are more loosely defined (for example, by "a significant breakthrough in Area X"). The economic evaluation is based primarily on the judgment of R&D managers, R&D producers, and implementation experts. Based on additional judgments of researchers, managers, and outside experts about the probability of success, R&D programs are devised to achieve a high probability of valuable results. The judgment might be broken down into the potential of the research area, the potential of the research team, and the potential of the research strategy. The logic is the same as before, but the analysis has quite a different character.

When all of an organization's research activities have been analyzed in these terms, they can be combined into a portfolio of activities and the whole portfolio can be profitably analyzed. We will mention only a few aspects of the problem here. R&D programs may compete for resources such as funds, researchers' time, and research equipment. They may complement each other by sharing research elements and facilities, and, more subtly, by reducing the risk of a whole product area. Also, most organizations desire a portfolio of R&D activities at various stages of maturity to provide them and their users with a steady supply of new accomplishments. Furthermore, there is always the issue of whether it would be better to risk resources on a single R&D venture or to use the same resources to support several smaller ones. An R&D portfolio analysis addresses all these questions.

Because it integrates the R&D activity with the needs and desires of the ultimate user and because it provides a vehicle for continual monitoring and updating of the R&D program, decision analysis can provide the control needed to manage the complex, dynamic, and uncertain R&D process.

17

USING DECISION ANALYSIS TO DETERMINE R&D'S VALUE

From Inside R&D's Monthly Report on Managing Innovation

Reprinted with permission from Inside R&D, March 9, 1983, pp. 3-4,
and April 13, 1983, pp. 3-4.

INSIDE R&D's monthly report on Managing Innovation

USING DECISION ANALYSIS TO DETERMINE R&D's VALUE
(Part I)

More and more, senior management worries about whether it's getting maximum value out of its R&D investment. Placing a value on R&D is difficult because R&D decisions are characterized by uncertainty, long time horizons, and a number of other characteristics that make them resistant to traditional financial evaluation methods.

In their work with such companies as Alcoa, CIBA-GEIGY, Exxon, Hoffman-LaRoche and Westinghouse, Michael Menke and his colleagues at Strategic Decisions Group, Menlo Park, CA, have found that these evaluation difficulties can be counteracted, if not actually eliminated, by the use of decision analysis.

Very briefly, this process involves constructing influence diagrams and decision trees that illustrate the chronological sequence of decisions and uncertainties affecting a particular project or projects. By providing a logical and consistent means of visualizing uncertain, complex problems, Menke considers decision trees ideal for describing to top management real decision problems that, to quote one R&D vp, are "too difficult to solve in my head and too important to solve in my gut."

Menke explains how decision analysis mitigates the factors that make R&D decision-making so difficult:

• Uncertainty is the hallmark of R&D, which is marked by both technical and commercial unknowns. Moreover, technical success can rarely be defined without considering the commercial environment. Unfortunately, most traditional valuation procedures do not explicitly incorporate uncertainty. However, the necessarily uncertain subjective judgments that experts (managers, engineers, scientists, consultants, etc.) are called upon to make in these areas can be quantified by using probabilities. These so-called subjective probability judgments have demonstrated their reliability as indicators of future technical and business success in many industrial situations, as Menke explained in "Managing Innovation," *Inside R&D*, 1/9/80.

After an influence diagram has identified the sources of uncertainty, sensitivity analysis ranks the importance of the various uncertainties, thus determining which ones are the most crucial to the problem at hand. Once probability has been used to help people quantify their judgments about the most crucial factors, a decision tree can be constructed in which the various branches represent different scenarios that combine the various technical and commercial possibilities.

In essence, the entire structure of decision analysis is designed to deal with uncertainty, Menke asserts.

• Long Time Horizons. Rather than ignore speculations about the success of a project that won't be observable for years, one can construct scenarios which indicate the value of a project in different possible futures corresponding to different degrees of technical and commercial success.

As for the serious problem of determining the future financial return necessary to offset an expenditure today, this is complicated by considerations of inflation as well as the inherent technical and commercial risks of R&D. Menke points out that a number of companies have overcome this difficulty by separating the true return desired by investors from both inflation and project risk. Without doing this, one will end up valuing only short-range projects, he warns.

Menke's recommendations: (1) strip away inflation by calculating costs and benefits in constant dollars; (2) treat risk not by applying a higher discount rate (as is often done) but by computing the actual probabilities of success and failure; (3) demand that a project achieve on-average a true, risk-free, inflation-free return of 3-5%, as measured by the expected net present value obtained by averaging over all the scenarios in the project's decision tree. This procedure avoids the short-range bias of traditional financial techniques, but does require that resources invested in R&D earn a sufficient return.

• Poor Communication Between R&D and Marketing. More than anything else, Menke insists, this arises because uncertainty is not treated explicitly, as described above, and the problems of risk and long time horizons are confused by allowing inflation and risk to enter into the setting of "hurdle" rates.

"The stated goal of decision analysis is to assign a value to the R&D project. The underlying goal is to use that value to try to surface all the critical issues and thereby make the communication more open, more honest, and more effective," Menke says.

• Many Alternatives. Lots of different R&D projects compete for limited resources, and any one project could be managed in many possible ways. Decision analysis provides a powerful way of determining the best development strategy for an individual project. Then, at a higher level, it can be used for managing a "portfolio" of projects, weighing one against another to determine the optimum mix for the organization.

• Complex Interactions. Projects are not independent of one another. For example: common aspects of R&D processes can be shared among several projects, thus reducing costs; the commercial value of two or more jointly successful projects may be greater than the sum of their individual commercial values; the technical success of a new (breakthrough) product may make obsolete current products or less sophisticated products under development; information pertaining to the success or failure of one project may influence the assessment of technical success on other projects; several parallel projects increase the likelihood that at least one will succeed technically; a successful R&D effort in your main area of business may increase the overall vulnerability of the business (by, for example, producing an improved product that depends upon a raw material from an unstable foreign source).

These and many other similar interactive effects can be accounted for in decision analysis, either as modifications to the probability factors or to the value factors at the end of the decision tree. On the other hand, they are very hard to capture with conventional financial analysis.

• The Dynamic and Sequential Nature of R&D. Menke characterizes R&D as an area of high uncertainty but low risk. The reason is that R&D decisions are usually made, and resources committed, in stages over a long period, rather than all at once. Consequently, it's difficult, for example, to determine the value of initial work on a group of related chemical compounds for a particular purpose when only one of these compounds will actually reach the marketplace--even if R&D produces two or three that are technically successful.

With a decision tree, however, one can examine, first, the decision on the various funding levels or directions of the applied research program and the different probabilities of achieving the desired results. This, in turn, leads to decisions about the kind of development program to undertake. Aim of the program, in turn, is to deliver a product or process that can be commercialized at a particular time, cost and level of performance.

This interplay of sequential decision with the probability of failure at each stage in the chain is, again, difficult to handle with normal financial evaluation methods. But decision analysis is eminently suited to the task.

• Difficult Value Issues. Besides worrying about market share, profits, production costs, and the like, management faces a host of more subtle issues. Can we change the corporate culture to compete in new business areas that R&D has opened up for us? Are we better off licensing this new development? How will it affect our image as a technical leader? And so on. Decision analysis allows a company to take a broader view of its R&D effort by incorporating such factors into a decision tree.

• Breakdown of Traditional Risk Criteria. Many of the traditional rules for calculating business profitability simply break down when applied to the complex, long-range uncertain world of R&D. And R&D people, traditionally suspicious of efforts to "quantify the unquantifiable," don't make things easier.

A prime example is the rule that risky R&D projects should require high hurdle rates. This rule breaks down--and is a continual source of friction between R&D and top management--because of the conflict between uncertainty and the long time horizon. The high degree of uncertainty encourages the financial types to raise the hurdle rate because of the perceived risk. But the longer the time horizon, the harder it will be to earn that hurdle rate. All too often, then, the application of discounted cash flow analysis will wash out long-range R&D in favor of the shorter range, "safer" projects.

(In the April 13 edition of "Managing Innovation," Menke will outline how the R&D decision analysis process can be applied.)

INSIDE R&D's monthly report on Managing Innovation

USING DECISION ANALYSIS TO DETERMINE R&D's VALUE (Part II)

Decision analysis is a way of dealing with many of the characteristics of R&D that make it so difficult to evaluate by traditional financial analysis. Last month (*Inside R&D*, 3/9/82), Michael Menke identified those characteristics (uncertainty, long time horizons, etc.) and explained how decision analysis provided the necessary techniques for evaluating alternatives despite those difficulties.

This month, Menke outlines briefly the actual decision analysis process as it would be applied in a practical situation by his firm, Strategic Decisions Group.

SDG's decision analysis cycle has the four stages shown in Fig.1.

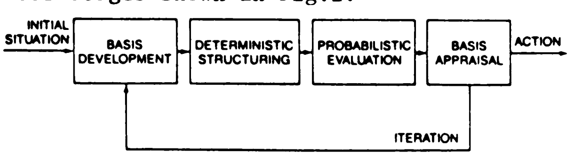

Fig. 1

(1.) In Basis Development, Menke explains, the task is fundamentally to identify the real decision. In this stage, a number of techniques are used to assess qualitatively the values, information and alternatives that are critical to making a sound choice. Drawing an influence diagram, for instance, is the final step of Basis Development. It is developed from interviews with key personnel and shows in a systematic way the important uncertainties and decisions that will affect achievement of the project goals.

(2). Deterministic Structuring treats the decision situation quantitatively and determines the critical uncertainties. It reveals, for example, how the different variables on the influence diagram affect the project's profitability, and how project cash flow may be increased or decreased as a result.

As with most financial evaluation procedures, discounted after-tax cash flow is used to determine a net present value for each scenario, but with one important difference: SDG separates inflation, technical risk and market risk from investors' real rate-of-return, ending up with a lower 3%-5% real discount rate that investors in any large company have every right to demand.

After developing a simple project-specific model to calculate the cash flow and net present value for any possible scenario, SDG then does an extensive sensitivity analysis to determine which of the many complicating factors (price, technical performance, development time, market share, etc.) creates the most uncertainty in the situation and is, therefore, most critical to making a good decision.

Sometimes the analysis will end here. You may not know the precise value of the project, but you will know that it is valuable enough to be worth continuing without further analysis. On the other hand, the possibilities of technical failure inherent in so many R&D situations will often make it desirable to proceed to the third stage.

(3.) In Probabilistic Evaluation, one obtains subjective probabilities for the most sensitive variables by interviewing experts (managers, engineers, scientists, etc.) regarding the likelihood of their future levels of occurrence. The way these interviews are conducted is critical to obtaining a reliable view of the projects, so SDG has developed a special interview process. (See "Forecasting R&D Success" in this column for Jan. 9, 1980.) This step is the most controversial aspect of the decision analysis process, but especially in R&D it is often the most essential.

These judgments are encoded for technical as well as marketing and commercial factors. The probabilities are then laid out on a decision tree and displayed as a risk-return profile, or "profit lottery," an example of which is shown in Fig. 2. This is the probability distribution of the net present value, reflecting all the scenarios described in the decision tree. Menke likes the world "lottery" because it connotes uncertainty--the range of dif-

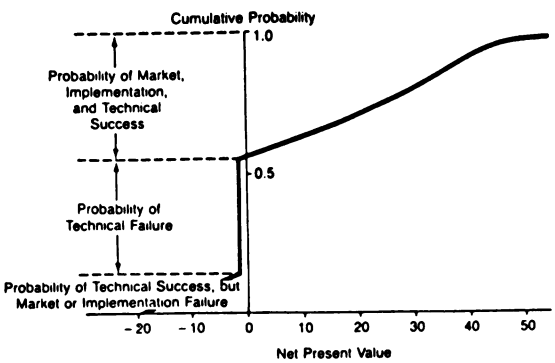

Fig. 2

The uncertainty associated with the value of a typical R&D project can be divided into three separate regions.

ferent possibilities, with their probabilities, that might be encountered if one went ahead with the project. Expected net present value, which SDG calls the Return of the project, is one of the key indicators of the project's worth. The curve, however, displays another important dimension--the full range of Risk, from the worst possible financial outcome to the best.

Menke observes that managers are usually not satisfied to work with a single expected or average value, particularly if the project commands a lot of the organization's resources. They want to adjust it for the degree of risk.

SDG defines the Risk of the project as one-half the variance V of the lottery. Given the Return which is the expected net present value, and the Risk, V/2, there is a simple and well established procedure which holds that if you measure the organization's ability to tolerate risk you can determine a value that is properly adjusted for the full degree of risk.

This so-called Certain Equivalent Value equals Return-Risk/Risk Tolerance, where Return and Risk are the average and the variance, respectively, of the lottery curve. From experience, Menke and his colleagues have found that Risk Tolerance for many organizations is about 1/6 of stockholder's equity.

Consequently, one can find an approximate Risk Tolerance in a company's annual report and plug it into the above equation to determine Certain Equivalent Value. In Menke's opinion, this is the best single estimate of the economic value of an R&D project to a particular organization.

(4.) Having found this value, Menke recommends that before making a final decision one ought to test the basis on which the evaluation rests. Once again, several techniques are employed to determine the economic value of obtaining more information, better alternatives, or a clearer statement of management's goals. Also at this stage the individual project should be reviewed for portfolio interactions with other projects or the existing business, which can modify their standalone value substantially. A future issue of "Managing Innovation" will explore SDG's approach to R&D portfolio management and resource allocation.

If the field of alternatives or the quality of information turns out to be insufficient, the Basis Appraisal will recommend not executing the decision immediately but, rather, obtaining more information or coming up with a better alternative. It is a self-correcting feature that only lets you proceed when the quality of the existing alternatives and information reduces the risk sufficiently to justify the project expenditures. On the other hand, the companies which have used decision analysis successfully have observed that it also reduces the common tendency to procrastinate in the face of uncertainty. In short, says Menke, it answers the final question: Should you act now or develop the decision basis further?

Readers desiring a detailed information packet on this decision analysis process should write to Michael Menke at Strategic Decisions Group, 3000 Sand Hill Road, Menlo Park, CA 94025.

* * *

18

SELECTING PROJECTS TO OBTAIN A BALANCED RESEARCH PORTFOLIO

Daniel L. Owen
Strategic Decisions Group

SELECTING PROJECTS TO OBTAIN A BALANCED RESEARCH PORTFOLIO

Daniel Owen

Strategic Decisions Group

ABSTRACT

This is a new methodology that can evaluate individual research projects and combine current projects and project opportunities into an attractive research portfolio. For individual research projects, it provides a framework for integrating and communicating information about technical and market uncertainties. The quantification of the characteristics of individual research projects assists in maintaining an appropriate balance between the riskiness of the research portfolio and its expected return. Empirical evidence suggests that many major laboratories may have a research portfolio that is too focused on process and project improvement. There are several explanations for this too conservative portfolio.

* This paper draws heavily on a presentation made to the Metallurgical Society of AIME, Atlanta, GA, March 1983.

SELECTING AND MANAGING YOUR RESEARCH PORTFOLIO

INTRODUCTION

Research and development (R&D) is characterized by uncertainty and long time horizons. There are two sources of uncertainty: the technical performance of the product and the market performance or profitability of the product if it is commercialized. Accounting for these two sources of uncertainty in the R&D decision-making process is made more difficult by the lack of communication between the research laboratory and the marketing organization. Researchers often complain, "They never tell us what they want, and they don't like what we do." In turn, the marketers lament, "They never ask us what we need, and we don't understand what they did."

Despite these uncertainties, Strategic Decisions Group (SDG) believes that decisions about R&D can and should be based on its ultimate contribution to corporate profitability. Our experience in applying decision analysis to important decisions involving uncertainty and distant payoffs, such as major capital investment decisions, strategic planning, and R&D management, convinces us that the contribution of research to profitability and the associated uncertainty in that contribution can be quantified.[1,2,3,4] The quantification of research alternatives assists the R&D project manager in recommending a research budget, in deciding which technology to pursue, and in determining whether to pursue a "backup" or competing technology. This quantification also assists R&D directors in recommending a laboratory budget, in selecting a strategic direction for the corporation's research, and in maintaining an appropriate balance between the riskiness of the research portfolio and its expected return.

The most important benefit of our approach to R&D management is not the recommendation of decisions, but rather the understanding that results from addressing a complex, uncertain problem with a logical, consistent methodology. In addition, the SDG approach provides a framework for communication among the parties involved in the decision, including the researchers, marketers, and project managers. It also allows laboratory directors to communicate clearly a carefully considered rationale behind a particular decision or recommendation to top management.

AN EXAMPLE DECISION ABOUT AN INDIVIDUAL RESEARCH PROJECT

A decision tree is an effective way of describing the sequence of decisions and uncertainties affecting an individual project. Figure 1 shows a decision tree (disguised for proprietary reasons) for a research project currently in progress at a major laboratory. The project may be described as research to develop an improved catalyst for use in automotive catalytic converters.

FIGURE 1

DECISION TREE PROVIDES A FRAMEWORK FOR INTEGRATING
TECHNICAL AND MARKET UNCERTAINTIES

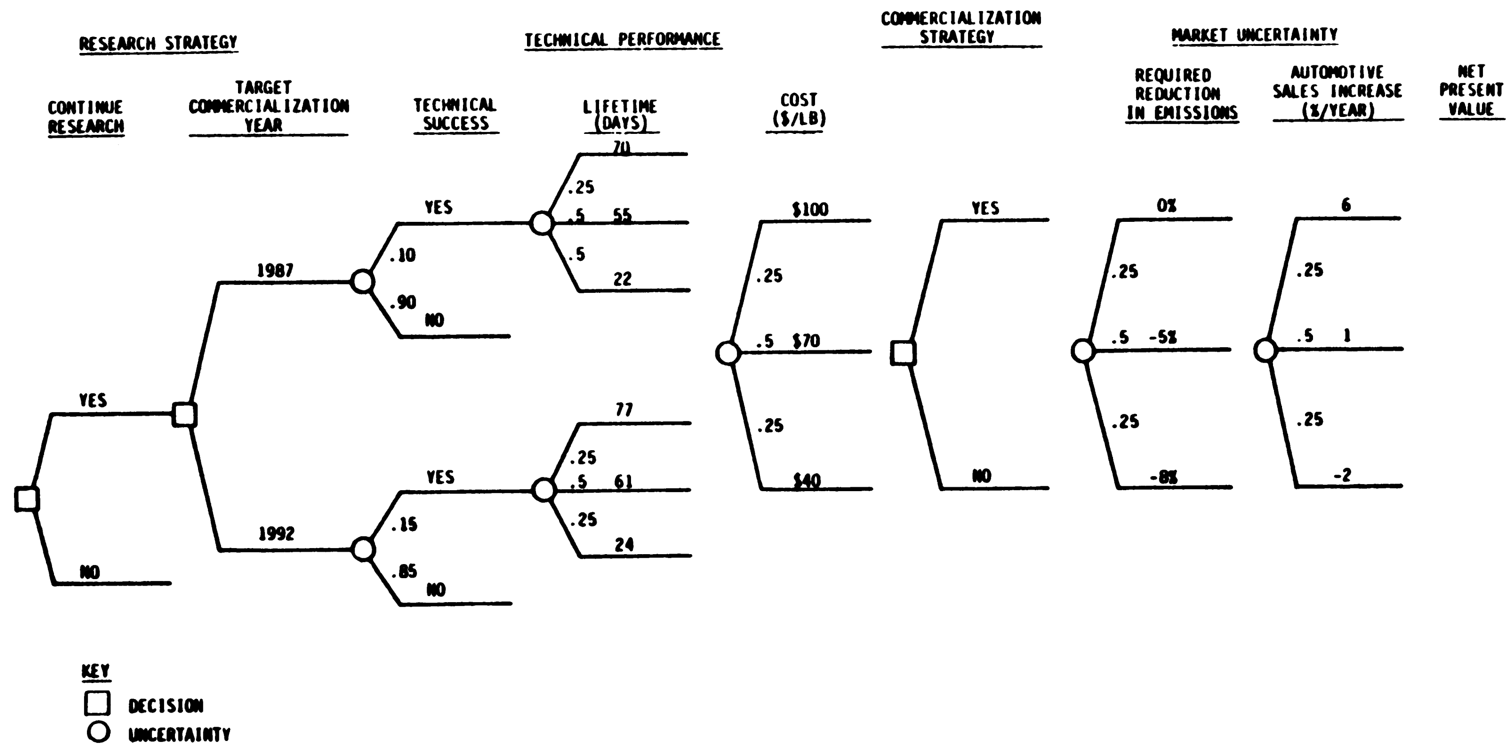

In Figure 1, squares represent decisions, which are under the control of the decision maker, and circles represent uncertainties, which cannot be controlled by the decision maker. First, the company must decide on a research strategy. Its options are to discontinue research or to continue research with two different target commercialization years. Commercialization in 1987 would require a parallel research program and a relatively larger budget. A 1992 target would allow time to work sequentially on problems. Different market conditions in each commercialization year would affect the value of the catalyst.

If the company decides to continue research, it must address important technical performance uncertainties surrounding technical performance. First, the research must establish whether or not the catalyst works; uncertainty about this is referred to as the "technical success" uncertainty. The probabilities of technical success have been assessed by the researchers as 10 and 15 percent, respectively, for 1987 and 1992. Researchers felt that a longer, more orderly program would allow more time for solving technical problems, thereby increasing the chance of technical success.

If the catalyst is technically successful, research must then establish the catalyst lifetime under certain extreme test conditions. The lifetime is currently uncertain and depends on the research strategy selected. The cost of the catalyst is also uncertain. The probabilities of catalyst costs of \$100, \$70, and \$40 per pound are 25, 50, and 25 percent, respectively. A node representing this uncertainty actually follows on each of the six branches of the lifetime uncertainty. For simplicity, only one node has been shown.

Once these technical performance uncertainties are resolved, the company can decide whether to commercialize the catalyst. This decision will depend on the technical performance outcomes. Commercialization will require an additional, substantial capital investment. If the company decides to commercialize the catalyst, the company must then address the important market uncertainties: the required reduction in emissions, and the number of cars sold, which determines the market size for the catalyst.

Several important observations about this decision tree can be made. First, it is an effective and concise way of communicating the decision alternatives, uncertainties, and values associated with a particular project. Second, it contains five uncertainties and is not too complicated. By using a careful sensitivity analysis, the five crucial uncertainties (from an initial list of about 25) that have the greatest influence on the value of the project were identified.

Finally, the information about the likelihood of various uncertainties and about the value of the project, given various events, is provided by experts within the client's organization or by its technical consultants. These probabilities are established in a formal interview process with individuals whose initials are usually indicated on the figure.[5] SDG's role is to assist the client by providing a logical framework for the client's information, values, and alternatives. Most often, however, the value of the project associated with a particular path through the decision tree is calculated from a business model partially or completely constructed by SDG.

Because the probabilities for each path through the tree have been specified and a value for each path can be computed, the expected value and uncertainty associated with each research strategy can be determined. For example, continuing research with a target commercialization of 1987 gives an expected net present value (NPV) of $35 million. There is a 10-percent chance of a value less than -$5 million or of a value exceeding $90 million. The 1992 commercialization has a higher expected value of $45 million, with a range of -$5 million to $105 million.

This result is somewhat unusual and counterintuitive. Usually, research projects have a market "window," during which the results of the project have maximum market value. If the results are not obtained by a particular time, the competitive advantage is lost, and the market value greatly diminished. For this project, the opposite is true: delaying the research results increases the project's value. The increased chance of technical success, longer catalyst lifetime, and improved market conditions make the 1992 commercialization a better research strategy.

This observation had important implications for the management of the project. For example, parallel activities on material development, fabrication, and manufacturing were being pursued in an effort to commercialize as quickly as possible. Because early commercialization was unnecessary, we recommended a reduced-cost, sequential research effort.

Another very useful result of this decision tree is the recommendation of the commercialization strategy given the technical performance results (see Figure 2). With the early commercialization research strategy, the catalyst should be commercialized unless the catalyst lifetime has a low value of 22 days. With the late commercialization research strategy, the catalyst should be commercialized unless the lifetime has a low value of 24 days and the cost takes on a high value of $100 per pound. The probability that commercialization will be desirable with the 1992 research strategy is 94 percent compared with 75 percent for the 1987 research strategy, because increased lifetime is likely to result from a longer research program.

As a result of the decision analysis, management changed its perception of the project's value. The main value of this research project derived from its "insurance" value to the company. Referring again to Figure 1, the company's current technology could be extended to meet up to a 5-percent required reduction in emissions. In that situation, the value of the research on the improved catalyst is small. However, a required reduction of 8 percent could not be met by an extension of existing technology, which gives the research results a very large value. This research project insures the company against the 25 percent chance that the current technology cannot be extended to meet future requirements.

In the simplest form, R&D projects can be characterized by four fundamental elements: R&D cost, uncertainty about technical success, uncertainty about commercialization cost, and uncertainty about the market value of the commercialized product or process (see Figure 3). The sequence in which decisions are made and uncertainties resolved is important and is as follows: the decision to pursue the research, resolution of the technical uncertainties, the decision to commercialize, and resolution of the market uncertainties.

FIGURE 2

THE LATER COMMERICALIZATION RESEARCH STRATEGY LEADS TO A GREATER PROBABILITY (94% VS 75%) OF COMMERCIALIZATION, GIVEN TECHNICAL SUCCESS, BECAUSE INCREASED LIFETIME IS LIKELY TO RESULT FROM A LONGER RESEARCH PROGRAM.

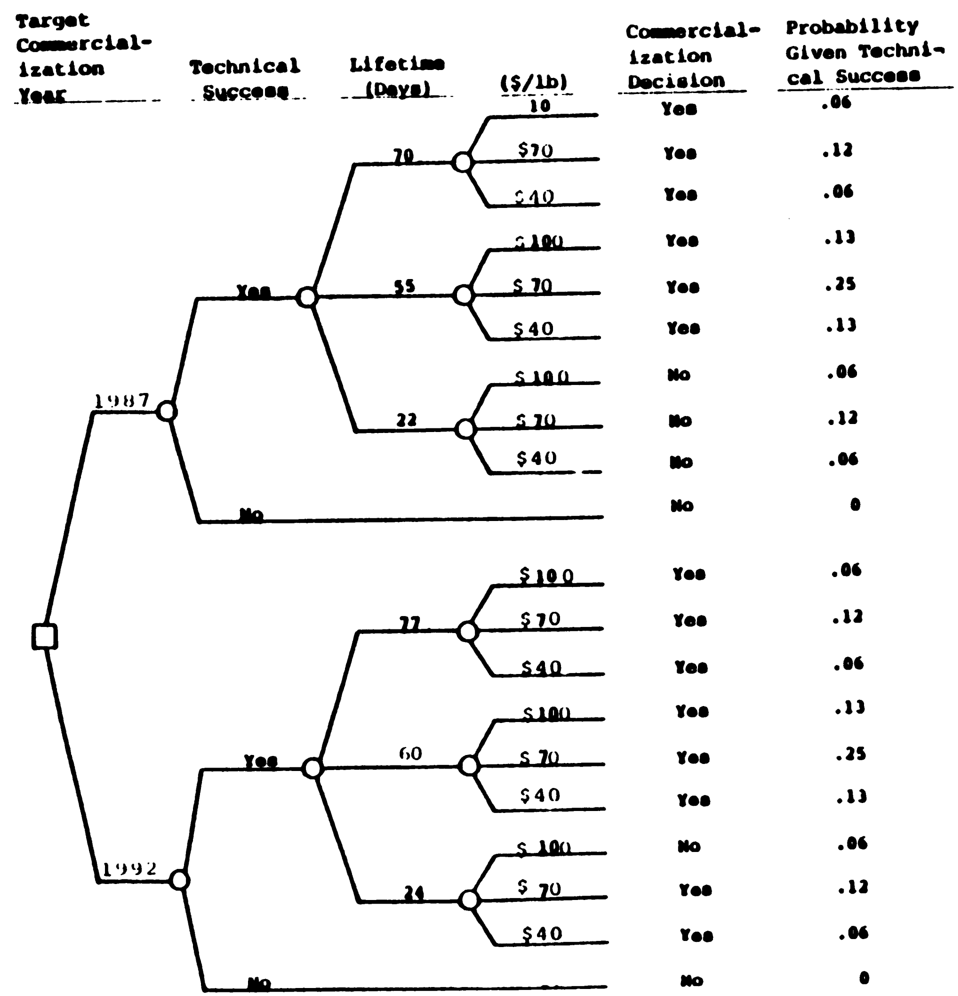

FIGURE 3

THE FUNDAMENTAL ELEMENTS OF INDIVIDUAL RESEARCH PROJECTS:
R & D COST, UNCERTAINTY ABOUT TECHNICAL SUCCESS,
COMMERCIALIZATION COST, AND UNCERTAINTY ABOUT MARKET VALUE.

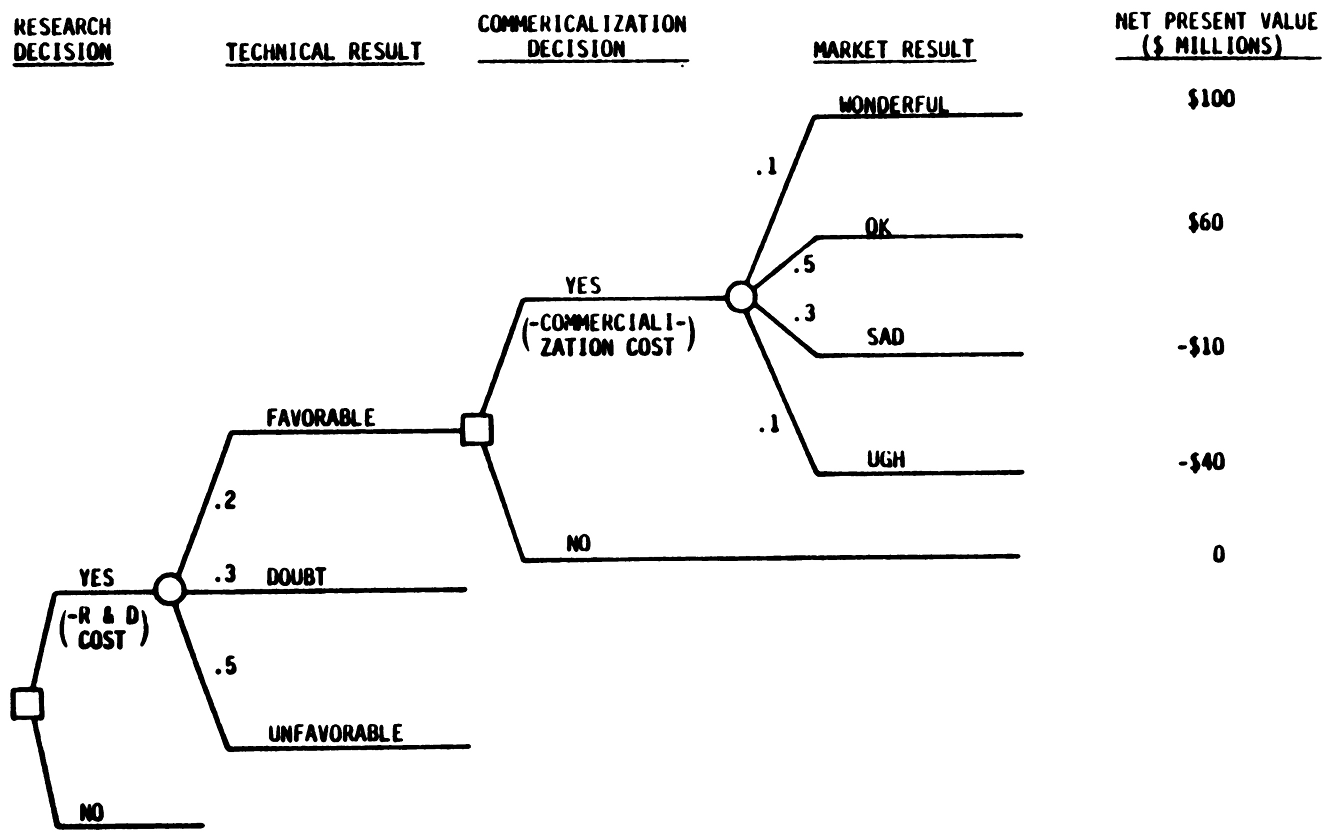

FUNDAMENTAL RESEARCH DECISIONS

Having demonstrated the application of the decision analysis methodology to a complex research problem in the previous section, this and several following sections provide a conceptual discussion of fundamental research decisions. Because the focus is on concepts, simplified examples are used.

The Decision To Begin A Research Program

Figure 4 shows a simplified decision tree for deciding whether to begin a $10 million research project on Technology A. If the program is pursued, the research results obtained will be favorable or unfavorable. The figure indicates that technical experts researching Technology A agree that the probability of favorable research results is 20 percent. This probability is assessed in a formal interview designed to minimize biases in the experts' judgments.[5]

Technical success is usually defined as the resolution of several precisely defined technical "showstoppers." The probability of technical success is the product of the probabilities of overcoming each showstopper. The probability of successfully solving each showstopper is assessed and then they are analytically combined to find the overall probability of success. Often, there is probabilistic dependency between the showstoppers, which must be included.

If the research results are favorable, there is a 95-percent chance that the technology will be successfully commercialized. An unfavorable research result reduces the chance of successful commercialization to 40 percent. In the latter case, existing technology would have to be employed in the commercialization; however, some benefit from the research program is anticipated even if the results are unfavorable. Of course, existing technology could also be employed if Technology A is not researched. The chance of successful commercialization is only 30 percent, however, because there are no benefits from the research program.

The expected value of each decision alternative can be computed by multiplying all the path probabilities following the decision by the values at the end of the paths (see Figure 5). For example, if the research result is unfavorable, the expected value of the decision to commercialize is $38 million (.4 x 110 + .6 x -10), compared with a value of zero from not commercializing. Favorable research results give an even higher expected value for commercialization. On the basis of expected values, the decision recommendation would be to commercialize whether or not the research results are favorable. This result could have important tactical implications for the marketing organization. Similarly, the expected value of the decision to research Technology A is $41 million, compared with only $14 million without the research. Therefore, the research should be undertaken.

Not all decision makers, however, want to make decisions on the basis of expected values. Because they are "risk averse," they view the likelihood of successful commercialization as more uncertain and believe

FIGURE 4

FUNDAMENTAL ELEMENTS CAN BE QUANTIFIED

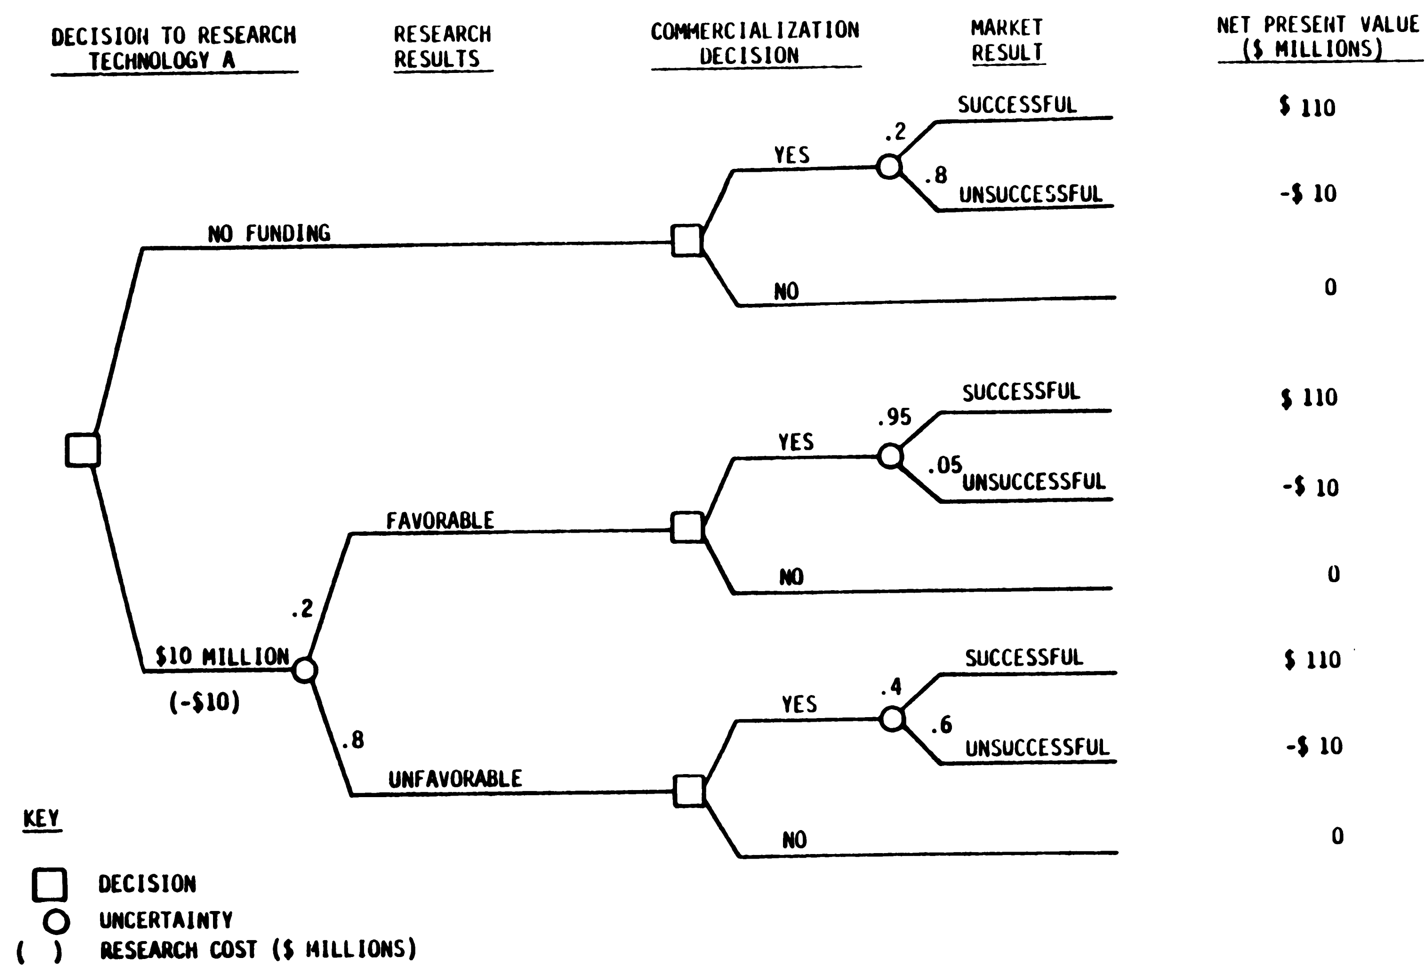

FIGURE 5

THE EXPECTED VALUE OF RESEARCHING TECHNOLOGY A SUGGESTS
THAT THE RESEARCH SHOULD BE UNDERTAKEN

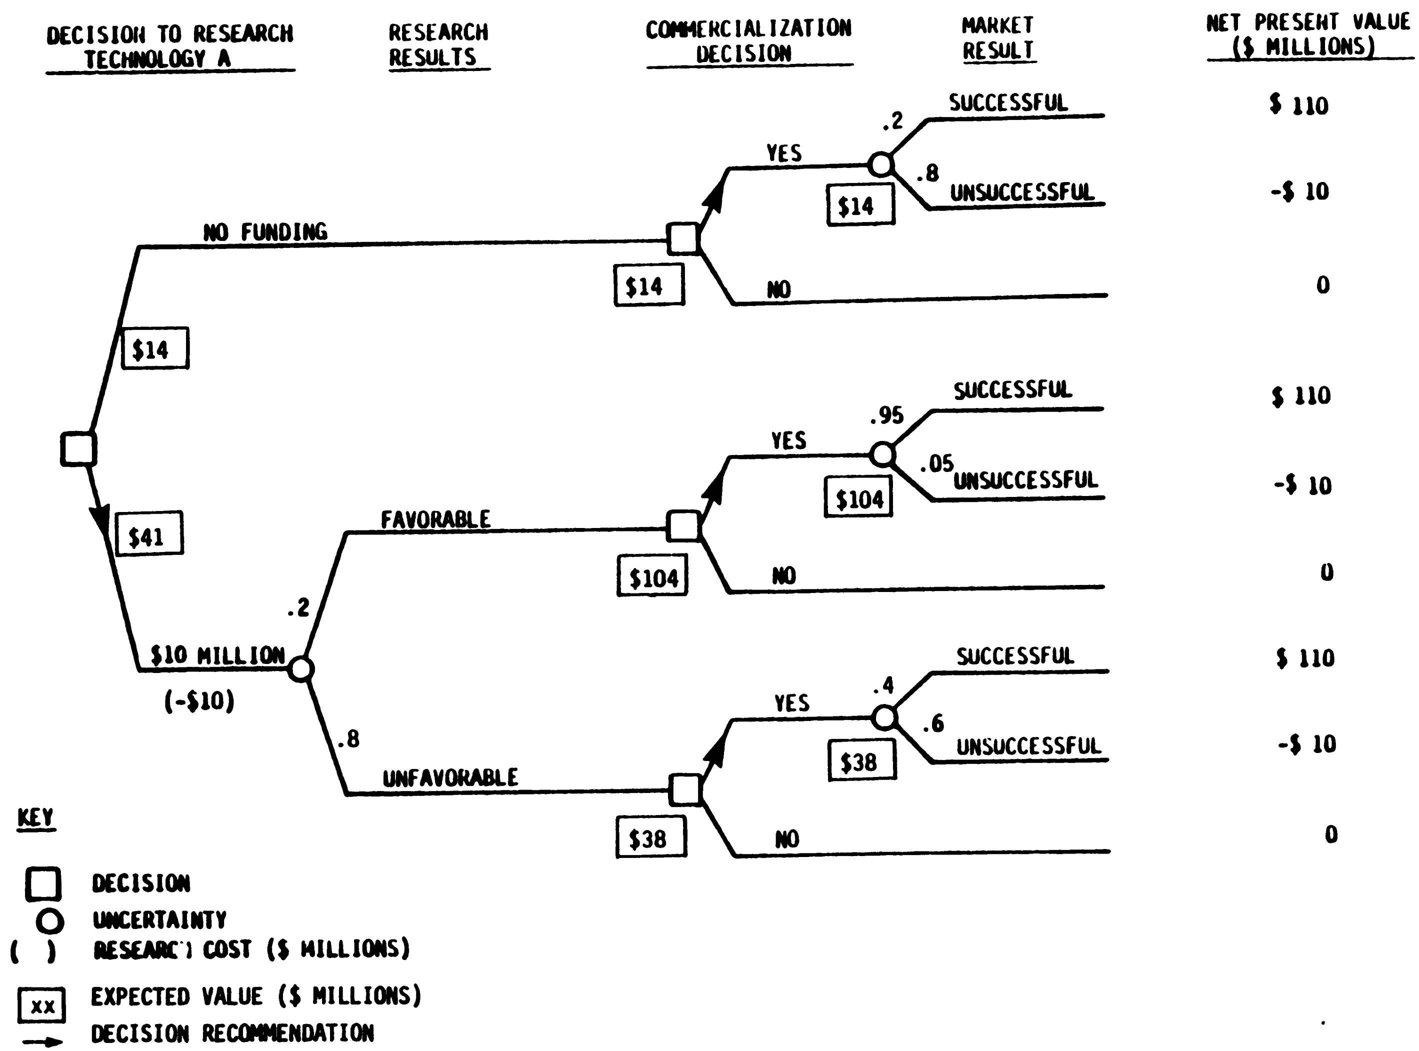

that the value of the commercialization option is somewhat less than its expected value. Incorporating a decision maker's (or corporation's) attitude toward risk into the decision recommendation is an important element of decision analysis.[6,7] However, because our current focus is R&D management, we have ignored risk aversion in this example.

The Decision About the Level of Funding

While researching Technology A is preferable, perhaps the funding level of the program should be increased to improve the likelihood of success and to achieve earlier results. The decision tree in Figure 6 shows two alternative levels of funding for the research program: $20 million and $10 million. The technical experts believe that the chance of favorable research results increases to 25 percent at the $20-million funding level. Marketing believes that the earlier commercialization resulting from the higher funding level would allow the company to gain a substantial edge over the competitors. Consequently, the NPV of the commercialized technology would increase by $10 million to $120 million if the research result is favorable. If the result is unfavorable, the NPV remains $110 million.

Using the same calculation as before, the expected NPV with the higher funding level is less than that with the lower level (see Figure 7). The additional funds do not sufficiently increase the chance of a favorable research result or sufficiently increase the market value to justify their expenditure.

Decisions About Research on Parallel or Competing Technologies

The problem that many R&D managers cite as the most perplexing is deciding whether to research several technologies to achieve a single research goal. These managers realize that the decision depends on the importance of the goal, the likelihood of success, and the incremental cost. However, balancing all these factors is difficult. One vice president for R&D described this problem as "too difficult to solve in my head and too important to solve in my gut."

A decision tree is a useful tool for thinking through the problem of "back-up" technologies. For example, in addition to Technology A, it might be possible to research an alternative technology, Technology B, that is identical to A in its marketability (see Figure 8). Figure 8 also shows a probabilistic dependency between Technology A and Technology B. If Technology A succeeds, then B is more likely to succeed, but its success is not guaranteed. Computing the expected NPV shows that pursuing Technology B in addition to A increases the overall chance of success sufficiently to justify pursuit of both technologies. In fact, pursuing both increases the expected NPV by $5 million. From the decision tree, we can also compute the probabilities of some uncertain events. For example, the chance of technical success increases from 20 to 32 percent with both technologies, and the likelihood of a successful market result increases from 51 to 58 percent.

It is easy to extend this decision tree to handle the more general case where additional value results of both technologies are successful because of some market synergy.

FIGURE 6

INCREASING THE LEVEL OF FUNDING OF A RESEARCH PROJECT MAY INCREASE THE CHANCE OF TECHNICAL SUCCESS AND THE VALUE

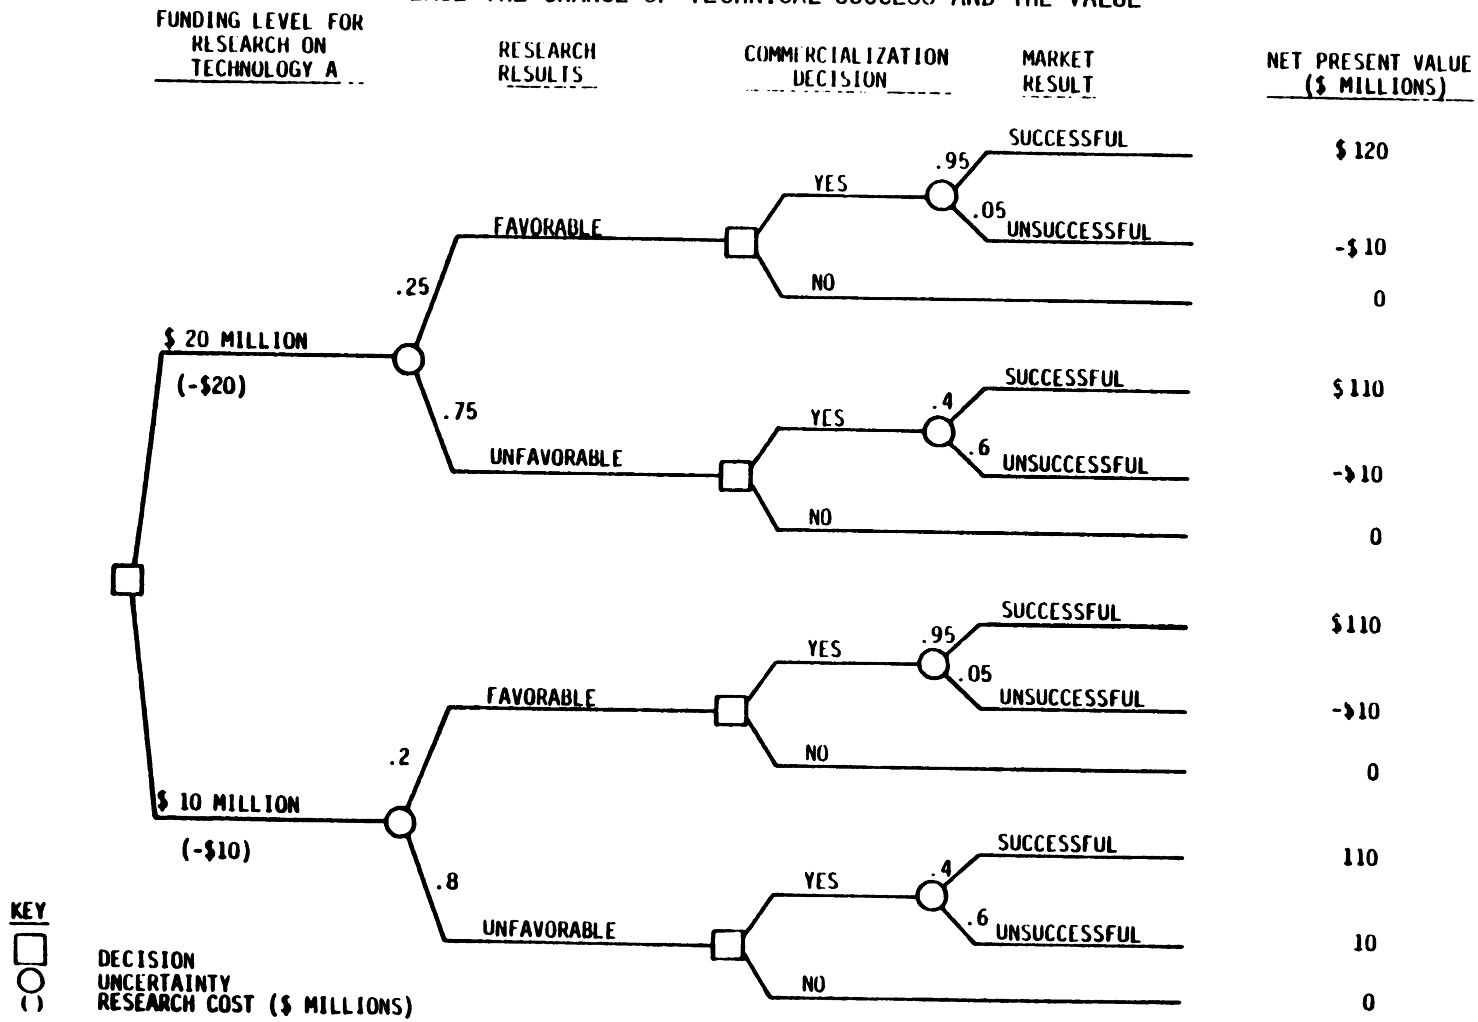

FIGURE 7

ADDITIONAL FUNDING FOR RESEARCH ON TECHNOLOGY A NOT JUSTIFIED BY THE INCREASED CHANCE OF FAVORABLE RESEARCH RESULT OR INCREASED VALUE

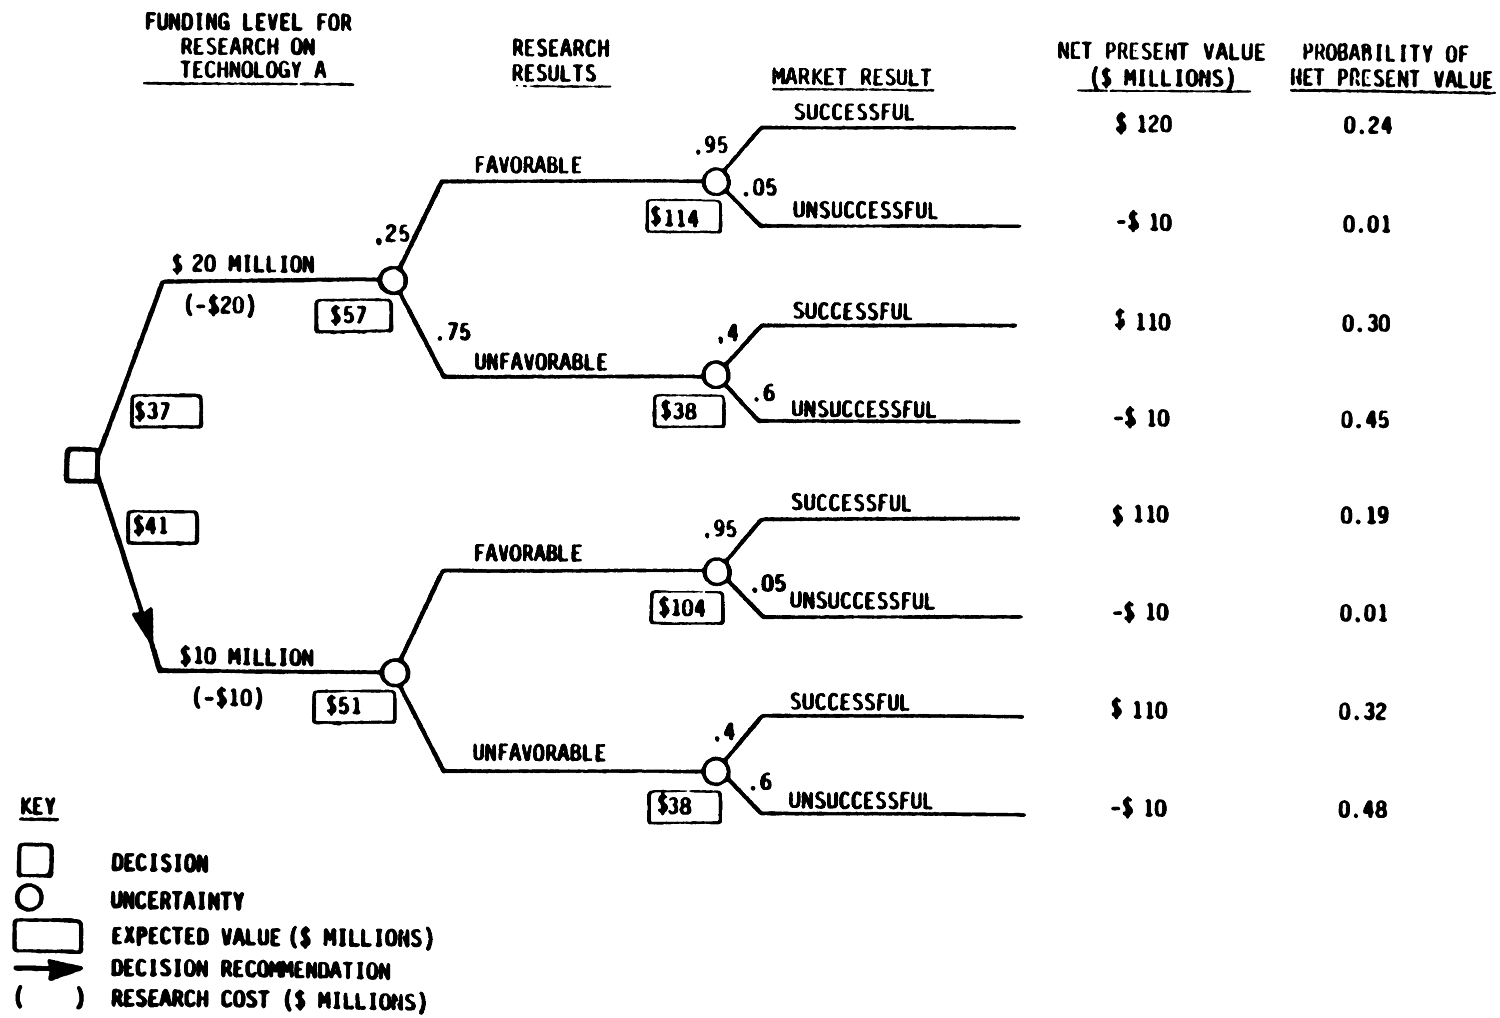

FIGURE 8

PURSUIT OF BOTH TECHNOLOGIES IS JUSTIFIED BY INCREASE IN EXPECTED VALUE

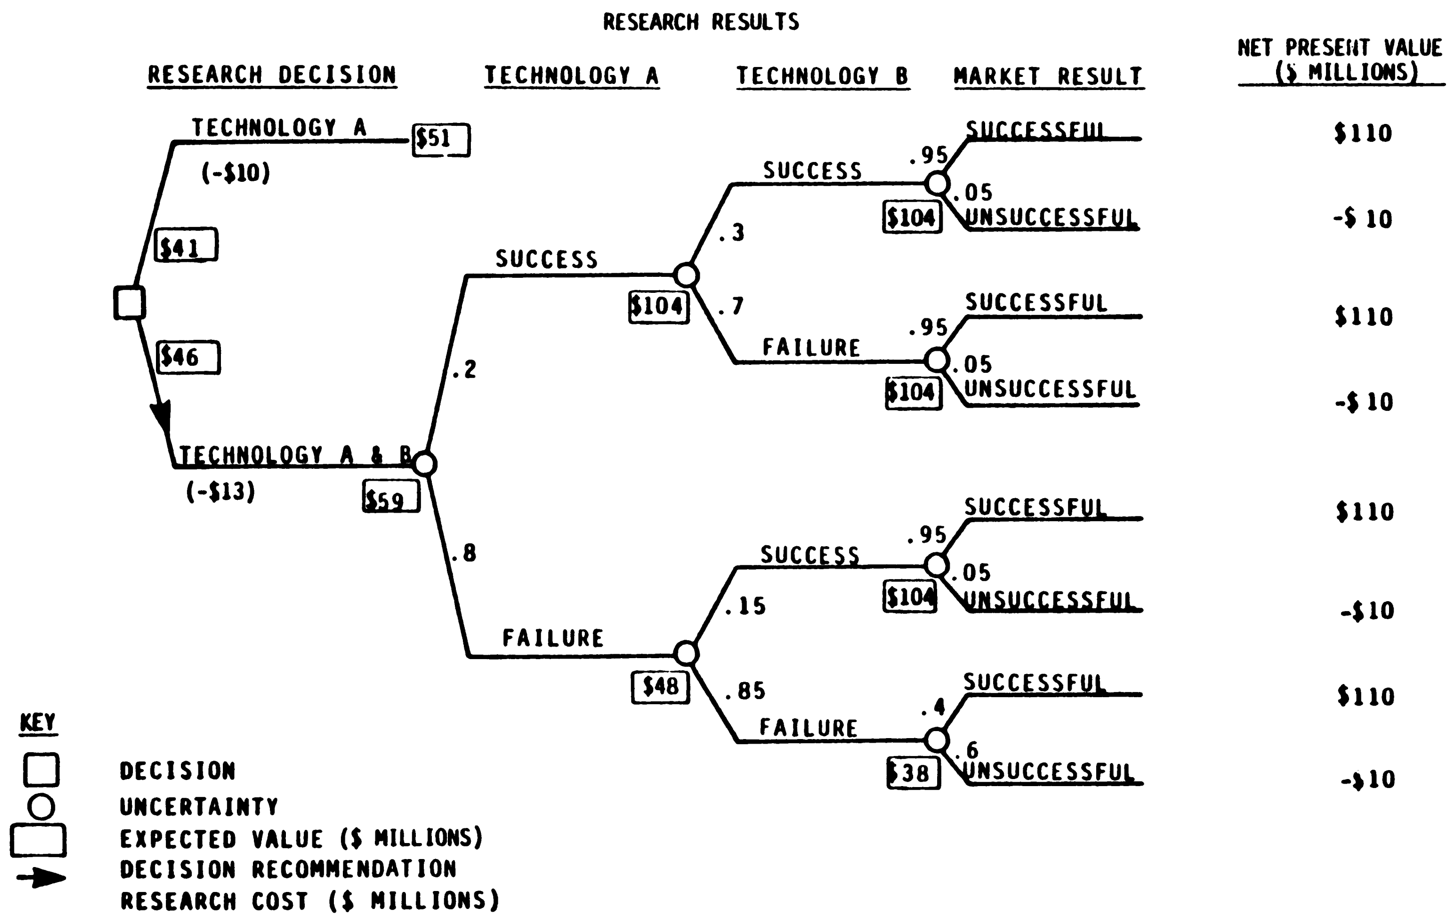

SELECTING A BALANCED RESEARCH PORTFOLIO

The research portfolio grid of Figure 9 helps to visualize qualitatively several aspects of the problem of selecting a research portfolio. To justify the budget, the laboratory director is tempted to select projects that have a high probability of technical success and, consequently, have a relatively low value. These are called "bread-and-butter" projects for the laboratory. They are often process and product improvements easily within the laboratory's capabilities, which will certainly not substantially alter the company's long-term profitability.

Projects with a high value (given technical success) and a high probability of technical success (silver bullets) are easily distinguishable even without quantitative methods. Unfortunately, such projects are very rare. More often, the laboratory director selects some "question-mark" projects, which are projects not likely to succeed but very valuable if they do. This high value frequently occurs because the project allows the company to introduce a revolutionary product or to take exclusive advantage of changes in the marketplace resulting from government regulation, complementary technical innovation, competitive actions, and so forth. Such projects usually demand state-of-the-art technology and are at the limit of the laboratory's capabilities.

For the laboratory director, the research portfolio problem can be described as selecting a balance between "bread-and-butter" and "question-mark" projects. Unless the probability of technical success and commercial value are carefully quantified, however, the laboratory director is likely to end up with a "bread-and-butter" and "turkey" sandwich.

SDG maintains that portfolio balancing can be best accomplished by placing quantitative descriptions of each project on the research grid. For example, the fundamental elements of the project described in Figure 3 can all be displayed on the research grid in Figure 10. A dot is positioned to show the probability of technical success and, given technical success, the expected value. The horizontal bar shows the range of uncertainty on the commercial value, given success, and the number indicates its budget requirement.

This research grid can be used to combine current and prospective research projects into a portfolio that enhances the value of the laboratory's research. Hyperbolas of constant expected net present value are shown on the research grid (see Figure 11). For example, projects E and G both have an expected value of about $10 million, though their values given technical success are very different. Projects that have the greatest expected value should be selected first (again, ignoring risk attitude for simplicity). Consequently, the most valuable research portfolio, consistent with a $40 million dollar budget constraint, will result from selecting projects A,B,C, and D. However, complications such as probabilistic dependency and market synergy require special treatment. From Figure 11, one can calculate that the total expected value of portfolio ABCD is approximately $275 million.

FIGURE 9

JUSTIFYING THE LABORATORY'S BUDGET AND
CAPITALIZING ON ITS SCIENTIFIC ABILITIES
OFTEN CONFLICT

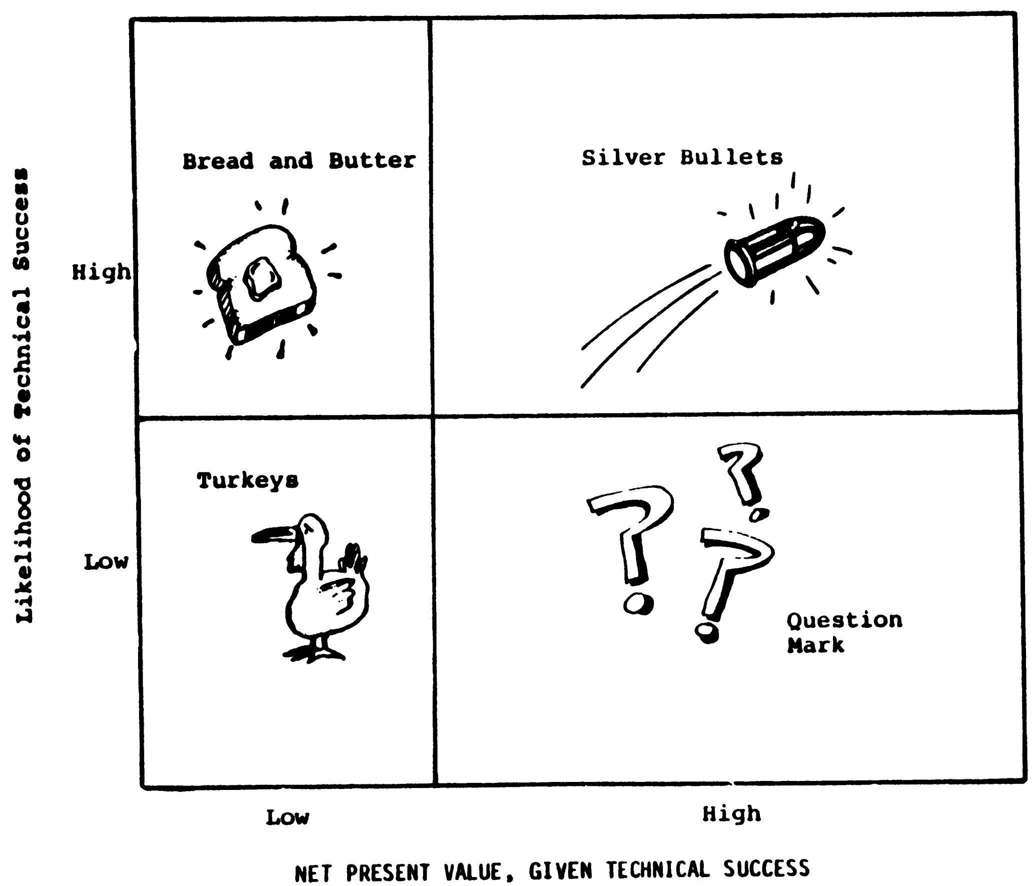

FIGURE 10

THIS RESEARCH GRID
CAN DISPLAY THE IMPORTANT PARAMETERS CHARACTERIZING A RESEARCH PROJECT

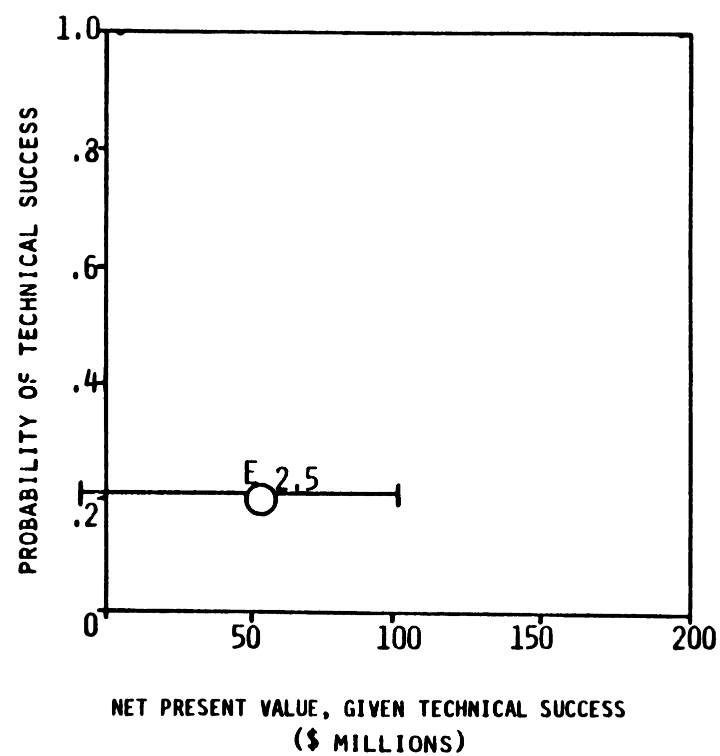

FIGURE 11

THE RESEARCH GRID MAY BE USED TO EVALUATE THE LABORATORIES' PORTFOLIO OF RESEARCH PROJECTS

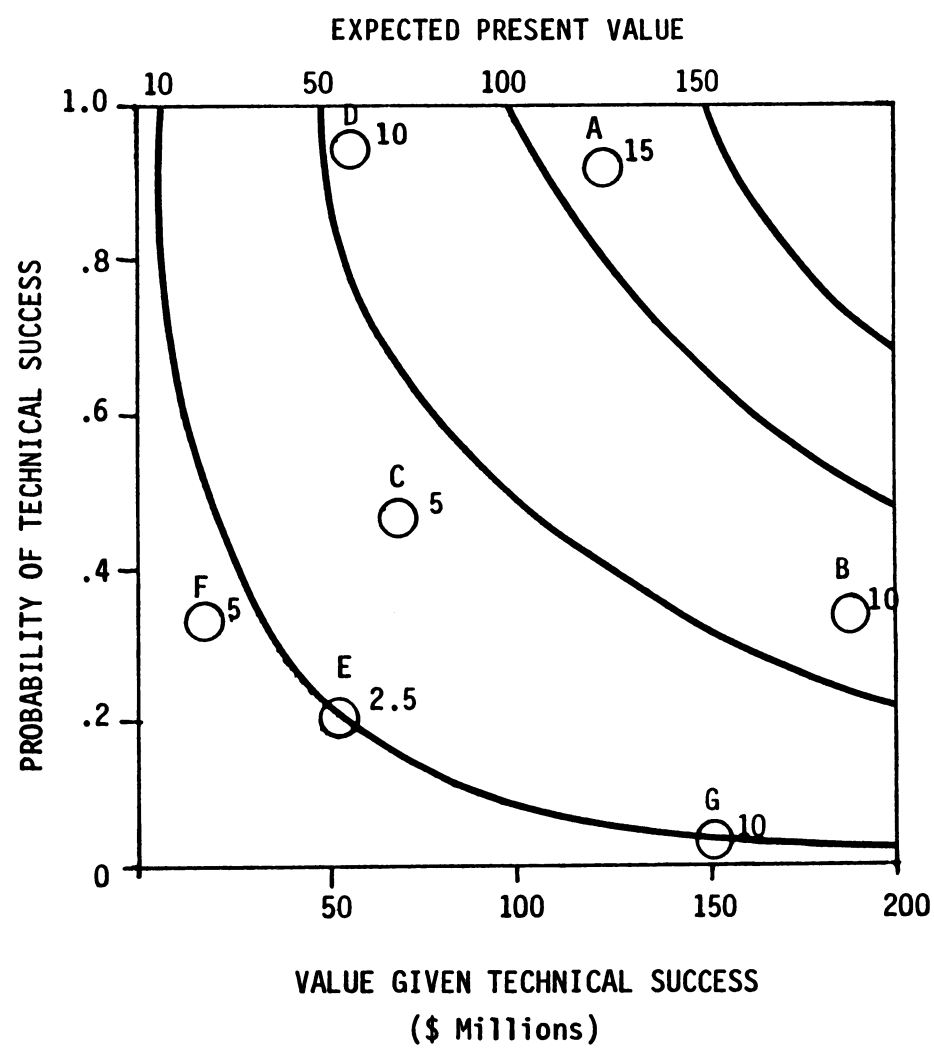

Project	Total Cost	Approximate Expected Value ($ Millions)
A	15	110
B	10	70
C	5	40
D	10	55
E	2.5	10
F	5	5
G	10	10

Another important feature of the grid is that the progress of individual projects in the portfolio can be tracked. Horizontal movement of project points over a period of time suggests that as the time for commercialization approaches, the perception of market opportunity has changed. A substantial drop in value could signal the need for an adjustment in the project budget. Project points should move vertically as time passes or disappear completely if research indicates that the technical problems cannot be overcome. When a project manager realizes that continued support of the project depends on increases in the probability of technical success, he efficiently allocates funds to those technical hurdles with relatively low probability. This incentive reduces the tendency of some project managers to work simultaneously on all aspects of the problem even when there is no time constraint.

Figure 12 shows the disguised research grid of the research portfolio for a major company. Early-stage projects have been in the laboratory for only a few years, and late-stage projects have been in the laboratory for many years. These empirical results are particularly interesting, because the evaluation of the projects was performed by trained professionals in decision analysis using a consistent methodology to evaluate all of the major projects of this laboratory with an annual research budget exceeding $50 million.

Probably the most interesting observation is that for both early- and late-stage projects, the probability of success decreases (almost linearly) as the expected value increases. Having made this observation, it is easy to suggest an explanation. Bread-and-butter research on process and product improvements does not require major advances in science or technology, and, consequently, is of relatively low value. Question-mark projects, on the other hand, require an innovation or scientific advance which leads to a low probability of success. However, if the project is successful, that innovation or advance will provide the company with a new and very valuable position in the marketplace. Often, successful question-mark projects lead to new business areas with high margins and few competitors.

Based on my experience with several major laboratories, most laboratories seem to have too many bread-and-butter projects and too few question-mark projects. There are several reasons why this observation is correct. First, in many instances, the laboratory's budget is based, at least in part, on the number of technical successes it can claim without regard for their value. Typically the laboratory director goes "downtown" at budget time armed with a list of projects that were successfully completed in the previous year. This system gives the laboratory director an incentive to pick projects that are likely to succeed, as opposed to those that are likely to be valuable. Budget allocation processes in which the operating divisions have a strong influence on the laboratory budget are also susceptible to a bread-and-butter focus.

A second reason laboratories may have too many bread-and-butter projects is that laboratory directors do not appreciate how many question-mark projects are required to ensure a single success. For example, suppose a laboratory has ten independent projects, five bread-and-butter projects with a success probability of .8 and five

FIGURE 12

PERIODIC REVIEW AND DISPLAY OF THE RESEARCH PORTFOLIO HELPS TO MAINTAIN THE BALANCE BETWEEN BREAD-AND-BUTTER AND QUESTION-MARK PROJECTS

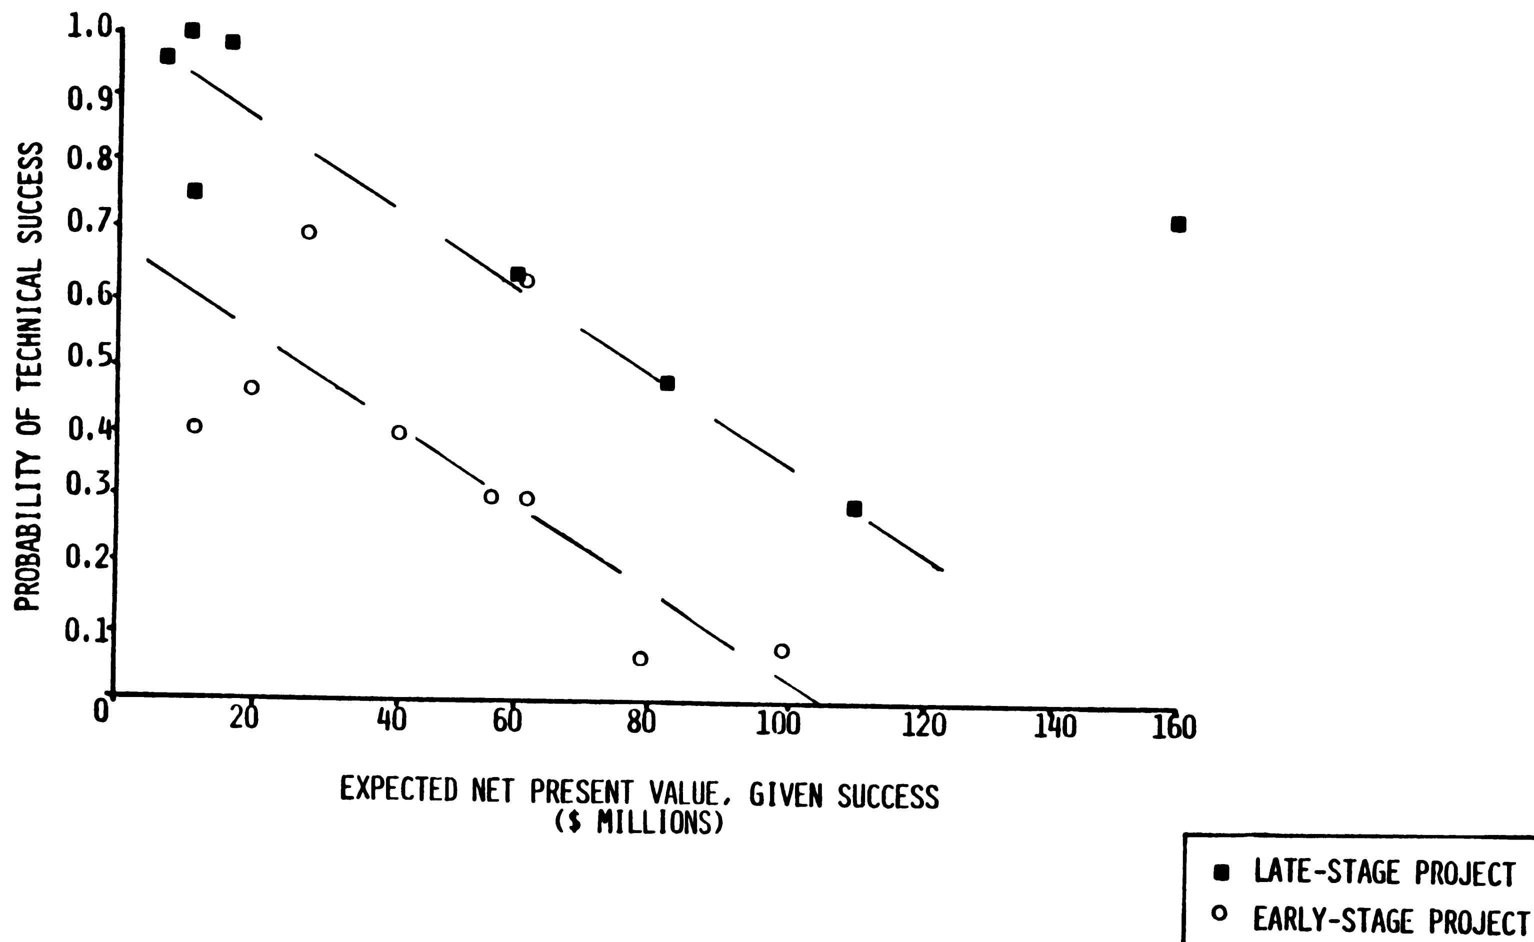

question-mark projects with a success probability of .2. The expected number of bread-and-butter successes is four, while the expected number of question-mark successes is one. While, the probability of no bread-and-butter successes is .03%, the probability of no question-mark successes is 33%. It takes a lot of question-mark projects to ensure that one is successful.

Finally, a laboratory may have too many bread-and-butter projects because the question-mark projects often have payoffs in the more distant future. The innovation, characteristic of question-marks projects, usually takes a long time to achieve and often requires a longer period for development and commercialization. If the company has too strong a focus on near-term results, question-mark projects may not be selected. The overly strong focus on the near-term can be a result of using an improper discount rate. For example, some companies increase the discount rate to account for risk, i.e. they use a risky discount rate. This practice is erroneous, because it implies that risk increases geometrically over time, which is not the case in general and certainly not the case for research. Of course, it also ignores the fact that different projects have different degrees of uncertainty. A proper method for treating risk is to quantify explicitly the uncertainty (as we have advocated in this paper), so that time preference and risk preference can be treated separately, the former with discounting and the latter with utility theory. Two major laboratories have asserted that removing the unjustified penalty on long-term research imposed by a too-high discount rate is a very important contribution of decision analysis.

Some people argue that it is risky to fund large numbers of low probability projects, because the probability of success is too low for each project. However, analyzing the portfolio shown in Figure 12 indicates that probability alone is not a good criterion for rejecting projects. Forty percent of the total expected value of the portfolio is contributed by projects with less than a 50-percent chance of success. While not funding question-mark projects may result in a less "risky" portfolio from the corporate perspective, the result is much less desirable. Funding only bread-and-butter projects will undoubtedly lead to a deterioration of the company's long term strategic position in the marketplace. Continued funding of research on the same processes and products leads to a reduction in the value of the research dollar as improvements become harder and harder to make, and the company's position becomes increasingly vulnerable to competitors that have a balanced research portfolio.

CONCLUSION

Technical and market uncertainties make it very difficult to establish the value of a research project to a company. Decision analysis is a proven methodology for addressing complexity and uncertainty in the management of R&D. Its most important benefit is in providing a framework for communication among the research organization, the marketing organization, business planners, and the corporate management.

Examples presented in this paper show the versatility of the methodology. One example shows how the value and uncertainty for a project were established and reveals the resulting strategic and tactical management recommendations. Decisions of whether to fund a project, of what should be the appropriate funding level, and whether to fund competing or "backup" technologies can also be addressed with decision analysis.

The quantification of the characteristics of individual research projects assists in maintaining an appropriate balance between the riskiness of the research portfolio and its expected return. Empirical evidence suggests that many major laboratories may have a research portfolio that is too focused on process and project improvement.

REFERENCES

1. R. A. Howard et al., "The Decision to Seed Hurricanes," Science, 176 (June 16, 1972): 1191-1202.

2. W. B. Gibson and M. M. Menke, "Meeting the Challenge of an Age of Uncertainty," Handling and Shipping, Presidential Issue, 1973.

3. J. E. Matheson, "Decision Analysis Practice: Examples and Insights," Proceedings of Fifth International Conference on Operations Research, Venice 1969, (John Lawrence, ed.) London: Tauistocle Publications, 1970, pp. 677-691.

4. M. M. Menke et al., "Evaluating Basic Research Strategies," Long Range Planning, 14, no. 3 (June 1981): 44-57.

5. C. S. Spetzler and C. S. Stael von Holstein, "Probability Encoding in Decision Analysis," Readings in Decision Analysis, ed. by R. A. Howard, J. E. Matheson, and K. L. Miller. (Menlo Park: SRI International, 1977):403-427.

6. R. A. Howard, "The Foundation of Decision Analysis," IEEE Transactions on Systems Science and Cybernetics, SSC-4, no. 3 (September 1968): 211-219.

7. C. S. Spetzler, "The Development of a Corporate Risk Policy for Capital Investment Decisions," IEEE Transactions on Systems Science and Cybernetics, SSC-4, no. 3 (September 1968): 279-300.

19

CALLING THE SHOTS IN R&D

Hans Ulrich Balthasar
Cilag, AG, Switzerland

Roberto A. A. Boschi
CIBA-GEIGY

Michael M. Menke
Strategic Decisions Group

Calling the shots in R&D

Subjective probability allows a reliable estimate of potential success five to ten years ahead

Hans Ulrich Balthasar, Roberto A. A. Boschi, and Michael M. Menke

On the basis of almost ten years' experience at Sandoz, a Swiss pharmaceutical company, these authors describe a method of evaluating the technical success of individual product development projects, basic research, and the R&D process as a whole. The forecasts come from a panel consisting of R&D line managers and specialists familiar with all aspects required for a project's success. The probabilities can be used to guide the pace and direction of R&D as well as to test the feasibility of long-range corporate objectives such as sales targets in long-term plans.

Dr. Balthasar is marketing manager for Diagnostica Products at F. Hoffmann-La Roche & Co., Ltd., Basel, Switzerland. He was formerly head of R&D planning at Sandoz Ltd., Basel. Mr. Boschi is in the R&D planning group at Sandoz Ltd., Basel. Mr. Menke manages the decision analysis activities of SRI International in Europe and has applied subjective probability measures to R&D problems in the pharmaceutical, office products, electronics, and nuclear industries.

Although all management functions must cope with uncertainty, R&D is generally agreed to be the function involving the largest number and widest range of uncertainties. Thus the R&D manager faces huge problems not only in selecting the most promising avenues for R&D effort and expenditure but also in attempting to ensure a steady flow of technically successful projects.

A flow of new products, while difficult to achieve, can pay off handsomely by effectively using expensive fixed developmental resources (e.g., safety testing, development engineering, and pilot plant facilities) and available marketing capacity (e.g., promotional department, medical representatives, and sales forces). In addition, a regular flow of new products can be important psychologically to provide the steady growth of revenues and profits by which outsiders are most likely to judge the quality of management.

The lengthening time scale of R&D work exacerbates the uncertainty and complexity of this management effort. Today, frequently five to ten years or more may be required to achieve technical success and an additional five to eight years may be necessary to realize full commercial potential in the pharmaceutical industry. Nevertheless, despite substantial uncertainties, it is possible to measure future success effectively. Here, we report how Sandoz, a large R&D-based Swiss pharmaceutical company, developed and tested such methods.

We describe forecasts made over the past seven years, compare them with the actual results, and conclude that subjective probability judgments are a reliable guide to future events and thus that management should use them as the basis for R&D planning and decisions.

Background of the approach

Given the uncertainties and obvious cost pressures in pharmaceutical R&D, Sandoz realized more than ten years ago that there was a great need to incorporate these uncertainties explicitly in the planning and decision-making process. After a preliminary review of available theoretical methods in the period 1967 to 1969, the company decided to assess subjective probabilities of technical success as a way of "quantifying the unquantifiable" in R&D. The use of probability allowed Sandoz to quantify the uncertainty in project duration and cost and the safety, efficacy, and market prospect of new products. It also made it possible to allow in a systematic way for the unforeseen developments that typify R&D work.

A major benefit of the explicit quantitative approach described here is that it synthesizes the opinions of a diverse group of experienced experts more effectively than alternative qualitative approaches. Another advantage is that many models and techniques to aid management in planning and controlling individual R&D projects require probabilities of success as one of the major inputs.[1]

A further advantage is that one can make a probability distribution on the output of the entire R&D process to depict the success of the R&D function in terms of the number of new products expected to be ready for market in each of the following years, providing a current indicator for intermediate to long-term results. Similar distributions can be developed for R&D expenditure, future sales of these compounds, and other parameters of interest to R&D and corporate management.

R&D environment at Sandoz

In 1976 Sandoz had sales of over 4 billion Swiss francs (about U.S. $2 billion) of which approximately 50% came from pharmaceutical products. The other 50% was divided among dyestuffs, agricultural chemicals, and food products. The R&D expenditure in 1976 was about 9% of sales. Pharmaceutical research at Sandoz has traditionally been oriented toward the search for new chemical entities that will be unique, effective, patentable, and safe. A general measure for the overall success of the R&D effort is a regular flow of marketable new products.

In the drug development process, the initial stage, sometimes called exploratory or preproject research, is the initial synthesis and screening procedure designed to identify therapeutically useful compounds. Although the activities necessary to promote a compound to the next stage can be accomplished in as little as six months, in practice as many as a thousand compounds may need to be prepared in order to find one preclinical candidate.

The next two stages, often called development, begin with detailed preclinical studies of the selected new chemical entities. These compounds are designated as projects.

The purpose of the preclinical stage is to determine safety and efficacy before a test of the new compound in man. Once the safety is established, detailed clinical investigations begin. The various clinical phases are designed to statistically establish tolerability, efficacy, and, it is hoped, superiority over current therapy. Long-range toxicity studies and production process development proceed in parallel. If these stages are successful, registration and marketing of the new drug follow.

For a project, technical success refers to the completion of all stages of the R&D activities, including the necessary permissions by local authorities of at least one major market country, up to the point at which the results of the development can be handed over as a complete package to the functional group responsible for the commercial activities, which is the definition recommended by the European Industrial Research Management Association (EIRMA).

In pharmaceutical R&D, about one project in ten goes as far as registration, so that for one technically

1. E.B. Pyle, III, et al., "Scientific Manpower Allocation to New Drug Screening Programs," *Management Science*, August 1973, p. 1433, and William E. Souder, "Analytical Effectiveness of Mathematical Models for R&D Project Selection," *Management Science*, April 1973, p. 907.

2. Irwin Kabus, "You Can Bank on Uncertainty," HBR May-June 1976, p. 95.

3. William E. Souder, "The Validity of Subjective Probability of Success Forecasts by R&D Projects Managers," *IEEE Transactions on Engineering Management*, February 1969, p. 35.

successful project, in the neighborhood of 10,000 new chemical entities must be screened. Unfortunately, a technical success does not automatically imply a commercial success.

Subjective probability method

Explicit numerical representations of expert opinion are a regular input to the evaluation of important management problems in the subjective probability approach. A substantial literature exists recommending subjective probability judgments as the most appropriate basis for management decision making under uncertainty. The literature is far more limited on how to transform expert opinion into the quantitative probability distributions needed to apply such decision analysis methods. The most serious drawback to wider practical application of these methods, however, is the lack of almost any objective or scientific proof of the validity and reliability of the subjective judgment of experts in practical applications.

There are many reasons for this shortage of convincing evidence. One is the confidential nature of many of the analyses hitherto performed. Another is that successful business executives have no motivation to publish or disseminate their methods. Also, the unique, nonrepetitive character of many of the strategic problems to which decision analysis has been applied prevents any controlled comparison. Finally, only a few organizations have sustained experience with the use of subjective probability and its application to strategic decisions, planning, and control.

An outstanding example of a successful repetitive application of subjective probabilities was reported recently by the Morgan Guaranty Trust Company, which has used this method for the past six years to forecast future (90- and 180-day) interest rates.[2]

Morgan Guaranty has had excellent results in terms both of forecasting performance and of decision (i.e., money) making. The work reported here, however, differs from Morgan Guaranty's work in two important respects: the industrial research character of the events that are under consideration as well as the longer time horizons (one to five years) of the predictions.

Monsanto Chemical Company has reported another very relevant, but less sustained, examination of the validity of subjective probability in forecasting technical success.[3] In a study of 11 projects to define or develop new products or processes, subjective probability appeared to be valid and reliable for indicating future technical success as well as superior to the narrative status-reporting methods more typically used to control R&D projects.

The Monsanto study, however, has the major drawback that it was done only once and lasted only 12 months. Moreover, Monsanto's projects were actually only project phases lasting 12 to 30 months, and its criterion for technical success was the opinion of its own management rather than that of an external body like a regulatory agency. The experience that we describe here is significant in that it provides sustained evidence about the validity of subjective probability methods to predict technical accomplishments in the intermediate to long term.

Calculating with the probabilities

Judgmental probabilities, though subjective, can be represented in an objective form, that is, as numbers on a scale from 0 to 1 (or 0% to 100%). For example, suppose that there are three current R&D projects, each having an assigned subjective probability of success of 0.4. In the future, 0, 1, 2, or 3 of these projects will in fact succeed, with probabilities 0.216, 0.432, 0.288, and 0.064 respectively. This is a probability distribution with an expected value of $0.216 \times 0 + 0.432 \times 1 + 0.288 \times 2 + 0.064 \times 3 = 1.20$.

In general, as in the example, the expected number of successful projects is precisely the sum of the success probabilities of the individual projects. Note that while the expected value is seldom an integer, and therefore is seldom achievable, it does measure the average result that would be achieved by a large number of similar research efforts.

Another useful number is the standard deviation. For projects whose technical success or failure is independent of each other, the standard deviation is the square root of the sum of each probability (p) times its complement (1-p), which is 0.85 in our example. The standard deviation is useful because together with the expected value it defines an interval within which a significant fraction of the actual results should fall. Therefore, comparing the actual results achieved (i.e., number of successful projects) to the expected value and standard deviations of the prior forecasts will allow us to gauge the validity and reliability of this application of subjective probability as a forecasting method.

Probability encoding

The first step in the subjective probability approach is to translate the subjective judgment of individuals into a suitable numerical form. This is called probability encoding. Techniques for eliciting one individual's opinion about the likelihood of a particular event have been well described in the literature.[4] These techniques, however, differ somewhat according to whether the event in question takes discrete values (e.g., a drug receives FDA approval or not) or continuous values (e.g., variable production costs). The probability estimates for individual projects discussed here fall into the former class. However, for both classes of events—discrete or continuous—the probability wheel described by Spetzler and Staël von Holstein is useful.

R&D forecasters seem to fall into two categories: those who feel capable of giving direct numerical probability assignments and those who have difficulty in making such judgments. Most people seem to fall initially into the second category. Furthermore, many persons who prefer direct numerical responses are later found to have little confidence in their initial numerical responses. For this reason, indirect response techniques like the probability wheel are generally a better way to begin encoding. Later, when the forecaster is more familiar with the process, he or she often prefers to assign probabilities directly to the events.

Rare events (for example, events with probabilities of 1% or less) pose a special problem, since none of the standard probability encoding techniques works well for them. In this case, experience suggests that probabilistic modeling is often more effective than direct encoding. In R&D, the outcome of an event of interest often depends on the sequential outcomes of a series of other events. The intermediate events may have probabilities high enough to allow effective use of standard encoding procedures.

The problem of encoding the probability of a single rare event is thus transformed into the task of encoding the larger probabilities of these events. Our experience suggests that, for probabilities above 10%, direct assessment is adequate; for probabilities less than 10%, a sequential probability model is usually better. The probability that the screening activity discussed earlier will yield a marketable product is modeled as the product of the probability that screening will yield a project and the probability that a typical project will succeed.

A final important issue in probability encoding is group assessment. In many situations, there are a number of persons whose experience and judgment should be considered regarding the likelihood of a future event. Achieving a consensus is especially difficult for probability distributions over a continuous range of values. However, for discrete events such as technical success, there is only a single number—the probability of success—to determine, and therefore the Delphi method[5] provides a useful way to search for a consensus. Nonetheless, the Delphi method rarely produces unanimity, and it does not relieve management of the final responsibility for the probability assignment to be used for management-planning and decision-making purposes.

How the approach works

After a two-year review of the available methods for R&D planning, Sandoz introduced a systematic approach in 1969 by assigning probabilities for the technical success of the most important projects in the late stages of development. Since 1970, probabilities have been assigned to important projects even at the earliest (preclinical) stages. Since 1972, probabilities have been assigned to all development projects, with monthly updating of all changes in project status (e.g., phase transitions and terminations). Since 1974, the whole project portfolio has been updated twice annually.

For the first few years the probabilities were assigned by one R&D manager, who consulted individual experts. Since 1973, the forecasts have been made by an expert panel consisting of about a dozen R&D department heads and managers, as well as some clinical specialists. The project manager acts only

4. Carl S. Spetzler and Carl-Axel S. Staël von Holstein, "Probability Encoding in Decision Analysis," *Management Science*, November 1975, p. 340.

5. B. Brown and O. Helmer, "Improving the Reliability of Estimates Obtained from a Consensus of Experts" (Santa Monica, California: Rand Corporation, 1962.)

6. H.U. Balthasar and S. Gutzwiller, "Steady State and Portfolio Concept in R&D Management," *R&D Management*, June 1975, p. 201.

indirectly in the forecasting process; that is, he serves as consultant to the expert panel and reviewer of its results to minimize the unavoidable bias and enthusiasm of a successful project manager. The more detached, but highly knowledgeable, functional R&D department heads act effectively as a filter to any natural biases with minimal loss of information.

Success of single projects

Whereas initially the probabilities were obtained by interviews, when the panel experts became sufficiently familiar with the technique, Sandoz introduced the use of questionnaires. Now, each panel expert receives in a first round a list of all projects in development and is asked to assign a probability for the technical success of each. Then the questionnaires are collected and tabulated.

In a second round each expert receives the list of projects with the probabilities given by all the other experts; for example: Project A: 0.5, 0.7, 0.5, 0.6, 0.5, 0.2, 0.5, 0.5, 0.7, 0.6, 0.2, 0.6 (0.6); Project B: 0.3, 0.2, 0.3, 0.3, 0.2, 0.4, 0.3, 0.4, 0.4, 0.4, 0.4, 0.4 (0.3); Project C: 0.7, 0.7, 0.6, 0.9, 0.7, 0.7, 0.9, 0.8, 0.7, 0.7, 0.9, 0.7 (0.8).

In parentheses appears the "consensus" probability proposed by the R&D planner. This proposal is not simply an average, but takes into account the relatively greater importance of some aspects (e.g., toxicity) identified by certain experts. With this additional information the individual panel experts are then asked to review and possibly to revise their forecasts prior to a meeting of the panel group. The probabilities do not change much in the second round, but they do have a tendency to show less dispersion.

At the group meeting, consensus on one probability is achieved for each development project. Here again, certain aspects carry more weight than others in the discussion; a purely mathematical resolution is neither possible nor desirable. While a consensus cannot always be easily achieved, panel experts have learned that management, for the further use of the probabilities, needs just one figure. So the panel reaches a compromise even in the most controversial cases, which occasionally requires a final decision by the head of research.

A typical schedule for such an estimation procedure is as follows:

November 1:	Questionnaires sent out.
November 20:	Experts turn in their answers.
November 30:	Project lists with experts' initial assessments are sent out for a second round.
December 20:	Meeting of the panel experts takes place.
December 31:	Updated portfolio listing is available with the latest consensus probabilities.

Within the two-month period, each panel expert spends perhaps two to four hours on the forecasting activity, including the meeting of the panel.

Overall R&D success

That portion of the total R&D effort that is oriented to the discovery and development of new products can be seen as a portfolio consisting of preproject research plus all the development projects in progress.[6] Since new active substances emerging from preproject research require a minimum of six years to reach the market, the output of R&D for at least the next six years is determined by development projects already initiated.

Using standard network planning methods, an expected completion period is assigned for each project, determined by the information available. Finally, as a result of the process just described, each project is assigned a probability of technical success. By comparing the expected success with the actual success over a period of years, the reliability of the subjective probability method can be tested and validated.

Analysis of results

For all projects started since 1971 a specific project life chart exists that plots the probability of technical success (PTS) over the time that it was an active part of the R&D portfolio. All the curves start with the project's inclusion in the portfolio and stop either with its successful completion (registration) or with its exclusion from the portfolio (dropped, stockpiled, licensed out). A number of typical project histories are summarized in *Exhibit I*.

Exhibit I
Life chart of four projects progressing toward technical success

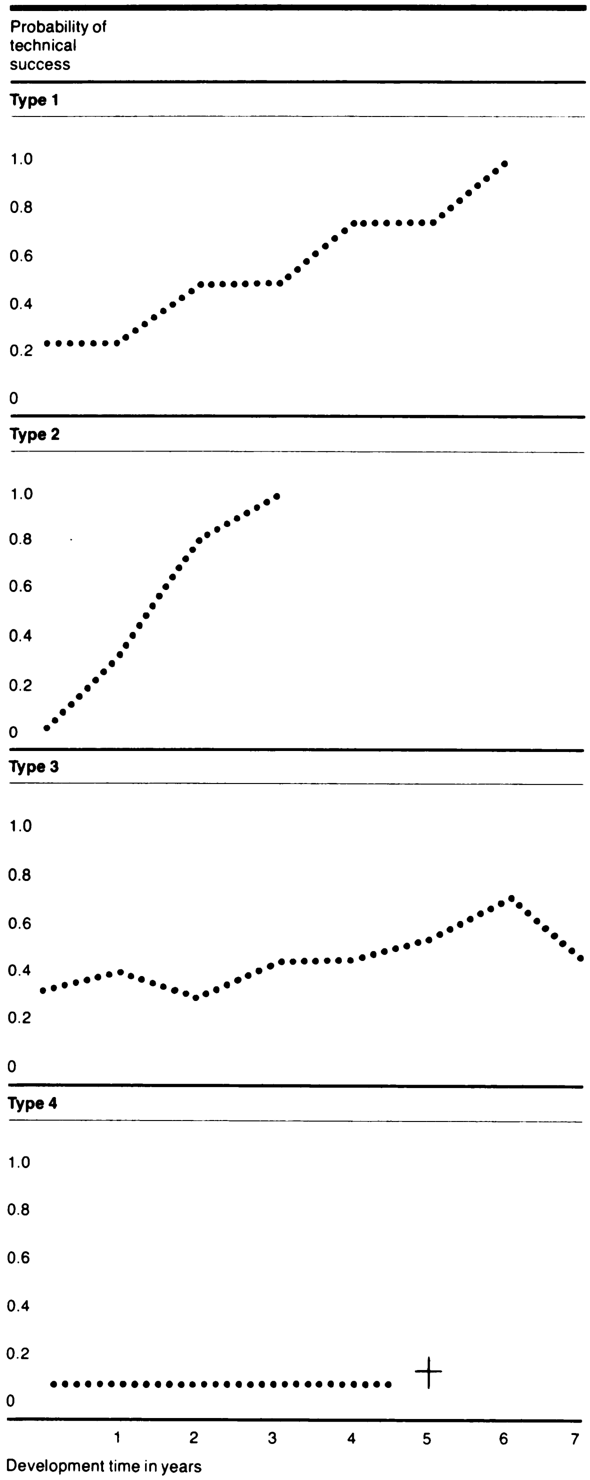

The four projects in *Exhibit I* illustrate common but distinct development patterns. Type 1 shows the normal career of a successful project, where the PTS steadily increases from the beginning to the end. Of course, this project may still be a commercial flop. Type 2 corresponds to a technically extremely successful project (priority). It is completed ahead of schedule. The completion of Type 3 is questionable and behind schedule. This implies a shorter commercial life (patent expiration, competition). Such a project may have been stalled for a time because of problems that could not be solved rapidly. The PTS of Type 4 is almost constant for many years. This is frequently the case with a hobby project. Such a project is normally dropped after a more or less long life.

The life charts of *Exhibit I* imply the following conclusions for management: if the PTS steadily increases, no action is required; whereas if it remains constant or decreases, the project should be reviewed and possibly terminated. The Monsanto experience referred to previously also reached this conclusion.

Progression of forecasts

Exhibit II shows the effectiveness of the management decisions taken by Sandoz on the basis of the success probabilities. The outlooks from December 1970, 1971, and 1972 showed erratic and declining expectations for future project success. After several years, however, the effects of management actions based on these probability forecasts can be seen in the outlooks from December 1973 and December 1974, which show a future prospect much closer to the desired steady state. Similar positive results continued in 1975, 1976, and 1977. This achievement is the kind of reward that can be expected by organizations that apply the probability method to forecast technical success on a dynamic basis.

Reliability of the method

Among projects in the portfolios of 1970, 1971, and 1972 that have been completed to date, there are from 5 to 20 actual results corresponding to each of the probability levels 0.1, 0.2, 0.25, 0.5, 0.75, 0.9, and 1.0. To test the reliability of the process, we have compared the probability assignments made before

Exhibit II
Progression of R&D success forecasts from 1970 to 1974

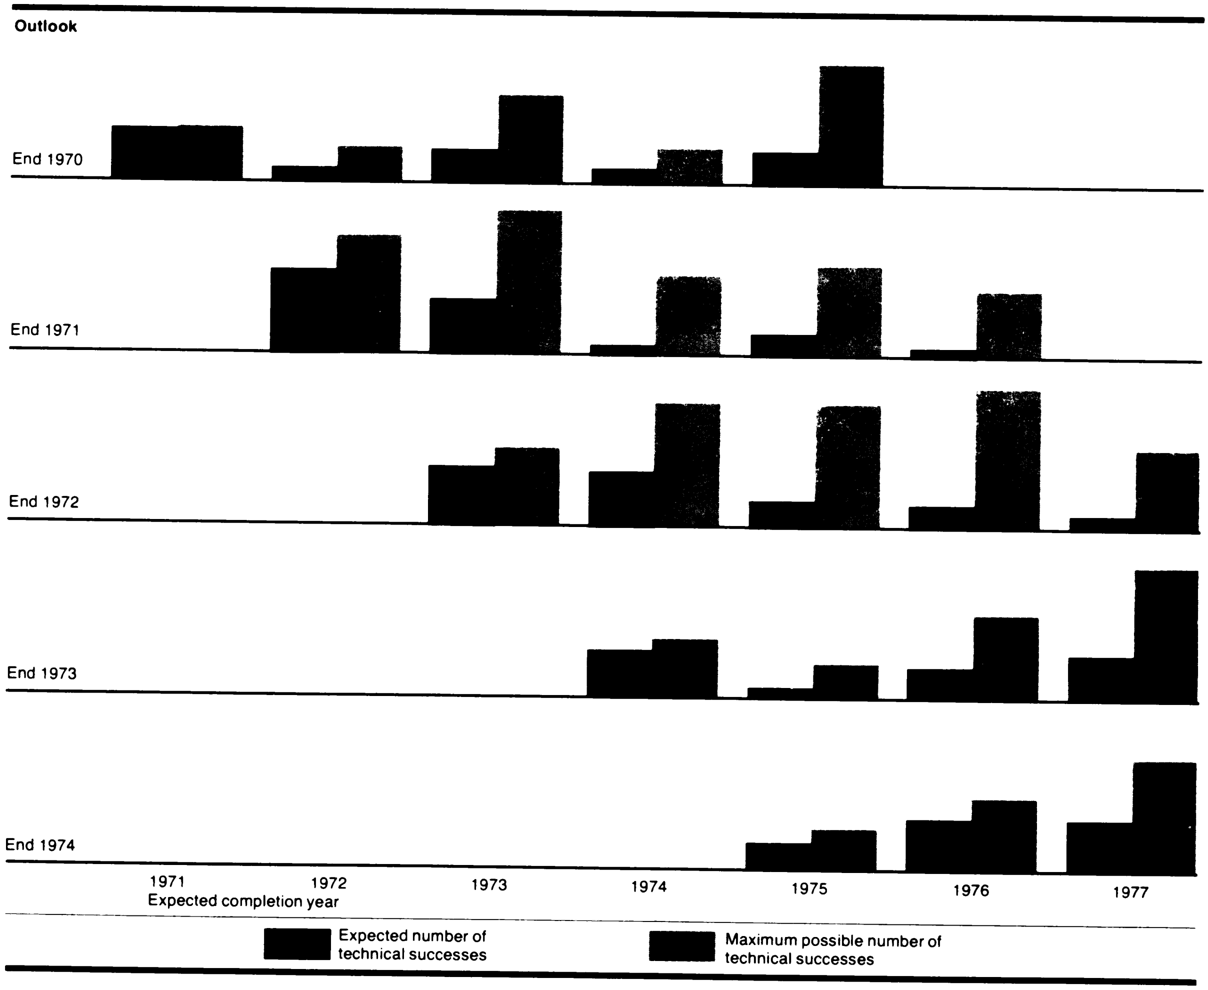

1973 with the observed frequency of success (see *Exhibit III*), - - - - - Points below the diagonal line indicate overoptimism, since the actual frequency of success is less than the prior probability; points above the axis, conversely, indicate pessimistic prior predictions. Thus this graph indicates that in the period 1970 to 1972 the success probabilities assigned *may* have been somewhat too high for low-probability (10% and 20%) projects but too low for higher probability (24% to 90%). Similar bias results have been reported by other observers.

The sample is small, but the considered data do not deviate significantly from the perfect calibration given by the diagonal of *Exhibit III*. This means that the expert judgment is statistically fully validated and that this statement should remain valid in the future. All points but one fall within standard deviation from the perfect calibration, and this indicates that early predictions using subject probabilities are a very useful guide to future results.

Success rate

Another way of judging the quality of the forecasts, which is especially appropriate, given the goal of achieving steady-state output from R&D, is to compare the number of projects in the portfolio that were expected to succeed with the number that actually did succeed. These comparisons for 1970 to 1973 are shown in *Exhibit IV*.

Exhibit III
Reliability of experts' judgment from 1970 to 1972

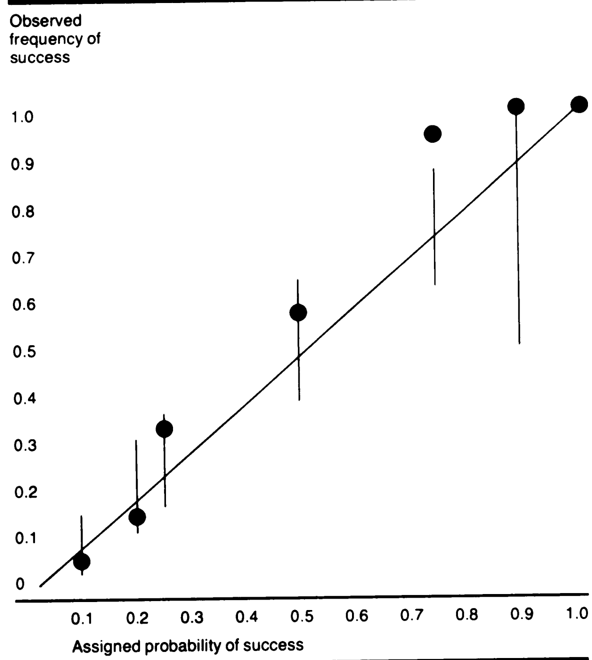

Each of the four outlooks compares the expectation with the actual result only for those projects that were included in the forecasting process. Thus actual results in successive years cannot be compared, since management was modifying the project portfolio and since the 1970 and 1971 portfolios did not include all development projects. The graphs show that successful products have been suitably predicted by the prior expected values.

A useful measure of predictability is to see how many of the actual results fall within the one standard deviation interval shown by the vertical bars in *Exhibit IV*. Assuming that all the squares not fully on the bars fall outside, then 12 of 18 actual results (67%) fall within the one standard deviation interval—quite reasonable for such distributions. Like the calibration test shown in *Exhibit III*, this finding gives confidence both in the approach and in the ability of R&D managers to use it successfully.

Another striking finding is that the uncertainty surrounding the success of the overall R&D portfolio of four and five years later, as shown by the standard deviations, is about the same as the uncertainty about results only one to three years away. This provides a quantitative and measurable example of the well-known portfolio effect, which achieves a relatively predictable result by combining a large population of individually uncertain events.

It is also reassuring that as far as they can be tested today, the actual results obtained in years 4 and 5 compare as well with prior expectations as the results for shorter time spans do.

Prerequisites for successful use

For a successful implementation of the subjective probability approach with an R&D organization, an important prerequisite is the combination of a promoter at management level who has appreciation of and an interest in the approach with an implementer in the R&D planning group who has technical knowledge of probability. Starting from this base, it is possible to inform the R&D managers and department and section heads about the approach over a longer period (several years) and to gain their active collaboration.

When the first attempts within a small circle have been found sound and workable, more persons can be gradually involved, up to the point where all people concerned with the projects' success are acquainted with the subjective probability method and the forecasts. By continually informing the panel experts about their estimates compared with the actual outcomes (survival rate of projects), their judgments can be improved.

Line managers are usually quite ready to give estimates. They are also interested in receiving good, realistic estimates. They see forecasting as a process of bringing to the surface their own, sometimes subconscious, judgments. They learn to express their feelings and opinions with more precision, which is educational as well as exciting. Our experience also suggests that this process of quantification has the effect of reducing the natural tendency toward personal biases.

Persons who are not in management positions, however, are much less interested in giving estimates. They prefer to remain neutral and avoid exposure. They are less willing to be held responsible for management decisions such as dropping a development project by assigning a low probability of success.

Exhibit IV
Comparing forecasts with subsequent results

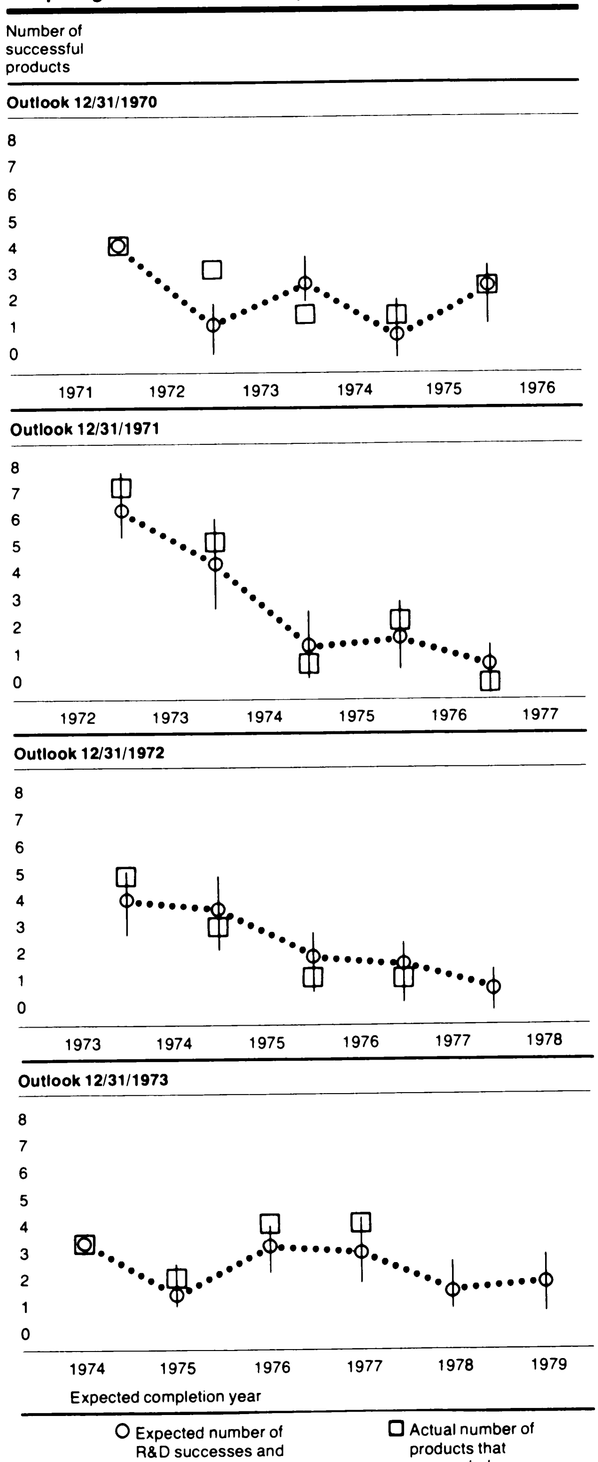

Costs & benefits

As for any new management method, there are two kinds of costs, namely, initial investment and operating cost. In our case, the operating cost has proved to be quite low: about three to four man-months per year, including the time of managers, planners, technical assistants, and secretaries; this is far below 0.1% of the overall R&D budget. Of this, only ten to twelve man-days are required from the R&D line managers, who are extremely busy people.

However, that time is not necessarily entirely incremental to their normal activities, because these are the people who would in any case spend considerable time discussing project priorities and chances for success. It is even possible that, if they determine a reliable consensus probability for the successes of 40 to 60 projects twice a year in the equivalent of no more than one man-day per person, this method may in fact be quite an efficient use of their time compared with more verbal, qualitative methods.

The initial investment may be considerably more of a deterrent to organizations wishing to introduce such a method. In the first place, at least one of the top line managers needs to become convinced that such a method is necessary and desirable.

Without his support it is unlikely that the method can be successfully established. Another initial cost is the development of, first, a willingness to try, then, an understanding of, and, finally, confidence in, the method among the R&D line managers. This method requires a balanced program of education and application.

As a primary benefit this method provides an advance indicator of the eventual success of both individual R&D projects and the overall R&D process to allow effective management action. Better decisions can be made regarding the individual projects, and better control can be exercised over the portfolio of development projects. Moreover, because of the portfolio effect, the uncertainty surrounding the overall success of the pharmaceutical R&D effort five or six years in the future is essentially no greater than that surrounding next year's results.

An important secondary benefit is the role of the probability encoding in clarifying individual opinions and resolving disagreements. If R&D managers focus primarily on the chances of technical success while marketing people focus on the likely commercial value of such success, not only is the possibility of reaching agreement among R&D managers enhanced but also the communication between marketing and R&D is increased.

This research was motivated primarily by the continuing efforts to develop advance long-range success indicators of the R&D program as well as to develop more useful methods for strategic planning and management control in a highly uncertain environment. The experience reported here shows that simple methods give forecasts that are valid as the basis for management decision in R&D.

Since the uncertainty faced by R&D is at least as great as that faced by other business functions, our evidence on the validity of subjective probability methods in R&D should be relevant to managers in all areas of business facing conditions of uncertainty. Therefore, this report on a real world case should be useful to all managers and executives concerned with improving their strategic and operational planning, decision making, and management control processes.

The value of opinion

Where there is much desire to learn, there of necessity will be much arguing, much writing, many opinions; for opinion in good men is but knowledge in the making.
Milton, *Areopagitica*

I dogmatise and am contradicted, and in this conflict of opinions and sentiments I find delight.
Johnson, *Extracts from Hawkins's Life of Johnson*

Opinion is ultimately determined by the feelings, and not by the intellect.
Spencer, *Social Statics*

So many men, so many opinions; every one his own way.
Terence, *Phormio*

He in whom the love of truth predominates . . . submits to the inconvenience of suspense and imperfect opinion; but he is a candidate for truth . . . and respects the highest law of his being.
Emerson, *Intellect*

What we have to do is to be forever curiously testing new opinions and courting new impressions.
Pater, *The Renaissance*

20

QUANTIFYING AND FORECASTING EXPLORATORY RESEARCH SUCCESS

Roberto A. A. Boschi
CIBA-GEIGY

Hans Ulrich Balthasar
Cilag, AG, Switzerland

Michael M. Menke
Strategic Decisions Group

Reprinted with permission from Research Management,
Volume XXII/No. 5 (September, 1979), pp. 14-21.

Quantifying and Forecasting Exploratory Research Success

Roberto A. A. Boschi, Hans Ulrich Balthasar and Michael M. Menke

This article describes how managers have improved the effectiveness of product-oriented basic research by using subjective probability to quantify and forecast research success.

Business planning cannot be done without forecasts. However, both planning and forecasting in R&D suffer from the obvious problems that research success is difficult to define and impossible to predict with certainty. Probability theory provides a method that is ideally suited to predict uncertain events like future research success. While this paper focuses on the pharmaceutical industry, the subjective probability approach is applicable to other R&D areas as well.

Subjective probability provides two important features, a scale upon which to quantify the uncertain future success of research and a method to forecast the degree of future success that is expected from the R&D activities. Since the probability of success in basic research is low and since the operational experience with this approach is still short compared to the time required to achieve success, it is not yet possible to validate this method statistically (as has been possible for drug development projects (*2*)). However, it is possible to draw the very important qualitative conclusion that this approach has been accepted by research managers in several companies as a helpful and practical planning method that allows them to provide meaningful inputs to the overall company planning process without sacrificing creativity, neglecting experience, or being forced to make unreasonable commitments.

Most pharmaceutical R&D operations involve three major stages or phases of R&D: 1) basic research, or screening as it is often called in the pharmaceutical industry, 2) the pre-clinical testing of the most promising individual compounds, and 3) the clinical development of those products that exhibit efficacy and safety in animal investigations. The latter two stages are usually called product development and generally consist of a large number of well structured projects (*8, 16, 17, 18*) to develop specific compounds into new drugs.

This paper concentrates on the basic exploratory research that takes place prior to the selection of specific compounds for development. We call this basic research activity "pre-project research." Here the chemistry, biology, medical and marketing departments discuss together into what areas research should be directed. The chemists have to know starting points for their chemical synthesis and make sure that there are a feasible series of new compounds. The biologists must make available an appropriate test battery allowing for the detection of active and safe compounds. This discussion between scientists and marketing people leads to a definition and ultimately a selection of research areas, within which the main effort will be concentrated. Once a research area portfolio is established, the chemists and biologists compile their specific research program and goals. These projects contain a chemical synthesis program and an ingeniously arranged battery of biological tests, as well as established selection criteria and decision rules (*16, 21*).

So far the process described is known to everybody in the pharmaceutical industry and that is how most companies do it. It is at precisely this point that the main problem occurs: How to measure progress? How to evaluate the working programs? How to know when to add resources or to drop a program? The measure and yardstick that one can now use in such a situation is the portfolio concept combined with assessments for each program of the probability of success, the time and cost

Dr. Boschi is a member of the R&D planning group at Sandoz Ltd., Basel, Switzerland. Dr. Balthasar is marketing manager for Diagnostica Products, F. Hoffman-La Roche & Co., Ltd., Basel. Dr. Menke is manager, decision analysis for SRI International in Europe. He is located in England.

needed to achieve such success, and the expected market reward for technical success (*3, 4, 6, 7, 9, 10, 14*). It is clear, although difficult, how to quantify time, cost and the potential market reward for product oriented research. The contribution of this paper is to describe a practical and workable method to quantify the probability of research success and to combine it with the other essential criteria.

Careful thought must be given to the definition of research success and development success, independent of ultimate commercial success. This can be done using the phased structure of R&D. As a first step, research management divides the full research and development process from the new idea until the finished product into a series of phases, like those listed above. These phases have operational time horizons which can range from six months up to 3 or 4 years. Thus the work is broken down into manageable pieces and proportions, separated by intermediate goals or so-called milestones. Management defines what activities have to be completed and what status has to be achieved at such a milestone (*8, 11, 16, 18*). The transition into the next phase only shall happen if all the activities and requirements of a phase are fully complete. The success of each phase can then be defined as the completion of all intermediate goals, allowing a transition into the subsequent phase. Thus, for example, the success of pharmaceutical pre-project research can be defined as having a compound selected for preclinical development, whereas the ultimate technical success of pharmaceutical R&D is acceptance for marketing by the various regulatory agencies (*2, 8, 18*). A similar definition of intermediate research project success has been used at the Monsanto Chemical Company (*19*).

This definition of intermediate success is essential for two reasons. One is that for events with probabilities less than 0.1, it is often more meaningful to assess their probability indirectly as the joint probability of a sequence of related events, each of which has a larger and more easily comprehensible probability. This is because human experience with low probability events is even more restricted than our limited experience with probability in general. The second reason is that in basic research, the time periods from the initiation of effort to a marketed product can be discouragingly long. If there were no intermediate goals, success could not be achieved within the typical time span of a researcher's job, leading to demotivation and demoralization. Research success must be defined within the time span of the researcher and the control span of the research manager.

The Portfolio Concept

Just as project management methods have proved useful for the management and control of single projects containing a series of activities, portfolio management methods are useful for the management and control of the full R&D program containing many different projects and research efforts. The portfolio concept has become widely accepted as a method for managing financial investments (*13*), but it can be equally useful for R&D management (*2, 3, 5, 9, 10*). Financial portfolio management was the source for the idea to develop the methods described in this paper. In R&D, however, it is important to distinguish between two parts of the portfolio (*3*): the research area portfolio and the development project portfolio.

The development project portfolio contains all compounds that have been selected for preclinical development up to submission for registration and to introduction into the market as a new therapeutic principle. Its successful application to the management of drug development has been described elsewhere (*2*). The research area portfolio contains all exploratory product-oriented research work being done in the pre-project stage. This work can be structured with a hierarchy of indication and working areas. Individual research programs with specific therapeutic goals within a particular working area are called topics. Figure 1 gives an example of such hierarchy with "indication areas," "working areas" and "topics."

The research area portfolio is selected by R&D management and should be updated annually, using defined criteria for the evaluation of the individual research ideas, the individual topic, working areas and overall indication areas. The evaluation and selection process leads to a conversion of research ideas into working areas and topics. Within these working areas and topics R&D managers and scientists can remain flexible and can adapt the approaches best suited to the available opportunities. Among the numerous evaluation and selection procedures published (*6, 7, 9*), one useful source is the checklists in the handbook of Heyel (*12*).

Assessing the Probability of Success

An essential, but challenging aspect of research is the role of chance. The characteristics of research preclude any attempt to describe *a priori* every possible future situation to which a specific research topic could lead, as might actually be possible for the development of a specific compound as a new drug. However, being aware of the exploratory nature of research, experienced researchers and research managers can take this into account in specifying how long they feel may be required to reach a specific goal, and indeed what chances there are to reach the goal within various time frames. The use of probability is thus a bridge, allowing one's past experience with research projects plus any special features of the current effort to be combined and summarized in a simple and useful way.

The key to effective application of the probability method is a clear definition of the event in question. The success of the pre-project research effort

LEVEL OF HIERARCHY OF RESEARCH AREA		RESERACH AREA EFFICACY GOAL
Indication Area	1.	Gastrointestinal Tract
Working Area	1.1.	Influence on Incretory/ Secretory Activity
Topic	1.1.1.	H_2-antagonists
Topic	1.1.2.	Hormones (gastrin, secretin)
Working Area	1.2	Ganglio-Muscular Approach
Topic	1.2.1.	Muscle tonus and/or peristatic activity in-/de-creasing agents
Topic	1.2.2.	Synaptical transmission
Topic	1.2.3	Ganglio-stimulant/-inhibitor agents
Working Area	1.3	Dietary Influence
Topic	1.3.1.	

Figure 1/ Example of an Indication Area with Working Areas and Topics

within a research area can be defined as finding a compound possessing enough promising properties to be selected for further development as a "project." This definition, while flexible, is precise enough that experienced research managers can assess the probability for such a success within specific time intervals. In order to develop as reliable as possible a view of their judgment several different types of questions can be asked, either by fixing time horizons and assessing the probability for a success within that horizon or by asking the manager to assess the time horizons corresponding to several probability levels for the success event. Ideally both types of questions should be asked to provide a consistency check on the judgments.

An example of the first type of question is to assess the probabilities that a research topic activity will yield a development project within, say, three years and six years time. In most private companies, if management does not see a reasonable probability of achieving research success within six years (i.e. promotion of at least one research discovery into the product development stage), then a reexamination or redirection of the research activity may be undertaken. The second type of question can be expressed as follows: "How many years T_1 of effort in Topic X are required to give 30% probability of producing a project? How many years T_2 to give 80%" T_1 and T_2 approximate the "earliest" and the "latest" point in time of a success event, as well as defining a time interval containing a 50% chance for success. Moreover, since the probability of achieving success within a very short time interval is very small, the probability of immediate success is zero.

The probability/time estimates can be obtained from the respective research managers using a structured interview of the type described by Spetzler and Stael von Holstein (*20*). As such they represent the managers' best, although not necessarily subjective, judgment regarding the chances for success as a function of time in the various research topics. When starting their work they give an estimate for the earliest and the largest point in time which they feel will be required to achieve success (a development project). Our experience in many companies has shown that researchers have a well developed sense for this type of estimate. They are able to define quite well the time horizon they assume for their work. They can say that for most programs if they do not have a clear indication of research success within, say, 6 years, they would abandon this research area and switch to another one.

Figure 2 illustrates the cumulative probability curves corresponding to the probability assessment and time estimates for three of the topics within a particular working area where each research topic is confident of achieving success. The horizontal axis gives the time in years, the vertical one the probability levels for achieving success by that time. To each topic corresponds a probability distribution for success over time, shown here as an S-shaped curve starting at the origin, passing through various points corresponding to the questions described above and aiming asymptotically at the probability level of 1.0.

The S-shape of the curves shown in Figure 4 is characteristic of research projects only recently initiated, where the expectation of success in the near future is very low due to lack of specific experience and knowledge. However, for research projects that have been under way for a sufficient period of time, the necessary expertise will have been developed so that the next discovery is just as likely next week, as the week after, as in any week this year. For such research projects the forecast of success versus time would look like Topic A in Figure 3. The shape of the curve is exponential, which is based on the assumption that for short time intervals the probability of a discovery within that time interval is proportional to the length of the interval. This model seems very reasonable for all research projects that have accumulated sufficient experience and expertise in that field.

Another more general aspect is that in research a topic may be initiated in a particularly risky class of compounds realizing in advance that no development project may ever be found. This view is contained in the forecast for Topic B in Figure 3, which says that there is one chance in three that this research effort will end without success (other than the new knowledge gained). Although a private enterprise company cannot afford to pursue knowingly too many such topics, the high risk may sometimes be justified by the size of the reward if successful. The shapes of the curves shown in Figures 2 and 3 are general enough to deal with the prospects of almost any basic research project.

The data base consisting of curves such as those

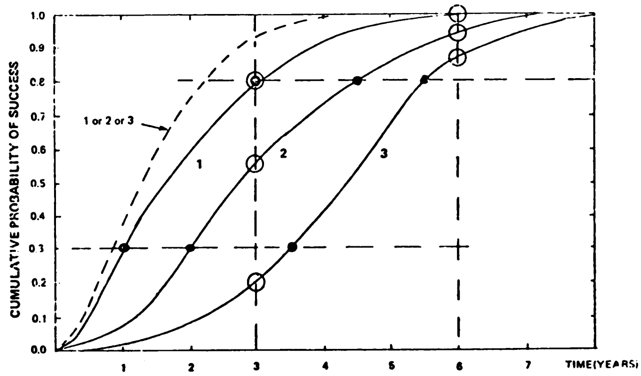

Figure 2/Graphical Display of the Probability of Success Over Time for Three Different Research Topics

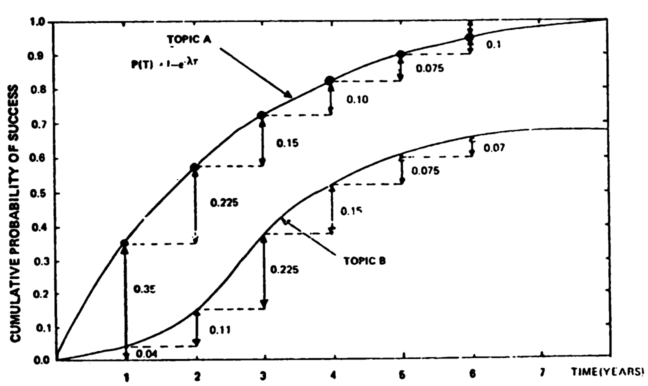

Figure 3/Forecast of Research Project Success for a Well Established Research Topic (A) and a New, Very Risky Research Topic (B).

shown in Figures 2 and 3 for each research topic allows a wide range of analyses for management planning, decision making and control. One of the simplest and most useful, already discussed in Reference 3, amalgamates this information at the level of working areas, indication areas, or the entire research portfolio to display risk profiles and time profiles for the corresponding research effort. For example, one could take a useful planning horizon like 3 years, and then group the topics according to their probability of success within that time span. Figure 4 shows how the resulting "risk profile" might appear. An analogous "time profile" can be generated by displaying the time required for each topic to reach a reasonable probability of success.

The displays shown in Figure 4 do not of course provide the complete information required for R&D decisions, but in practice we have found that they introduce an important new dimension into management discussions regarding the R&D portfolio. They present graphically the time-risk tradeoff and as such appear to give managers a much sounder basis for discussing these issues and incorporating them in the R&D decision process. Moreover, their value as a control indicator becomes significant when they are monitored over a period of 3 or more years.

Expected Times to Success

Since the curves in Figures 2 and 3 are obtained directly from the judgment of the responsible researchers, there is no *a priori* reason that they should possess any particular functional form, although for research topics that are already well established there are good theoretical and practical reasons why they should approximate a distribution in the Binomial (Poisson) family. Nonetheless, if they appear symmetrical it is possible to use a known symmetrical distribution, such as for example the Gaussian error function, to approximate them and thus to derive an estimate for the average time before a project will emerge from that research topic in terms of the directly encoded times T_1 ($p = 0.3$) and T_2 ($p = 0.8$). Similar analyses can be made for non-symmetrical distributions, such as the log-normal or Poisson.

Of perhaps greater interest to marketing is to know the probability that an indication area or working area will yield at least one project from its topics within various time horizons. This can be simply reformulated as one minus the probability that none of the topics will yield a project by that time. For any time, this quantity is simply the product of all the distances from the top of the graph down to the different probability curves. Then in the example of Figure 4, assuming only the three topics shown in that particular working area, the probability of at least one project emerging within three years time is $1 - (1 - 0.8)(1 - 0.55)(1 - 0.2) = 0.93$, assuming that the success of these three topics is technically independent. The dashed curve in Figure 2 summarizes the probability that 1 or 2 or 3 will succeed as a function of time. Clearly the combined curve will always be above and to the left of that for the leftmost topic. Using such combined distributions, research and marketing management can determine the probability to get new preclinical candidatges in important areas like antibacterials or cardiovascular drugs within specified time horizons. Such information is of course very important to the commercial management in order to plan the required production, marketing and financial resources.

Connecting "R" Data with "D" Data

For research management it is interesting to have a look at pre-project work alone as well as project work alone. However, it is also important to see the profiles of the entire R&D operation, and especially how they change over time (*16*). Whereas in the individual analyses one can calculate the distributions only by counting the projects or the

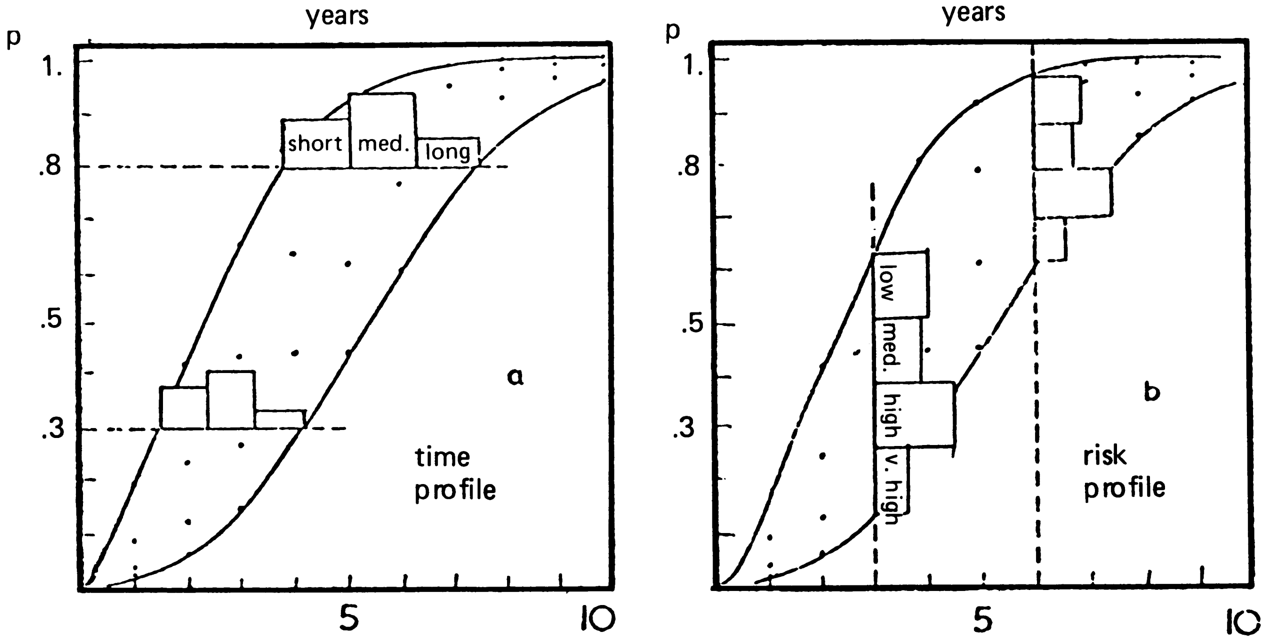

Figure 4/Display of "Time" and "Risk" Profiles; Possible at the Working Area, Indication Area, or Total Research Portfolio Level

topics, before adding "R" and "D" together, both need to be brought onto a common scale. In the case of resources, this can be done by weighting the work done for both the research topics and the development projects with the costs allocated to them over the considered period. The costs then measure the resources allocated to the various pieces of research. While this task requires a very careful definition of the various aspects of R&D work, it is conceptually straightforward (*10, 11, 12*).

Attaining a common scale to measure combined R&D success is more subtle. For the research topics the success definition must be extended: instead of taking the generation of a project as the only yardstick for research success, by analogy to the development project the achievement of a registered product is also taken as the research criteria. This means that the time horizon must be extended by the average 6 to 7 years necessary to bring a project through all the development phases to the registration goal and the probability must be reduced by a factor corresponding to the average probability for a new development project to succeed. In practice this factor is about 0.1, because an average compound just beginning preclinical investigations has about a 10% chance to reach the market. Thus, if a topic is assesed to have a 50% chance to yield a project within 2 years, then there is a 5% chance to have a product from that research effort within eight or nine years.

After this recalculation the same distributions that have been described above for the pre-project research alone and in References 2 and 3 for product development alone can be plotted for the full R&D operation. The impact of the pre-project part of the research in terms of new products can of course only be seen in years 6 and beyond.

Having estimates of the relative chances for success for both the R and D parts of the R&D operation allows management to make some simple but useful consistency checks on the resources allocated to research and to the subsequent development activities. If the overall success probability for R plus D is for example 0.16, it makes a great difference whether this is due to 0.2 x 0.8 or 0.4 x 0.4. The former needs more R activity to feed the D; the latter requires relatively more development effort. Moreover, if the nature of either R or D changes over time (e.g. due to increased regulatory requirements) then a corresponding change in the other activity may be required to maintain a proper balance of effort. This was discussed more extensively under the heading "steady state" in Reference 3. Thus the probabilities of technical success for both R and D — *monitored over time* — provide very important management control information for the R&D director (*3, 7, 19*).

Linking Research Success to Market Success

The discussion above does not make the final essential extension of the research portfolio output to commercial results. This requires in particular the ability to estimate the monetary sales figures associated with research results many years in the future. However, given the product orientation of pharmaceutical pre-project research, as illustrated by the medical/therapeutical efficacy goals of the research topics shown in Figure 1, there is a sound basis upon which to make long range predictions for sales corresponding to the successful accomplishment of the basic goal. Experience in the pharmaceutical industry has shown (*1, 3, 4*) that for areas where a reasonable, though not optimal, drug therapy is currently available, rather accurate forecasts can be made, which may however be conservative in that additional indications may be found. For areas where no reasonable drug therapy is presently available, the actual market potential of a therapeutic breakthrough is often underestimated.

Accepting these limitations, the pharmaceutical industry has developed some highly effective long range sales forecasting methods which allow a sufficiently accurate assessment of the ultimate commercial success to be expected from the majority of industrial pre-project research. Abt et. al. (*1*) have described in detail a methodology used by CIBA-GEIGY for assessing the peak sales (in both volume and monetary terms) corresponding to specific (but variable) assumptions about R&D results as well as future competitive product introductions and overall market trends.

Studies of many pharmaceutical product life cycles have also indicated that from 4—7 years is typically required for a new product to reach its peak sales potential in a particular market area, and for all but the most successful products the sales tend to remain at this peak level for a similar period of time. This common finding is summarized in the generic product life cycle shown in Figure 5. Here the word generic is used to indicate that for different values of the many parameters this same life cycle

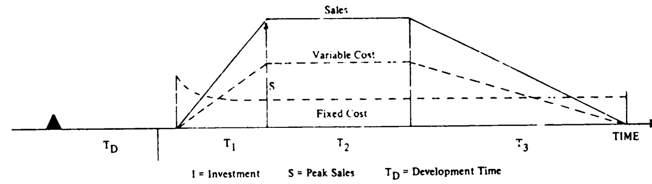

Figure 5/Generic Product Life Cycle Model.

can be used to represent nearly any new pharmaceutical product. Given the degree of uncertainty associated with products that will be introduced only 5 or more years in the future, it is clear that no greater degree of accuracy would be warranted.

Such a generic life cycle, with appropriate parameters for introduction date, peak sales, and probability of technical success can clearly be evaluated for each project in a development portfolio, as has been discussed in Reference 3. Table 1 shows how the expected sales results from a portfolio of new products under development have been assessed. The expected value method (*15, 22*) is simply to multiply the probability of successful development by the peak sales forecast for a new product corresponding to the therapeutic efficiency goal of that research topic. The resulting "expected" peak sales figures (Column D), which represent the mean or average values in the usual sense of the word, are hypothetical for any single

Table 1/*Development Project Portfolio with Expected Sales*

Status: December 1979 A: Year of Registration; B: Sales in the Peak Year; C: Probability of Success; D: Expected Sales (BxC)

Preparation	A	B	C	D	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990
INDICATION AREA 1															
Project A	1985	30.0	.10	3.0	.0	.0	.0	.0	.0	.0	.0	.6	1.2	1.8	2.4
INDICATION AREA 2															
Project B	1985	50.0	.10	5.0	.0	.0	.0	.0	.0	.0	.0	1.0	2.0	3.0	4.0
Project C	1983	100.0	.50	50.0	.0	.0	.0	.0	.0	10.0	20.0	30.0	40.0	50.0	50.0
Project D	1983	100.0	.10	10.0	.0	.0	.0	.0	.0	2.0	4.0	6.0	8.0	10.0	10.0
Project E	1984	100.0	.10	10.0	.0	.0	.0	.0	.0	.0	2.0	4.0	6.0	8.0	10.0
INDICATION AREA 3															
INDICATION AREA 4															
Project F	1983	60.0	.10	6.0	.0	.0	.0	.0	.0	1.2	2.4	3.6	4.8	6.0	6.0
Project G	1982	100.0	.10	10.0	.0	.0	.0	.0	2.0	4.0	6.0	8.0	10.0	10.0	10.0
Project H	1983	50.0	.10	5.0	.0	.0	.0	.0	.0	1.0	2.0	3.0	4.0	5.0	5.0
Project I	1982	100.0	.10	10.0	.0	.0	.0	.0	2.0	4.0	6.0	8.0	10.0	10.0	10.0
Project J	1980	90.0	.75	67.5	.0	.0	13.5	27.0	40.5	54.0	67.5	67.5	67.5	67.5	67.5
OTHERS															
Project K	1980	50.0	.75	37.5	.0	.0	7.5	15.0	22.5	30.0	37.5	37.5	37.5	37.5	37.5
Project L	1982	40.0	.25	10.0	.0	.0	.0	.0	2.0	4.0	6.0	8.0	10.0	10.0	10.0
Project M	1981	150.0	.10	15.0	.0	.0	.0	3.0	6.0	9.0	12.0	15.0	15.0	15.0	15.0
Project N	1982	50.0	.50	25.0	.0	.0	.0	.0	5.0	10.0	15.0	20.0	25.0	25.0	25.0
TOTAL PORTFOLIO					.0	.0	21.0	45.0	80.0	129.2	180.4	212.2	241.0	258.8	262.4

Table 2/*Derivation of Future Sales Expectation from a Specific Research Topic*

TOPIC A: Gastrontestinal Tract, Improved Efficacy. Peak Sales Estimate if Successful 50 Million. Time Required for Development 5 Years. Probability of Development Success 0.2

	Year of Research Success	Probability of Success in that year	FUTURE ANNUAL SALES EXPECTATIONS FROM RESEARCH TOPIC A: 1-5	6	7	8	9	10	11	12	13	14	15
Sales Expectation if research success occurs in year 1 and development is successful			0	10	20	30	40	50	50	50	50	50	50
Sales Expectations from Research Success in any Given Year (assuming successful development)	1	0.35	0	3.5	7.0	10.5	14.0	17.50	17.50	17.50	17.50	17.50	17.50
	2	0.225	0	0	2.25	4.5	6.75	9.00	11.25	11.25	11.25	11.25	11.25
	3	0.15	0	0	0	1.5	3.0	4.5	6.0	7.5	7.5	7.5	7.5
	4	0.1	0	0	0		1.0	2.0	3.0	4.0	5.0	5.0	5.0
	5	0.075	0	0	0			0.75	1.5	2.25	3.0	3.75	3.75
	6+	0.1	0	0	0				1.0	2.0	3.0	4.0	5.0
Sales Expectation if year of research success is uncertain but development is successful			0	3.50	9.25	16.50	24.75	33.75	40.25	44.50	47.50	49.0	50.00
Sales Expectation when development has only a 20% probability of success			0	0.7	1.85	3.3	4.95	6.75	8.05	8.90	9.50	9.80	10.00

project, but their sum for the total develoment portfolio will be quite reliable if the probabilities and peak sales are reliably assessed. References 1 and 3 indicate that this is the case. Therefore, following our discussion in the previous section, "Connecting Pre-Project Data with Project Data," an analogous method can be developed for the pre-project research portion of the R&D portfolio in 4 steps. The result is a simple but reliable method to indicate the range of sales values that can be expected in the time horizon 6—15 years from basic research projects underway at the present time.

First, for each research topic the probability of yielding a project in each of the next several years can be determined from the curves shown in Figures 2 and 3. Second, a peak sales estimate for a successful compound from this research topic must be provided, considering market conditions expected ten or more years ahead. Third, the sales expectations from initiating a successful development project in each of the next several years must be developed using a reasonable set of product life cycle parameters and deferred by the time required for development. Fourth, these sales expectations must then all be reduced by the average chances for development success for a new compound just beginning development in that particular indication area.

The points shown on Topic A of Figure 3 above the Time values of 1,2,... 6 years indicate that there are corresponding probabilities of 0.35, 0.225, 0.15, 0.10, 0.075 and 0.05 to get the first success from this topic in each of those years. This is the first step. The next two steps are illustrated by Table 2, which shows how the uncertainty in the time foreseen before achieving the first development project from the research program described as Topic A translates into a spreading of the corresponding sales expectation over time. Moreover, the row called Sales Expectation (if Year of Research Success is Uncertain but Development is Successful) must be reduced by the probability assigned for development success, in this case 0.2, to produce the row called Sales Expectation when Development has only a 20% Probability of Success. These figures, starting five years from the current year, could then be entered as a row describing the expected sales from Topic A into an extension of Table 1, which can now accommodate both research and development projects on an equal footing.

It must be emphasized that the individual rows of Tables 1 and 2 are not meaningful by themselves. Development projects either succeed or fail and research success (as defined here) occurs in a specific year. However, for an R&D portfolio containing a number of different projects whose probabilities of technical success are independent, the sum of the expected sales from all projects in a given year is a useful and important figure, about which an operationally meaningful deviation can be calculated. Thus, for not only development, but also research as well as the total R&D portfolio, it is possible to develop a systematic and meaningful forecast of the future sales to be expected as a result of these activities on the basis of a few simple quantities that can be reliably estimated today.

Applications and Results

Several applications on a large scale over several years have shown that the method is accepted by both researchers and research managers (*4, 6, 14*). It has been found practical, understandable and inexpensive. The probability profiles and time horizons do indeed show the status of research and development in a company. At a single point in time they are of limited use to management, but after one or two years, when time series of such figures are available, then much more importance can be assigned to these data because they show clearly the trends in R&D.

By selection of appropriate research areas and development projects and by conscious decision regarding the resources allocated to such efforts, management can actively influence the time- and risk-profiles of the R&D portfolio. The method therefore allows an R&D organization to react quickly and in the desired direction. In the case of the development of project portfolio, the trends have already clearly shown the positive impact on the operation from management decisions based on the probability forecasts (*2*).

This type of R&D management does not exclude satisfaction for the researchers. On the contrary, to the scientific satisfaction of having done creative chemistry or innovative biology is added the reward of research success when a new project is selected and the experience of the pre-project researchers of performing a job not only accepted, but highly esteemed by the other functions. Development, marketing and finance can become witnesses of the success of the pre-project research departments within their company. With this short-range yardstick the success of the basic work in the R&D operation becomes fractionally measurable. Without the techniques described in this paper the ultimate success (i.e. a market product) is a rare event coming to surface only seven to fifteen years after the initiation of research work. This means that there is no real control and little correlation between research effort and commercial success, often leading to frustration for the basic researcher and sometimes even leading him to focus on "private" objectives.

References

1. Abt, R. et. al. "The Dangerous Quest for Certainty in Market Forecasting," *Long Range Planning, 12* (*2*), pg 52, April 1979.
2. Balthasar, H.U., Boschi, R.A.A., Menke, M.M., "Calling the Shots in R&D," *Harvard Business Review*, Vol. 56 (3), pg 151 (May—June 1978).
3. Balthasar, H.U., and Gutzwiller, St., "Steady State and Portfolio Concept in R&D Management," *R&D Management*, Vol. 5, No. 3, June 1975, ppg 201—207.
4. Bender, A. Douglas, et. al. "Stimulation of R&D Investment Strategies," *Omega, 4* (1) pg 67 (1976).
5. Channon, Derek, F., "Use and Abuse of Analytical Techniques for Strategic Decisions," Manchester Business School (1977).
6. Cochran, Michael A., et. al. "Investment Model for R&D Project Evaluation and Selection," *IEEE Transactions in Engineering Management, EM-18* (3), ppg 89—100, August (1971).
7. Cranston, R.W., "First Experiences With a Ranking Method for Portfolio Selection in Applied Research, *IEEE Transactions on Engineering Management, EM-21*, ppg 148—152, (1974).
8. Djerassi, Carl, "Birth Control after 1983," *Science, Vol. 169*, pp 941—51.
9. Freeman, P., and Gear, A.E., "A Probilistic Objective Function for R&D Portfolio Selection," *Operational Research Quarterly, Vol. 22*, No. 3.
10. Girkinger, W., "Konzeption eines integrierten Planungs- und Budgetierungs-systems," *die pharmazeutische Industrie, 40*, Nr. 9 (1978), pg 903 ff. and Nr. 10 (1978) pg 1028 ff.
11. Hettinger, Wm. P., "Research Planning," *Research Management*, Vol. X, No. 4, (1967).
12. Heyel, C., *Handbook of Industrial Research Management*, Reinhold Book Corporation, New York (1968).
13. Markowitz, H.M., *Portfolio Selection: Efficient Diversification of Investments*, New York, Wiley (1969).
14. Pyle, E.B., et. al., "Scientific Manpower Allocation to New Drug Screening Programs," *Management Science*, Vol. 19, No. 12, August (1973).
15. Raiffa, H., *Decision Analysis: Introduction Lectures on Choices under Uncertainty, Addison-Wesley*, Reading Massachusetts (19680).
16. Schmidlin, P.E., "Standpunkt und Probleme der forchenden pharmazeutischen Industrie," *Arzneim. und Forschung/Drug Research, 28 (11),*, pp 1260—64 (1978).
17. Siegenthaler, Kurt E., "Projects in Pharmaceutical R&D," Project Management Institute, PM15, Toronto, Proceedings (1973).
18. Smith, Robert N., "Fahrplan der klinischen Prufung," Roche Nachrichten, *5*, pg 6 (1978).
19. Souder, William E., "The Validity of Subjective Probability of Success Forecasts by R&D Projects Managers," *IEEE Transactions on Engineering Management*, Vol. *EM-16*, No. 1, February (1969).
20. Spetzler, Carl S., and Stael von Holstein, Carl-Axel, "Probability Encoding in Decision Analysis," *Management Science*, Vol. 22, No. 3, November (1975).
21. Thesing, Jan, *Drug Research in Pharmaceutical Industry: Requisites and Potential*, Medizinisch Pharmazeutische Studiengesellschaft E.V., Frankfurt/Main (1977).
22. Wagner, H., *Principles of Operations Research, with Applications to Managerial Decisions*, Englewood Cliffs, N.J., Prentice-Hall (1969).

21

EVALUATING BASIC RESEARCH STRATEGIES

M. M. Menke
Strategic Decisions Group

J. Gelzer
CIBA-GEIGY

J. P. Pezier
Investment Intelligence Systems Corporation

Reprinted from Long Range Planning,
Volume 14, No. 3, pp. 44-57, 1981.

Evaluating Basic Research Strategies

M. M. Menke, Resource Planning Associates, U.S.A., J. Gelzer, CIBA-GEIGY, Switzerland and J. P. Pezier, SRI International, Europe

A model has been developed to describe the pharmaceutical research and drug innovation process. The model originally served to show the range of results that could arise from research in the field of infectious disease chemotherapy over a sustained period of time. The main goal of the model was to illuminate the risk/reward/time tradeoff faced by management in selecting a particular research strategy. The results obtained with this model were one of several considerations influencing the research strategy decision of CIBA-GEIGY in 1976 to adapt and strengthen its research effort in infectious disease chemotherapy.

This model has so far been applied only for the comparative evaluation of infectious disease chemotherapy research strategies. However, there is no conceptual reason why it should be limited to this field of drug research or indeed to pharmaceutical research at all. This type of analysis should be applicable wherever clear, product-oriented research goals can be identified and defined, where some reasonable financial reward can be attached to achieving these goals, and where an assessment--albeit subjective--can be made of the likelihood of achieving these goals.

The Difficulties of Evaluating Research

Today quantitative analysis is very commonly applied to evaluate marketing,[1-3] production[1-3] and investment[2-4] alternatives as an input to the management decision-making process. Quantitative analysis has also made some headway in the evaluation of specific new products prior to a market introduction decision, a manufacturing capacity decision, or even design decisions in the late stages of development.[1,2,5-7] Until recently, however, there have been rather few reports of successful* applications of quantitative methods to evaluate research and early development decisions. It appears paradoxical that those most involved with the application of the scientific method have been reluctant to extend the use of scientific methods from the laboratory to the manager's office.

Numerous explanations can be postulated as to why management science has made the slowest progress among managers whose responsibility is to manage science. Scientific research in application-oriented industrial laboratories is very complex and by necessity a multi-disciplinary undertaking. Luck, hunch and serendipity often appear to be the principle explanation for discovery and success.[8]

In most areas of research today up to 15 years can be expected to elapse between the initiation of a new research program and the achievement of commercial rewards in the marketplace;[12,13] moreover, the failure rate of 'product' ideas at the research stage is very high. Statistics recently published for pharmaceutical and agrochemical research indicate that 7–10 years are required only for *development*—after research has discovered and identified a promising product candidate.[13,14] Depending upon the 'screening philosophy' prevailing, up to 1000 preparations must be evaluated and tested to reveal one such candidate. Only 1 out of 10 of the 'preclinical' product candidates generally reach the market[13,15-17] implying that about 10,000 research compounds are theoretically required to yield one market product. Which of these 10,000 will succeed is of course unknowable, and yet without such information definitive sales value and production cost assessments are not possible. Thus traditionally

*By successful we mean here actually used as an important input to the management decision-making process. An exception to the above generalization is the sustained experience of Smith Kline and French Laboratories documented in references 9–11 these aspects of the complexity, long-time scale, and uncertainty have provided excuses as to why the quantitative evaluation of research programs and research program segments was either impossible or at best impractical.

Despite their traditional distaste for scientifically structured and formal evaluation procedures of research programs and strategies, research directors are coming under increasing pressure from management to answer such questions as how much research can be afforded, are the efforts focused on the right projects, and will the money invested 'pay off'? This pressure is due partly to the steadily rising cost of R & D,[13,18,19] partly to increased regularity requirements in the chemical and pharmaceutical industry,[17] as well as to incessant pressure on current profit margins from competitive mechanisms and governmental price controls.[20] Therefore not only research expenditure, but also research output (or productivity) have become major management concerns. In response to this concern, a small, but growing, number of organizations have found it possible to yield a substantial insight into the cost/risk/reward relationship in research by using methods of sufficient sensitivity that allow ample scope for incorporating complex subjective judgment, while still providing pertinent information for senior management planning and decision making.

Because of the complexity of the research process, as well as the crucial role of chance in research,[21] the development and use of any model of the research process require the intimate involvement of researchers as well as research managers. This is particularly true for defining sensible research strategy alternatives, specifying the complex mix of resources (human as well as physical) required to pursue these strategies, assessing the corresponding financial costs, defining the goals associated with these strategies (and the innovative character of achieving these goals), and finally assessing the probability of achieving these goals over time. The model of the research process and the associated judgments must also take into account a number of general features of research, including the minimum viable size for particular research efforts, the possibility of diminishing returns as individual research programs grow too 'large',* and the relationship of research productivity not just to level of effort, but also to cumulated effort (i.e. the experience of the research team), the state of the art,* the knowledge base, and the innovative potential of individual researchers. All of these features taken together imply that it is impossible that a model of the research process developed without the active participation of the researcher will reflect reality whatsoever. The interesting experience of Smith Kline and French Laboratories[9–11] indicates that their successful modeling efforts were indeed initiated and developed within their research department.

Given the constraints imposed by the nature of research, as discussed briefly above, it is nevertheless possible to summarize the thinking of experienced researchers about the structure of the research process with a remarkably small number of basic assumptions and to derive useful and intuitively reasonable quantitative results. The three principal assumptions underlying the research model elaborated jointly by SRI/CIBA-GEIGY and presented in this paper are:

(1) The probability of a product concept arrival (i.e. in this case the discovery of a substance or compound destined to become a new drug) in a given time interval is proportional to the length of the interval, for intervals that are small compared to the mean time between arrivals. Furthermore, the numbers of arrivals during non-overlapping time periods are independent;†

(2) The level of effort can influence the arrival rate, but not the peak sales value of the individual arrival (which, however, depends upon the direction of effort);

(3) The arrival rates and peak sales potentials for products discovered in the next 15 years can be estimated from historical experience modified by the current and anticipated industry environment.

The first of these describes the probabilistic nature of the discovery process, which can be quite different for pharmaceutical research than for research into commercial nuclear fusion. The second of these recognizes a fundamental distinction between the probability of research success and the value of that success. A clear separation of these two concepts is vital to any sound evaluation procedure for R & D.[23] The third, however trivial it may seem, reasserts the time-honored principle that past experience must provide the basis for future prediction. Even in areas where past

*We refer here in particular to the lack of visible output (in the form of innovative drugs) from the many large, state-supported pharmaceutical research institutes outside of the U.S., Western Europe and Japan, despite high level of academic qualification of individual researchers. This may be due not only to size, but to related communication and motivation factors of these organizations. In this regard pharmaceutical research differs dramatically from high energy physics, where today pioneering research requires very large research teams. In any case, it is obvious that individual characteristics, such as morale and spirit of the research personnel, must be carefully considered in assessing the probability of research success.

*We refer here to the need to be able in the first place to detect novel, interesting properties of preparations under investigation, through the innovation, design and development of appropriate experimental models and techniques (see references 13 and 22).

†With this assumption, the output of pharmaceutical research can be described by a Poisson model. This model ignores the point that one innovation may provide the lead for a series of related innovations, which in reality creates a clustering of innovation but may not significantly alter the long-run success rate over a 10–20 year period.

Table 1. Example of a therapeutic segment with indication areas and research topics

	Research areas/efficacy goals	Time	Risk	Hierarchy level
4.	Infectious diseases			Therapeutic segment
4.1	Bacterial infections			Indication area
4.11	Antibacterial agent, very broad spectrum	long	low	topic
4.12	Antibacterial agent, active only against gram negative	long	low	topic
4.13	Antibacterial agent, gram positive improved potency	medium	low	topic
4.14	Anti-enterotoxin: active immunization 'vaccine'	long	medium	topic
4.2	Virus infections			Indication area
4.21	Anti-viral agent, broad spectrum (e.g. Myxo)	very long	high	topic
4.22	Anti-hepatitis vaccine	open	high	topic
etc.				
	⋮	⋮	⋮	⋮

experience is severely limited (e.g. an effective curing anti-cancer drug), such experience must provide the starting point for future prediction, although the uncertainty must increase as the experience base shrinks. Equally however, it must be recognized that the environment in which industry operates is not static. Therefore it would be inappropriate to assume that past success rates or sales value distributions would remain routinely valid in the future.*

Definition of a Research Strategy

To evaluate both the probability and the value of research success, a research strategy must aim to achieve goals that are meaningful both to research managers who must assess the likelihood of achieving these goals and to business managers who must assess the value of achieving these goals. The importance of clear meaningful research goals have been stressed in several other publications.[24,25] These references also describe a method for grouping individual research programs (called in reference 25 'topics') goals into so-called Indication Areas and Therapeutic Segments. Research success—as distinct from development success and commercial success—can then be defined as achievement of the research goal, leading in pharmaceutical research to the identification of a clinical drug candidate which in turn calls for the initiation of a drug development project. An example of a therapeutic goal hierarchy that could be used to define an antibiotic research strategy is shown here as Table 1.

In addition to the particular therapeutic goal statement attached to each topic, the likelihood and value of research success are influenced by the technical/scientific approach used to pursue the goal. A very clear and useful way of categorizing the technical approaches used for pharmaceutical research has been discussed by Berde:[26] 'The principle task of the research chemist (and microbiologist)* in the pharmaceutical industry is to prepare new compounds (preparations)* for biological evaluation. These compounds (preparations)* may be grouped in three classes':

- ☆ natural products;
- ☆ derivatives of natural products (including so-called semisynthetics);
- ☆ synthetic chemicals.

'There are three possible ways in which the research chemist and pharmacologist (biologist)* may co-operate in the search for new drugs':

- ☆ random screening;
- ☆ molecular modification (lead optimization using structure/activity relationships);
- ☆ the rational approach (target drug design or natural substitution).

For the dimension of the search procedure we prefer to emphasize the biological aspects by using the phrases Study Structure/Activity Relationships and Natural Substitution instead of Molecular Modification and Rational Approach, since the former terms convey more about how the other members of the research team co-operate with the compound providers.

A pharmaceutical research strategy can be defined by allocating scientific resources to various combinations of these classes of preparations with the corresponding search strategies in pursuit of specified goals. These research efforts may be

*In particular, each company must ascertain whether the success rate for a particular research area is declining (because of already achieved solutions and more demanding regulations or because of changes of attitude in society) and how the marketing performances will evolve over time.

*The terms in parenthesis ( ) are not taken directly from Berde[26] but have been added here in order to generalize and include antibiotic research.

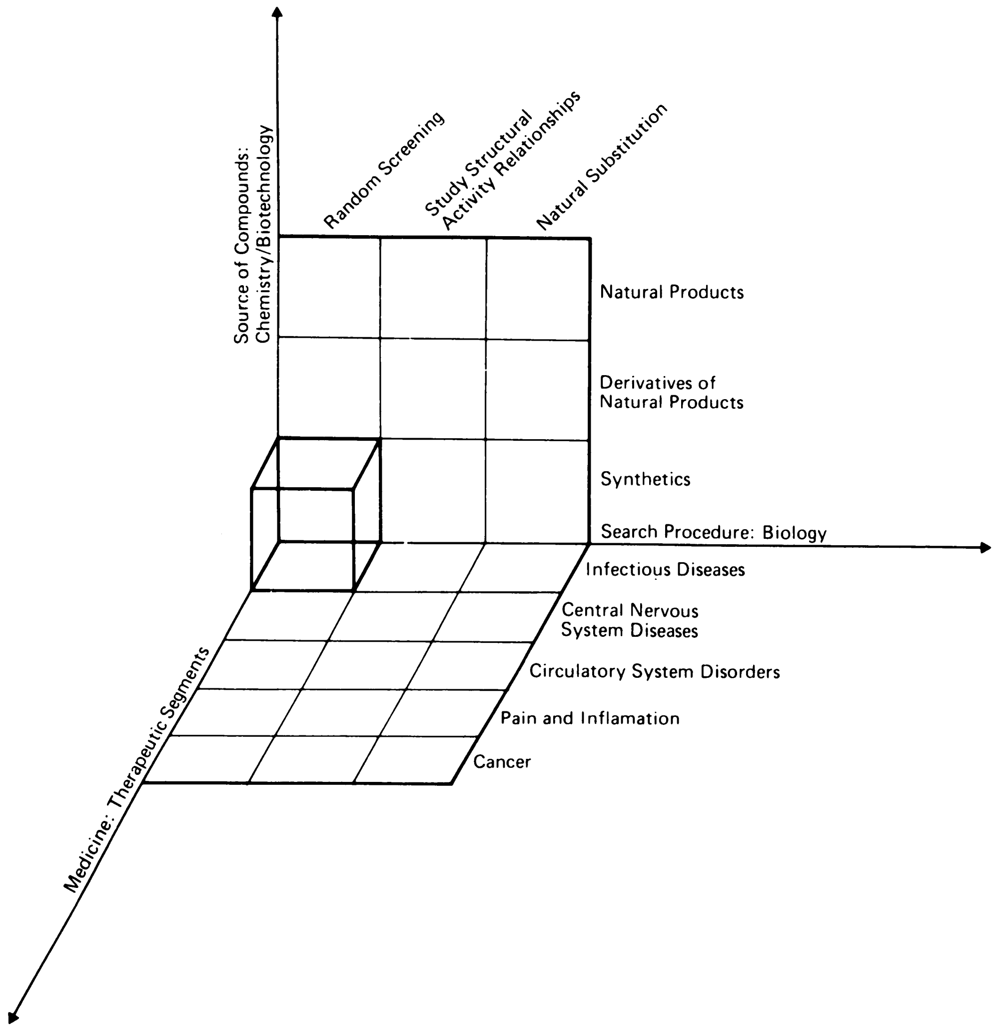

Figure 1. Dimensions of research strategy

pursuing rather broad therapeutic goals, as may be likely in the case of blind random screening, or more specific and very well-defined therapeutic goals, as would be likely in the case of molecular modification/lead optimization or the rational target drug design/natural substitution.

Figure 1 illustrates the interaction of these three main dimensions of a research strategy. The axes of the diagram correspond roughly to the main stages of drug R & D: chemistry and biotechnology/microbiology providing the sources of potential new drugs, biology searching for significant model activity (whether *in vitro* or *in vivo* test models) and finally clinical pharmacology and clinical research establishing the desired therapeutic effectiveness and usefulness in human diseases and disorders. A research strategy then is the allocation of scientific effort to particular cells or volumes in this three-dimensional diagram. The resulting research strategy can then be quantified by assessing an arrival rate (probability of success per unit time) and a value distribution to each cell corresponding to a specific therapeutic goal, source of preparations, and search strategy. In the case of random screening, it may be deleterious to allocate efforts to only one or a few specific therapeutic goals due to narrow scope or the very general nature of a given screen. It may thus be essential to assign likelihoods and values to one or more entire therapeutic areas, according to the construction and concept of the screen.

In Figure 1 we have used the term 'Natural Substitution' instead of 'Rational Approach' or 'Target Drug Design' since our understanding of this approach would stipulate that it tends to work back from the way that nature solves similar problems (e.g. receptor site structure specificity), rather than forward from the case of potentially interesting compounds. A similar interpretation of this approach is provided by Laurence and Black,

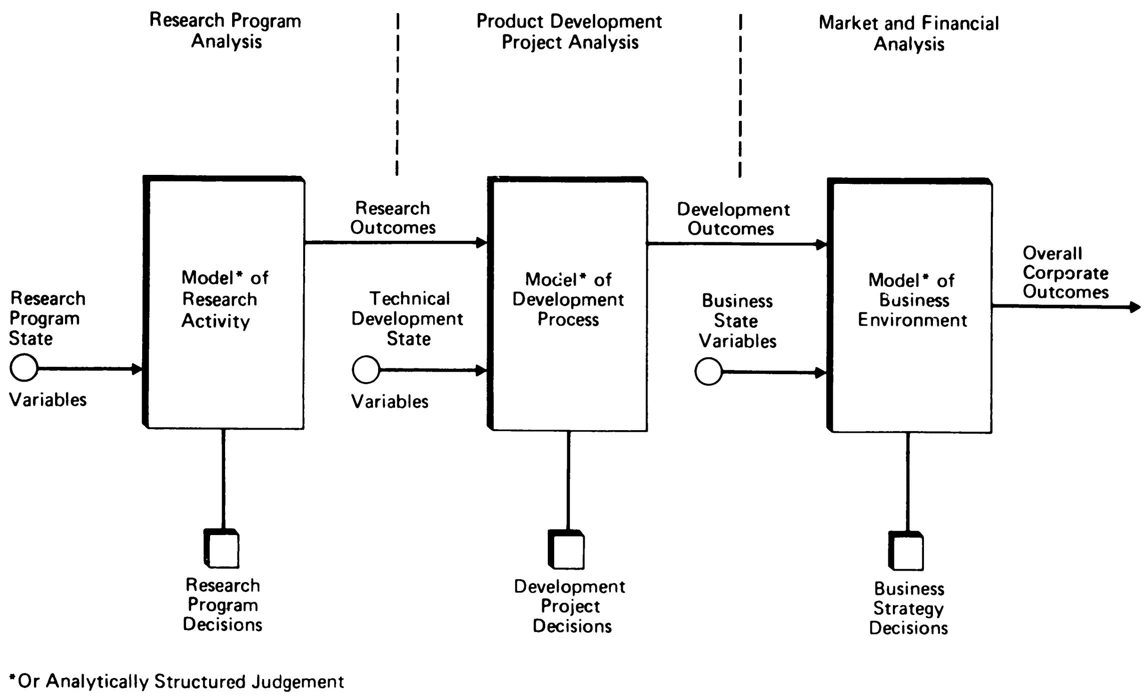

Figure 2. Co-ordinated sequence of research, development and business analysis

who describe this research approach as 'starting with physiological control processes and setting out to find substances which can annul or mimic them'.[27]

Describing and Evaluating Research Strategies

As explained above, research is only the first step of a long and arduous chain of events required to find new drugs and bring them as new products to patients. For purposes of simplicity, one can divide this chain into three stages as shown in Figure 2. Research strategy decisions do not lead directly to commercially useful results, but rather to research results which can potentially be developed into commercially useful results. In the drug industry typical research results would be preparations suitable for drug development, or ideas and principles enabling more productive research or more profitable development. Both types of result can have significant business implications, but the benefits are more readily apparent in the case of a specific compound where the subsequent development requirements and potential business results can be anticipated.[12,16]

The starting point for generating useful business results from research is the allocation of resources to possible research programs. This is called in Figure 2, research program decisions. Table 2 describes three research strategy alternatives that were considered in the infectious disease segment analysis.*

Strategy One is a description of a feasible antibiotic research program, comprising both random screening and more targeted efforts. Strategy Two eliminates random screening altogether, substantially reducing costs and focuses on molecular modification of known interesting molecules. This could be a highly efficient way of carrying out research, but presumably less innovative than continuing random screening since there is little or no likelihood to encounter intramurally substances with really novel structural and biological properties. In economic terms, Strategy Two offers lower costs and lower risks, but at what loss of potential reward? Strategy Three involves both organizational and direction of effort changes designed to obtain enhanced productivity for a modest increase in cost.

In addition to the annual costs, one lower and one higher than 'Strategy One', each strategy can be described in terms of two important productivity measures—the number of substances screened annually and the number of derivatives specially produced and evaluated annually. Reliable averages for these data can be estimated from the staffing and organization of each strategy. Table 2 emphasizes that Strategy Two abandons the expensive

*Numbers are purposely disguised to protect proprietary information.

Table 2. Research strategy descriptions

Research strategy alternative	Total annual Cost ($ millions)	Average number of substances screened annually	Average number of special derivatives produced and evaluated annually	Market MTBA* (years)
One	10·5	6000	260	5·6
Two	7·2	0	400	3·7
Three	12·5	8000	360	3·5

*Mean time between market arrivals for compounds with peak world sales potential greater than $8m/year.

screening activity altogether, resulting in its substantially lower cost. The information in Table 2 provides the basic data which determines the costs, rewards and risks of each research strategy, but the 'best' decision is far from evident given this data alone. At this point no strategy can be either singled out as best or eliminated as clearly inferior.

Up to this point we have described the effort of each strategy in terms of costs and product-oriented accomplishments, namely the substances and derivatives prepared and tested each year. However, as stated above the success rate is very low for potential new drugs to reach the market. Here we must invoke the first principal assumption in the section, 'Difficulties of Evaluating Research', namely that the probability of a product concept arrival in a given (short) time interval is proportional to the length of the interval. In other words, any new substance is as likely as any other to be the 1 in 10,000 that reaches the market. However, a crucial subjective judgment must be made by research management at this point regarding the appropriate future frequency of discovery, as well as the uncertainty of this frequency.

There are two complementary ways to approach this judgment, top-down and bottom-up. Top-down, one can use global reviews of industry or company past research productivity to see how many substances were screened or derivatives produced in order to yield a market product. Such surveys can be found, for example, in reference 28 and were checked against company internal sources. The complementary bottom-up approach is to organize the infectious disease research into indication areas and topics like those shown in Table 1, and then to assess the time required for the first development compound to emerge from each research topic, using the methodology explained in reference 25. This method works well for ongoing research programs and for minor modifications of ongoing programs, but it requires more effort than readily available for a first analysis. Accordingly most of the judgments used in this analysis on the frequency of future research success were made on a top-down basis, modifying past experience to consider current conditions.

Several points should be noted. The success rates for screening and the success rates for molecular modification differ by at least an order of magnitude, but it is logical to expect that research groups actively involved in screening have the advantage of being first to evaluate derivatives with a higher innovative potential (and thus potentially higher sales). The assessment of a success rate (in terms of average number of compounds to be evaluated for each market product) provides a parameter for the Poisson model of the research process that follows from the first basic assumption. This must be done separately for each strategy and for each type of activity (e.g. substance isolatior as opposed to efforts in semi-synthetics) within the strategy. The application of the Poisson model will yield an overall Mean-Time Between Arrivals (MTBA) for each strategy and its sub-categories. This determines the probability of arrivals of new products for each research strategy; the next section discusses how to treat the other big uncertainty, namely the sales value of these product arrivals.

Describing Pharmaceutical Markets and Products

The methodology of the previous section allows one to determine how frequently market products will arrive, on average, in each strategy, but does not address the crucial issue of the monetary value to be attributed to these arrivals. This again is highly uncertain. In this section we shall outline the methodology used and illustrate it with data published in the literature. Figure 3 shows data presented by Schwartzman[29] on the 1972 sales of New Chemical Entities (NCEs) introduced in the United States during the period of 1962–1968. Since a pharmaceutical product generally requires about 5 years after market introduction to reach its peak sales, this graph provides an approximate distribution of peak sales for NCEs. Furthermore, since the U.S. buys about 25 per cent of the world's pharmaceuticals, the value axis has been multiplied by four to show the approximate distribution of peak world sales for these products.

This distribution can be used to estimate the relative chances for products emerging from pharmaceutical research to attain different sales values. In our analysis we used four classes of peak world sales values: $0–8m (too small), $8–24m

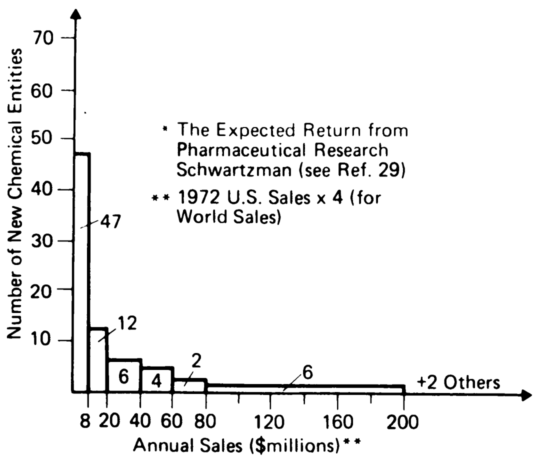

Figure 3. 1972 sales of new chemical entities introduced in U.S. 1962–1968

(small), \$24–80m (medium), and \$80m+ (large). Too small here means below the size or sales volume where a large company might consider it worthwhile to commit a development effort and introduce the product. Thus one sees immediately that more than half of these possible products are *a priori* excluded from consideration. For the other classes, it is evident that a small product is more likely than a medium one, which in turn is more likely than a large.

Using these relative frequencies and other factors particular to a given organization's position in the chemotherapy/antibiotic market, Table 3 shows the resulting MTBAs for each value class.

By adding the reciprocals of the MTBAs for individual value classes one obtains the reciprocal of the MTBA for the combined class. As an example, $\frac{1}{7}+\frac{1}{10}+\frac{1}{22}=\frac{1}{3\cdot5}$, yield the 3·5 year MTBA for Strategy Three products with a presumptive peak sales greater than \$8m/year. The relative frequency of large products in Table 3 has been reduced relative to the Schwartzman data for reasons specific to the organization studied and the nature of the anti-infectious market. Also the frequency of large products in Strategy Two is only $\frac{2}{3}$ of that in Strategy Three, although the frequencies for medium and small discoveries appear the same, reflecting the higher innovative potential of Strategy Three.

It will always be very dangerous to attempt to quantify the results of research in monetary terms. Perhaps the greatest danger is the failure to understand the latent need for innovative products like xerography, sophisticated pocket calculators and drugs like benzodiazepines or cimetidine. Moreover, society may change its value judgment regarding a specific benefit/risk ratio (in industrialized regions an example may be occurring in the field of contraception). Research based firms

Table 3. Mean time between arrivals (MTBA) for products of various peak sales potentials (\$ millions)

Strategy	\$0–8* (years)	\$8–24 (years)	\$24–80 (years)	\$80+ (years)	\$8+ (years)
One	3·3	10	14	35	5·0
Two	2·5	7	10	33	3·7
Three	2·7	7	10	22	3·5

*This class is considered too small to justify market introduction.

may also be over-optimistic at times about their ability to implement new ideas, not only in the technical sphere but also in the market.[20] But it should not be forgotten that drugs found or developed for one particular therapeutic indication frequently yield entries into other important therapeutic applications or areas of innovation.[30] In summary, research intensive organizations must guard against the two important tendencies to overestimate the likelihood of research success while underestimating the value of that success when using specific figures like those in Figure 3 and Table 3.

To relate the peak sales to sales in other years and to the year of discovery, some general information regarding pharmaceutical product life cycles is needed. Studies of the actual sales patterns of a great many pharmaceutical products indicate that the life cycle can be adequately approximated by a 'generic' trapezoidal model shown in Figure 4.

This model is specified by five parameters: the development time (T_D), the peak sales (S), the number of years required to reach peak sales (T_1), the number of years of peak sales (T_2), and the length of the product decline period (T_3). History indicates that T_D, T_1 and T_2 are approximately equal to 5–7 years. T_3 is similar, or longer in the case of very successful products. We call this curve a generic* product life cycle since it is used to represent those as yet undiscovered products that will carry the results of R & D out into the marketplace.

The major uncertainty in describing approximately what the sales will be over time is of course the magnitude of the peak sales, which reflects both the technical innovativeness of the new drug and the commercial skill of the company in marketing the drug. Once a particular compound or product has been identified for a target market, this question can be approached by standard market research and modeling methods.[1] However, even for research areas where specific products have not yet emerged, general studies of the Schwartzman type[29] are likely to be available which provide useful information on

*The use of the word generic here should not be confused with a generic (or non proprietary) drug

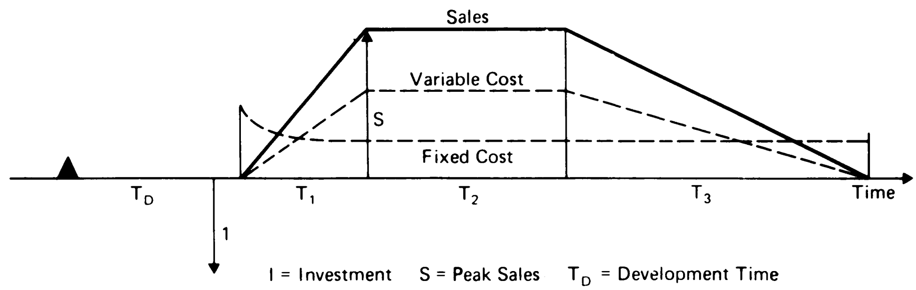

Figure 4. Generic product life cycle model

the value and likelihood of commercial success. This was the approach used to assess the potential rewards of research success in this analysis.

The Financial Consequences of R & D

The question is often raised whether R & D, and research in particular, should be considered an investment, and therefore possibly analyzed as an investment. The arguments against a rigorous financial analysis of R & D include: (1) the long time required to achieve commercial success precludes meaningful quantification, (2) the high risk cannot properly be taken into account, and (3) the ultimate business success sought by product/process directed research is not under the complete control of the researchers, but may be more due to skillful production and marketing. Another common argument is that today's research is paid for with profits from the success of past R & D, and the goal of today's research must therefore be to generate profits to pay for future research (rather than simply to pay back its current cost, with interest).

The authors are convinced that a meaningful quantitative evaluation of the financial return on R & D is possible, but only using methods that specifically address the three concerns above. The decision analysis approach, with its clear separation of the time value of money (i.e. discount rate), the probability of technical success, and the probability of commercial success, allow us to do this.[1,5,23,31] Many proposed methods to evaluate R & D involve the use of so-called 'risk-adjusted' discount rates, which increase the 'risk free' discount rate on the grounds that this should account for the technical risk. However, in most cases the resulting high discount rates combine with the long time scales of R & D to generate negative values for R & D. This is clearly not a viable approach. Meaningful answers are only possible using the same 'risk-free' discount rate that the organization would use for all other resource allocation decisions (e.g. capital investment) and then treating the risk separately and explicitly as discussed below and in References 5 and 23.

A proper methodology must be able to value correctly the results of successful R & D over time and then assess and independently consider the likelihood of such success. The combination of these measures will yield meaningful results, as we shall show. Another issue is that the risk in a major research effort comprising many research programs or in a total R & D portfolio is relatively much less than on an individual project basis. The evaluation methodology must also take this into account. The previous section has explained how to evaluate the financial consequences of a new drug given a peak sales estimate and has displayed in Figure 3 how the MTBAs for different peak sales classes differs. To transform this information into an estimate of the financial consequences of research on a year-by-year basis, a probability model is needed. One could develop year-by-year (or for other timeperiods) an explicit assessment of whether or not small, medium and large products had been discovered in that time period, with probabilities derived from the Poisson Model with appropriate MTBAs. However for five or more time periods, this approach would develop a very large number of possible scenarios, each requiring a separate financial evaluation.

A simpler approach to find the *expected value* of various research strategies is to use the probability model as if a certain fraction of small, medium and large products were discovered every year. This approach will correctly portray the expected value development, but will not show the risk associated with many years without the discovery of or the opportunity associated with early discovery of one (or more) large products. For a 10+ year research effort, this is not a bad approximation for small or medium products whose MTBAs are of the same order of magnitude as the research commitment, but it clearly suppresses the risk associated with whether or not a large sales product will be discovered within the time horizon of the strategies.

With this proviso, a model was developed that transformed the assumptions in Table 3 using the life cycle in Figure 4 (along with further assumptions on investment and expense ratios) into the *expected* sales, profit and cash flow developments

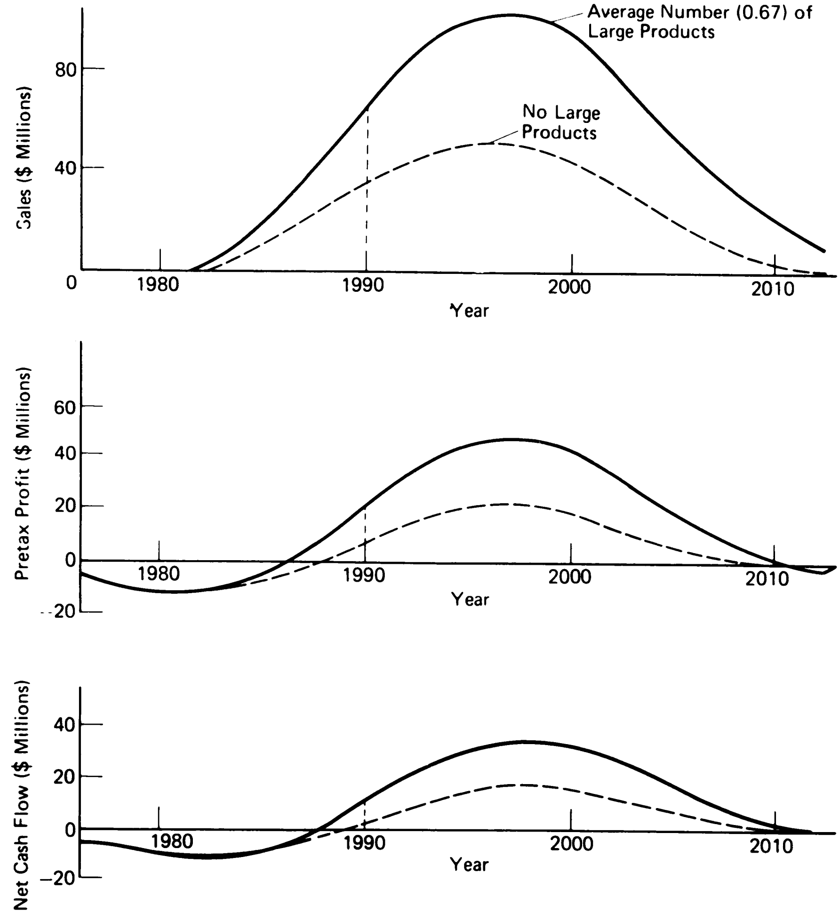

Figure 5. Expected sales, profit and cash flow developments from research model

from a research program of either limited or unlimited duration. The expected results are shown in Figure 5 for a 15 year commitment to research Strategy 3. For an unlimited commitment the curves build up to steady-state levels which they maintain. Expected sales in Figure 5 begin about 7 years after the beginning of the commitment and peak in 1997, about 7 years after the commitment period ends. The solid and dashed lines show the very significant effect of discovering either an average number of large products (0·67 in 15 years) or none during the commitment period.

Using this research model to develop expected cash flows as in Figure 5, one can develop net present value (NPV) measures by discounting the cash flows of each strategy. For many companies, given an appropriate risk-free discount rate, the NPV represents a cash sum equivalent in worth to the entire cash flow. Thus it can be used as a single figure of merit for ranking strategy values under certainty. Table 4 shows how the three strategies of Table 3 compared on discounted profit and net cash flow for three different discount rates. One sees immediately that the financial measures are reasonable, even though the research results stretch out over a very long period (40 years in this case). Figure 6 shows that the NPVs are positive up to discount rates in the range 18–22 per cent, which means that the strategies have internal rates of return in this range.

Many assumptions are needed to reduce the results of long-range research strategies to a financial criterion such as NPV. It is essential to test the sensitivity of the criterion to determine whether or not robust conclusions can be drawn and to identify which sources of uncertainty about assumptions create the greatest uncertainty in the criterion. The most important findings from the research model sensitivity analysis are shown in Table 5, which highlights the overriding importance of the large product discovery to research profitability. While not unexpected, it is important that the sensitivity analysis does demonstrate that the model has intuitively reasonable behavior. It is well known in the drug industry that for many companies a few products provide most of the profits, and Table 5 mirrors this industry characteristic. More surprising might be the *relative* insignificance of a 20 per

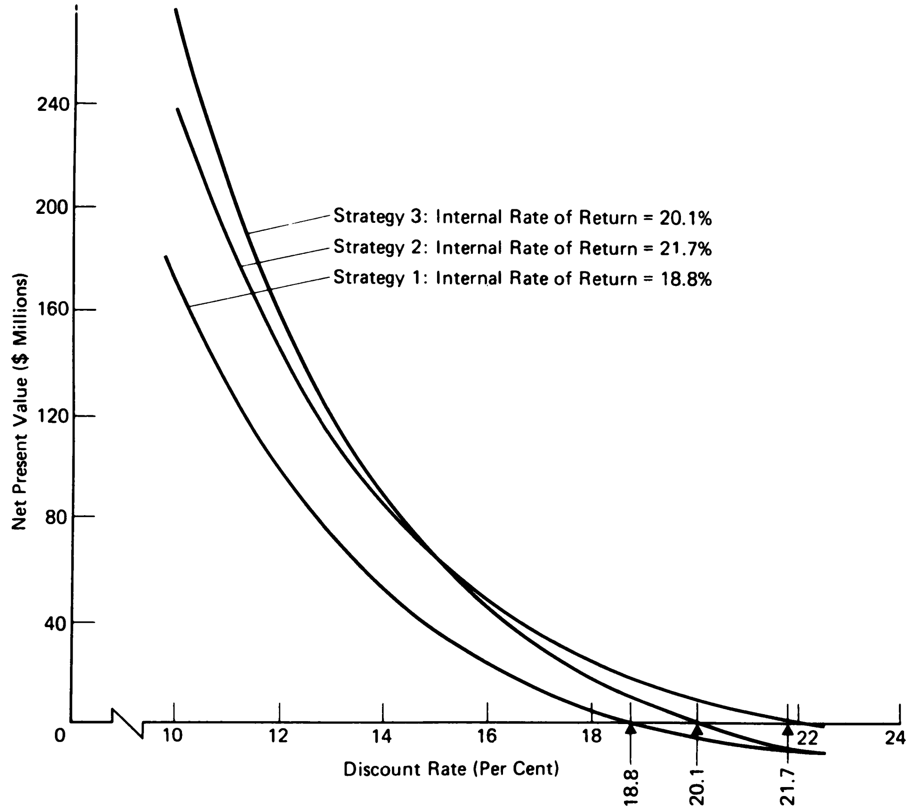

Figure 6. NPV as a function of discount rate for representative cash flows of the three research strategies

Table 4. Discounted research model results

Discount rate (%)	Strategy 1 Pretax profit	Strategy 1 Net cash flow	Strategy 2 Pretax profit	Strategy 2 Net cash flow	Strategy 3 Pretax profit	Strategy 3 Net cash flow
10	290·9	173·4	390·2	236·5	446·4	268·0
15	86·8	36·7	135·6	65·4	144·4	65·0
20	16·6	−5·2	43·0	9·3	37·1	0·4

Table 5. Research model sensitivity analysis

Variable	Variable range	Range of NPV Strategy 1	Range of NPV Strategy 3
Number of large products discovered in 15 years	0–Average	160	214
Success rate in random screening for new chemical entities	1/5000–1/20000	98	152
Success rate in evaluating special derivatives	1/450–1/900	68	95
Fixed and variable product cost	−10%, +10%	58	86
Decline in rate of innovation relative to historical industry averages	0, 20%	50	76

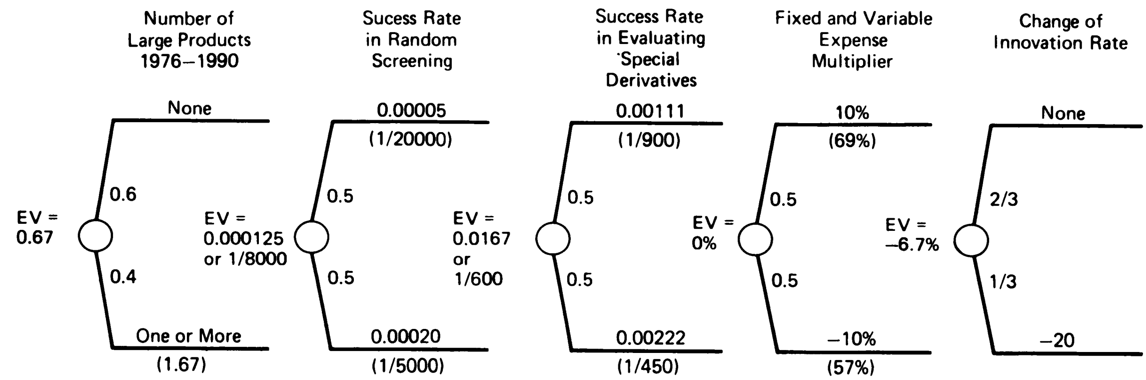

Figure 7. Probability tree for research results

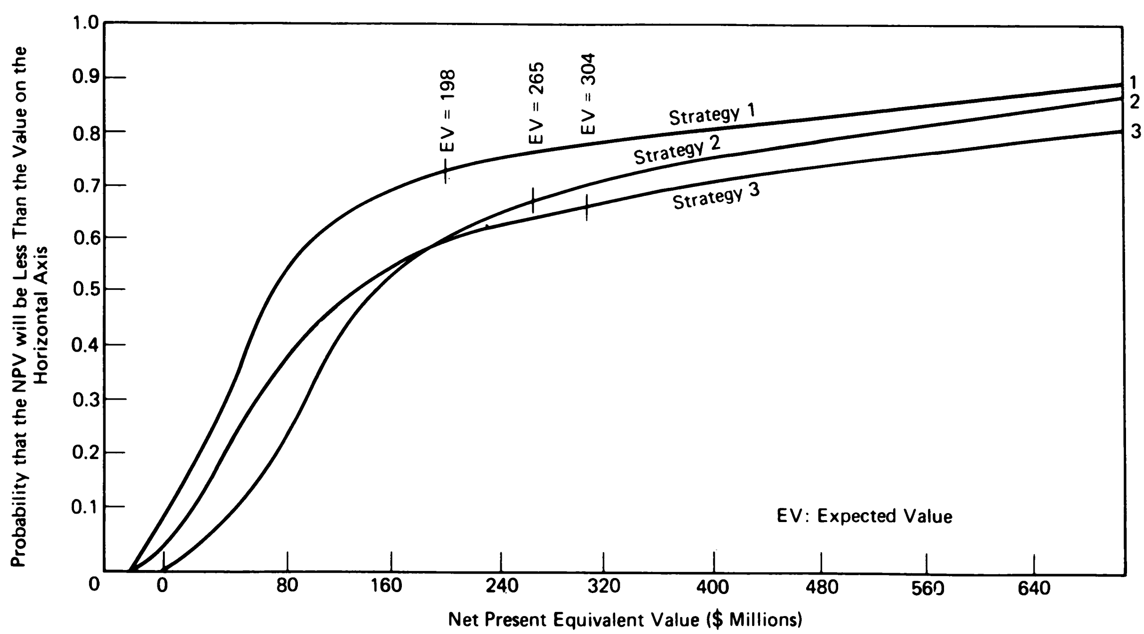

Figure 8. Lotteries for three research strategies (innovation 1976–1990, 10 per cent discount rate)

cent decline in the rate of innovation relative to the so-called 'golden years' of antibiotic discovery in the period 1950–1970. However, the magnitude of uncertainties due to any one of these variables alone in relation to the absolute NPVs at 10 and 15 per cent in Table 4 indicates that while a research program is probably justified, the risks are large enough to warrant a more detailed examination.

Risk and Opportunity Analysis

Here we show a highly simplified representation of the full range of financial consequences that may result from a management decision to commit the company to a research program for an extended period of time. Following the order of importance determined by the sensitivity analysis, Figure 7 shows a probability tree that develops a range of results to which the research program might actually lead. The probability tree of Figure 7 generates 32 ($=2^5$) combinations of assumptions, or scenarios. Each scenario can be interpreted by the research model discussed above to yield appropriate sales, profit and cash flow streams and their corresponding net present values.

To summarize the consequences management faces in pursuing a particular strategy, the probability distribution over a profitability criterion like NPV can be displayed as in Figure 8. This probability distribution (often called a 'profit lottery') displays the range of net present values, from worst to best, that result from the different scenarios generated by the tree of Figure 7.

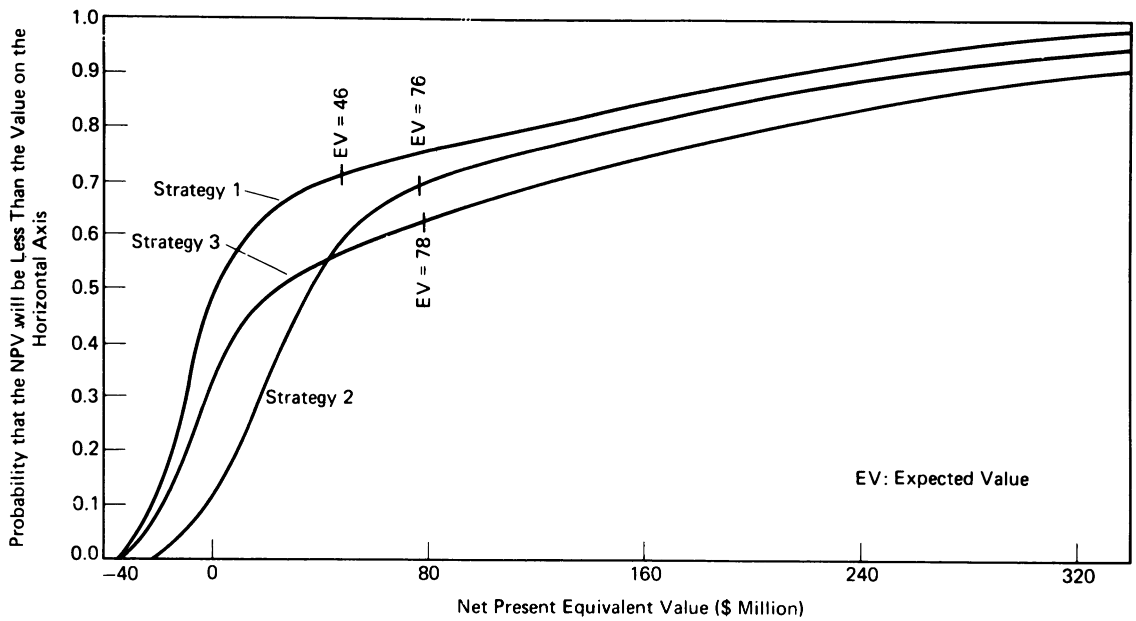

Figure 9. Lotteries for three research strategies (innovation 1976–1990, 15 per cent discount rate)

Figure 8 displays the probability distribution over profitability in the so-called cumulative form, that is any point on one of the curves shows the probability that pursuing the associated research strategy will lead to an NPV less than or equal to the value on the horizontal axis. Conversely, one minus that probability is the probability that pursuing that strategy will lead to an NPV greater than the value on the horizontal axis.

One of the most significant statistics describing the probability distribution, the expected or average value, is indicated. The expected value is the average result the company could expect if it pursued many independent programs whose NPVs had the same probability distribution. Note that since the distributions are not symmetric the expected value differs from the most likely (steepest) value as well as the median (50/50) value. Also the expected value has the distinct advantage that it takes into account the extreme points of the distribution, which the most likely and median do not. Consideration of the extreme possibilities is especially important for a reasonable evaluation of R & D projects.

The distributions in Figure 8 merit further discussion. All three strategies have at least a 90 per cent chance to provide a positive NPV (discounted at 10 per cent) and the worst possible outcome foreseen is about S20m negative NPV with Strategy 1. There is also a 60 per cent chance that the NPV will be less than about S190m (S100 for Strategy 1). On the other hand, if a very big product is discovered, NPVs greater than S500m are entirely possible. The skewed shape of all three curves is a visible manifestation of the riskiness of research. Stated as simply as possible, the odds are that research (over a sufficiently long period) will bring modest but worthwhile gains; however there is an outside chance (here 30–40 per cent) that research during the same period will make a major breakthrough resulting in very high profitability.

Note that the lottery for Strategy 1 is entirely above and to the left of Strategies 2 and 3. This means that for any level whatsoever of NPV, Strategies 2 and 3 offer a higher probability to exceed it. Although we are not guaranteed a better result by taking Strategies 2 and 3, they definitely offer a higher probability of achieving a better result. This effect is called stochastic dominance and it implies that whatever management's feeling is with regard to risk, Strategy 1 is worse than either 2 or 3.

The choice between Strategies 2 and 3 is not so obvious. Strategy 3 offers a higher expected value, but has a significant probability of yielding a worse result than Strategy 2. The higher expected value comes from the higher probability that Strategy 3 offers to reach an NPV in excess of S190m. This is a reflection of the continued screening research in Strategy 3 as opposed to the cheaper but less innovative restriction to molecular modification in Strategy 2. The choice between Strategies 2 and 3 confronts management with a risk/reward trade-off: are they willing to pursue a riskier strategy (3) which on average leads to a higher reward? There is no correct answer to such questions, but the theory or risk preference can be used to offer some insights as to what management should prefer given various levels of risk tolerance. However, for a company like CIBA-GEIGY which is actively researching a number of therapeutic areas, as well as several other fields of chemistry, the expected value should be a good guide to the relative worth of a single therapeutic area.

Figure 9 illustrates how management's time preference (expressed by the discount rates of 10 and 15 per cent) interacts with the riskiness of the strategies. At the higher discount rate of 15 per cent the future benefits of research discoveries are reduced, which implies that Strategies 2 and 3 have virtually identical *expected* NPVs. In this instance the 'best' solution should be clear, since Strategy 1 is still stochastically dominated and Strategy 2 is as valuable as Strategy 3 on average but less risky. Thus if 15 per cent were the appropriate risk-free discount rate, management should pursue Strategy 2, unless they are essentially risk neutral.

Conclusions

The results reported, while still preliminary, did play a role in CIBA-GEIGY's decision on how to proceed and how to select among various strategic alternatives in the research area of anti-infectious chemotherapeutics. The numerical results, even had they been derived with much deeper study than was possible in this project, are not to be slavishly followed. The following results, however, considerably increased managements confidence in their decision:

- ☆ In spite of the long time scales (15+25=40 years), the economics of research are sensible in the discounted cash-flow sense. The internal rates of return on the research expenditures fall in the range of 15–25 per cent, given the pricing assumptions used in this analysis.
- ☆ Strategy 1 is stochastically dominated by Strategies 2 and 3 and therefore should never be preferred.
- ☆ The sum of cash flows for Strategy 3, ('increased effort' including screening) is larger than for Strategy 2 ('lead optimization'), however the pay back period is longer with Strategy 3. At 15 per cent discount rate the expected net present values for Strategies 2 and 3 are approximately equal.
- ☆ Strategy 3 is more risky than Strategy 2 because it depends more critically on the discovery of 'big selling' products. The choice between Strategies 2 and 3 depends therefore essentially on management policy regarding time and risk preference.

The insight into research profitability gained from this quantitative evaluation of three alternative research strategies can and did help design improved strategies beyond those analyzed.

The research modeling effort passed the primary test of reasonableness by analytically capturing intuitive reasonable characteristics of the research process. As could have been predicted *a priori* (but was not), the most sensitive variable in the research problem turned out to be the number of 'big-selling' products that will emerge from the 15 year research program (the probabilities of discovering 0, 1, 2 . . . such products are 0·51, 0·34, 0·11 . . . for Strategy 3). Having been able to capture so many aspects of the research strategy decision, we feel that the valuation approach presented merits further application, in particular for the relative comparison of alternative research areas; not only in pharmaceutical research but in all business areas where the future growth and renewal of the business depend upon a flow of product concepts, ideas and discoveries from the research process.

References

(1) R. Abt *et al.*, The dangerous quest for certainty in market forecasting, *Long Range Planning*, **12**, 52–62 (1979).

(2) W. G. Byrnes and B. K. Chesterton, *Decisions, Strategies and New Ventures*, George Allen and Unwin, London (1973).

(3) Joseph W. Newman, *Management Applications of Decision Theory*, Harper & Row, New York/San Francisco/London (1971).

(4) R. F. Egger and M. M. Menke, Decision analysis at Hoffmann-La Roche, to be published in *Planning Review*.

(5) H. U. Balthasar, R. A. A. Boschi and M. M. Menke, Calling the shots in R & D, *Harvard Business Review*, **56**, 151–160 (1978).

(6) R. G. Hudson, J. C. Chambers and R. G. Johnston, New product planning decisions under uncertainty, *Interfaces Special Practice Issue*, p. 82 (1977).

(7) Th. Benelmans, Strategic planning for research and development, *Planning Review*, **12**, 33–44 (1979).

(8) Karl H. Beyer, *Discovery, Development and Delivery of New Drugs*, Spectrum Publications, New York/London (1978).

(9) E. B. Pyle *et al.*, Scientific manpower allocation to new drug screening programs, *Management Science*, **19** (1973).

(10) Michael A. Cochran *et al.*, Investment model for R & D project evaluation and selection, *IEEE Transactions in Engineering Management*, **EM-18** (3), 89–100 August (1971).

(11) A. Douglas Bender *et al.*, Simulation of R & D investment strategies, *Omega*, **4**, 67 (1976).

(12) W. M. Wardell, The drug lag revisited: comparison by therapeutic area of patterns of drugs marketed in the United States and Great Britain from 1972 to 1976, *Clinical Pharmacology and Therapeutics*, **24**, 499–524 (1978).

(13) E. F. Jenny, Forschung ans des Sicht des Forchers, *Die Pharma Ind.*, **41**, 15–27 (1979).

(14) J. R. Huber and Jeanette Zwick-Holm, *The Swiss Chemical Industry*, Information Service of the Swiss Society of the Chemical Industries, Zurich (1977).

(15) H. J. Bein, E. Eichenberger, G. Jung, C. Popzer, W. Stamm and M. Taeschler, Comparison of experimental and clinical findings; an analysis of clinical trial preparations in the correlation of adverse effects in man, S. B. de Baker, Amsterdam (Eds.) *Experta Medica* 1971, pp 179–183 (Experta Medica International Congress Series No. 220).

(16) J. Thesing, Trendumkehr in der Arzneimittelforschung, *Naturwissenschaften*, **64**, 601–605 (1977).

(17) J. Gelzer, Governmental toxicology regulations: an encumbrance to drug research, *Archives of Toxicology*, **43**, 19–26 (1979).

(18) R. W. Hansen, The pharmaceutical development process: estimates of development costs and times and effects of proposed regulatory changes, *Issues in Pharmaceutical*

Economics, R. I. Chien (Ed.) Lexington Books, D. C. Heath & Company, pp. 151–181 (1979).

(19) Editorial: Bleak prospect for drug research, *The Lancet*, p. 1335, 24 and 31 December (1977).

(20) B. G. James, A look into the future of the pharmaceutical industry, *MM & M*, pp. 13–21, July (1978).

(21) Jan Thesing, *Drug Research in the Pharmaceutical Industry: Requisites and Potential*, Medizinisch Pharmazeutische Studiengesellschaft e.V., Frankfurt/Main (1977).

(22) J. Gelzer, Gedanken zur Zukunft der Pharmazeutischen Forschung, *Chemische Rundschau*, **31**, 5 (1978).

(23) J. Matheson, R & D decision analysis, SRI International, preprint, Menlo Park, CA (1977).

(24) H. U. Balthasar and St. Gutzwiller, Steady state and portfolio concept in R & D management, *R & D Management*, **5**, 201–207 (1975).

(25) R. A. A. Boschi, H. U. Balthasar and M. M. Menke, Quantifying and forecasting basic research success, *Research Management*, **22**, 14–21 (1979).

(26) B. Berde, Industrial research in the quest for new medicines, *Clinical Experimental Pharmacology and Physiology*, **1**, 183–195 (1974).

(27) D. R. Laurence and J. W. Black, *The Medicine You Take*, Fontana, New York/London (1978).

(28) J. Berdy, Recent developments of antibiotic research and classification of antibiotics according to chemical structure, *Advances in Applied Microbiology*, **18** (1974).

(29) D. Schwartzmann, *The Expected Return from Pharmaceutical Research*, American Enterprise Institute, Washington, DC (1975).

(30) R. G. Denkewalter and M. Tishler, Drug research—whence and whither, *Progress in Drug Research*, **10**, 11–31 (1966).

(31) Michael M. Menke, Strategic planning in an age of uncertainty, *Long Range Planning*, **12**, 27–34 (1979).

22

SOCIAL DECISION ANALYSIS

Ronald A. Howard
Department of Engineering-Economic Systems
Stanford University

Reprinted with permission from Proceedings of the IEEE, Vol. 63, No. 3, March 1975, pp. 359-371.

Social Decision Analysis

RONALD A. HOWARD, SENIOR MEMBER, IEEE

Abstract—How can a diverse free society find decision mechanisms that are logical, efficient, and timely? This is a problem that has challenged man at least from the time of Plato's *Republic*. We see today factions of society arguing over alternatives rather than over values or probabilities. Adversary proceedings encourage people to advocate extremes rather than a careful balance of several considerations.

Decision analysis, a logical procedure for balancing the many uncertain, complex, and dynamic factors that characterize a decision, offers promise of a new and valuable procedure for social decisions. The decision analyst creates an extrapersonal explicit model of the decision under consideration. Information on possible alternatives, uncertainties, relationships, or preferences can come from different groups and still be represented within the same decision model, with the implications for the decision apparent to all. One can imagine a society where decision making has become decentralized, where distinct bodies are responsible for creating social alternatives, assessing the probabilities of various outcomes for each alternative, and setting the preferences of society. Once the alternatives, information, and preferences were established, society would make the decision using only the principles of logic.

Applications to automotive pollution, hurricane seeding, and nuclear safety demonstrate the approach.

I. Introduction

It MUST come as a surprise to most people that there is a logical way for an individual to make a decision. The essence of uncertainty seems to imply that any attempt to quantify the decision-making process must be doomed. Yet the developments of the last 10 or 20 years have shown that it is not only possible to define a good decision, but also that individual and corporate decision makers can implement a logical decision-making process in the important decisions they must face.

The question we propose to examine here is whether this process can be successfully adapted to the social decisions that we seem to encounter in increasing numbers as our society becomes more interdependent. We shall first describe the logical decision-making process we call decision analysis and show how it has been used to aid the individual decision-maker. Then we shall consider the important challenges to this procedure posed by social decisions. We shall propose a procedure for such decisions, and illustrate its use in selected social problems.

II. Decision Analysis

Good decision making is just common sense. The difficulty was that we did not know what common sense meant until we parsed the decision-making process. More complete descriptions are given elsewhere [1]–[3]; here we shall only summarize. A good decision is the logical consequence of what you can do, what you know, and what you want. "What you can do" are the alternatives available to you. Finding alternatives is the most creative part of the decision process. "What you know" means the knowledge of relationship and of magnitude, the information you bring to the decision-

Manuscript received September 19, 1974. This work was supported by the National Science Foundation under Grant GK-36491.

The author is with the Department of Engineering-Economic Systems, Stanford University, Stanford, Calif. 94305.

making process. Finally, "what you want" refers to the preference you have for the various consequences of the decision, consequences that may be distributed over time or uncertain.

The concepts we have just stated are accepted by all who work on the problem of making better decisions; however, interpretations of the concepts have varied. One group of interpreters we might call the decision theorists. The decision theorists proceed by capturing the personalistic view of the decision maker. In principle, this means having the decision maker assign a joint probability distribution on all the consequences of each alternative, and then specify a multiattribute utility function on those consequences. The best alternative is then the one with the highest expected utility. While this procedure seems simple, it rapidly exceeds the assessment capability of decision makers on any but the simplest problems.

Another group of interpreters whom we shall call the decision analysts also believe in the importance of incorporating judgment and preference. However, they spend more of their effort in constructing an extrapersonal explicit model of the decision under consideration. This model captures the important relationships of the decision problem and is open to the inspection of all who might contribute to the decision. The extrapersonalization of the decision model allows information to be collected from many disparate points of the organization: lawyers have a place to assign legal probabilities; metallurgists a place to assign distributions on material strength. While the decision maker is ultimately responsible for preferences, he can focus his attention only on the parts of the preference structure that cannot readily be delegated to others.

The creation of extrapersonal models by decision analysts was a natural consequence of applying decision analysis to the major decision problems of corporations and government agencies. Explicit models were a practical tool developed because of their usefulness in analysis and implementation of the decision. However, now that they exist, we begin to see how an entirely new structure for social decision could be built upon them. The explicit decision model provides the crucial step in dividing responsibility for the creative, informational, and preferential aspects of a decision among as many individuals or groups as required.

III. Social Decision

Now we turn to the problem of social decision and consider how decision analysis might be helpful. Most social decisions are presently made as the result of a legislative process. The legislators, primarily trained as lawyers, pursue their work through hearings intended to clarify the nature of the decision. However, these hearings often become adversary processes since the proponents of various alternatives are motivated mainly by a desire to see their favorite alternative selected rather than by a desire to illuminate the issue. Possibilities become confused with likely consequences. Bad outcomes are extensively discussed by opponents with the purpose of making their probability appear higher. From the SST to nuclear power plant siting, we have all seen and been affected by the results.

This system requires legislators to create alternatives, examine their implications, and value the outcomes for society. We can now ask whether we wish all these functions to reside in one group and, further, whether this group is properly trained to perform them. What might have been feasible in a simpler environment may no longer be adequate. The same forces that required corporations to seek more effective decision procedures are acting on society as a whole.

IV. Social Decision Analysis

Let us now assume the existence of a large supply of decision analysts capable of structuring social decision problems to the point where an explicit decision model was available. It would then be possible to constitute various bodies charged with the responsibility of providing the requisite inputs to this model.

One body of people would be responsible for the creation of new alternatives for solving society's problems. Anyone could suggest alternatives to this body for consideration. The results of its deliberations would be a set of certified alternatives for dealing with any societal problem.

A second body would be responsible for assessing the probabilities of various outcomes for each alternative. This would require capturing structural knowledge available on the alternatives and assessing probabilities on uncertainties. The work of this body would be heavily aided by decision analysts and subject-matter experts. The result would be a certified probability distribution on the outcomes of any alternative selected.

A third body would be concerned with characterizing the preferences of society. This body would be heavily influenced by the desires of all citizens. Voting procedures could be developed for determining citizen preference. However, the values of society should change slowly: major value changes would require the same care currently devoted to a constitutional amendment.

The values that would be set would be, for example, the value of a life, of health, of recreation, of cultural attainment. Outcomes extending over many years would require setting a social time preference, which would be based strongly on intergenerational considerations. While there are good arguments that society as a whole should be risk indifferent, other risk attitudes could be specified. The net product of this third body would then be a certified set of preferences to be used in evaluating outcomes.

Once the certified alternatives, information, and preferences were established, society would make the decision using only the principles of logic. If the resulting decision appeared incorrect, it could only be changed by a change in the certified inputs.

It would be possible to have different certifying bodies for different areas of social decision making. However, the fundamental nature of the deliberation of the preference body make it likely that only one such body should be constituted.

We shall now consider examples of situations where several interrelated factors, uncertainties, possibly high risks, and social concern can combine to produce a social problem worthy of the type of approach we have been discussing. We shall present these analyses in very brief summary form. Readers who desire a detailed examination of information sources and modeling choices should consult the references.

V. Automotive Emissions—The Problem of Externalities

Automotive emissions provide an excellent example of what economists call externalities.

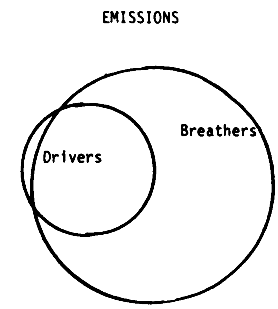

Fig. 1. The single basin model.

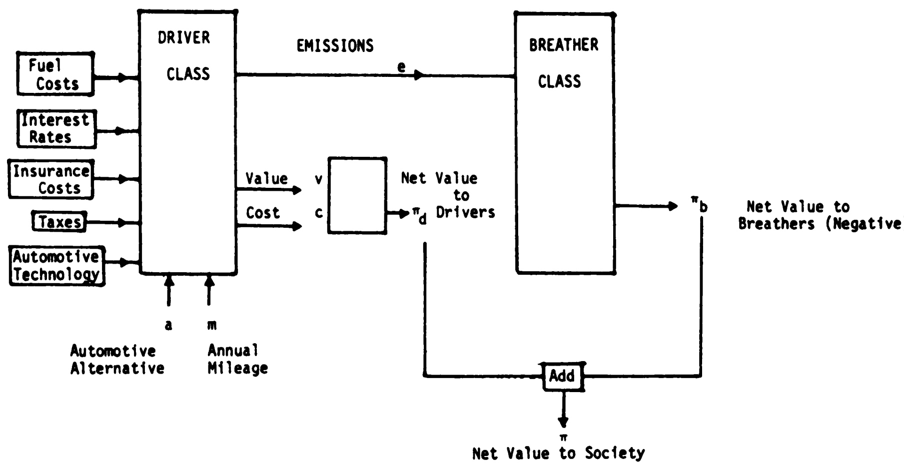

Fig. 2. The driver-breather model.

A. Externalities

An externality arises when the actions of one individual have effects on other individuals for which the other individuals are not economically charged or compensated. Externalities can be favorable, as when a world-famous violinist practices in his backyard on a Sunday afternoon and his playing is enjoyed by his neighbors, or detrimental if the same violinist practices at three o'clock in the morning.

B. Drivers and Breathers

Drivers of automotive vehicles create an externality when the emissions from their vehicles are noxious either directly or indirectly to the other individuals in the society. We shall restrict our attention to a single topographical basin where virtually all the automotive emissions are caused by people who live in the same region. Now we consider two classes, the "drivers" who create the emissions and the "breathers" who are affected by them. The term "breathers" may emphasize the human physiological effects of the emissions, but we also include individuals who are additionally concerned by aesthetic and more indirect effects such as the influence of emissions on plants; thus breathers are virtually all the human inhabitants of the region.

Fig. 1 shows the essence of the single-basin model. Drivers create emissions affecting breathers. Note that most of the drivers are breathers themselves; we are characterizing roles rather than individuals. However, typically there are many nondriving breathers, including children and the very elderly. There are also drivers who are not, or only partially, breathers, for example, those who live outside the region but who drive through it.

C. The Driver-Breather Model

Our first task is to formulate a model of externality that will allow us to see the effects of policy changes. Such a model for the case of emissions appears in Fig. 2. We imagine each member of the driver class to be faced by an economic environment described by fuel costs, insurance costs, interest costs, taxes, and automotive technology. In the face of this environment, the driver must select an automotive alternative a and an annual mileage m. We shall consider these decisions to be made on an annual basis.

Automotive Alternatives: The automotive alternative is a specification of which automotive system the driver will use for the coming year. It includes the complete spectrum of automotive conveyances each driver might buy or rent, physical modifications to those conveyances, and maintenance policies for the conveyances. For example, one alternative might be to buy an old car and do as little maintenance as possible. Another might be to buy a new no-lead car using a catalytic converter.

Cost: When the automotive alternative is selected and the annual mileage m is specified, the annual cost of operation c can be computed. Capital costs of purchase are annualized by simple or sophisticated methods.

Emissions: Associated with the (m, a) choice will be an emissions level e. While there is controversy over the consequences of various types of emissions, we shall simplify that

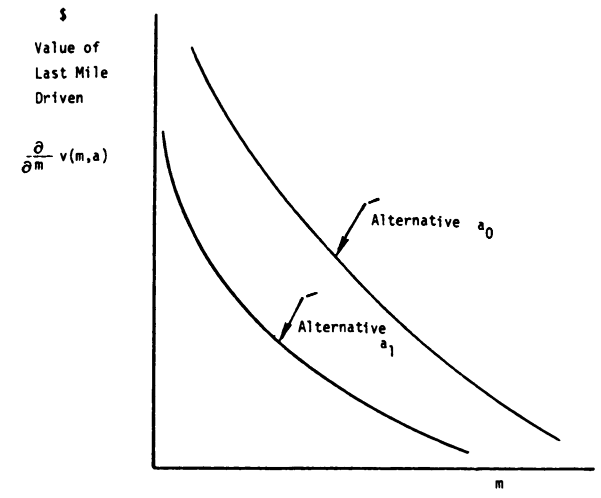

Fig. 3. The driver's marginal value of additional annual miles driven.

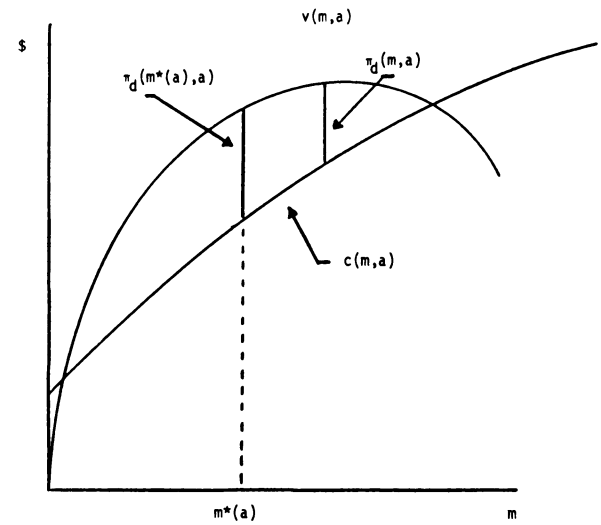

Fig. 4. Value and cost of mileage for a given alternative a.

question here by measuring the total weight of all pollutants characterized as harmful. More detailed treatment is, of course, possible.

The Driver's Decision: How does a driver decide which (m, a) choice to make? We assume that he selects the choice that maximizes the difference between his value for the choice $v(m,a)$ and the cost of the choice $c(m,a)$. This difference we call the net value to the driver $\pi_d(m,a)$.

Marginal Value of Annual Mileage: To determine $v(m,a)$ we begin with the driver's marginal value of additional annual mileage driven shown in Fig. 3. In theory, we construct this value for each automotive alternative by asking the driver how much he would pay to be able to drive 1 mi during the year, then another mile, and so forth. Presumably, the first miles are very valuable indeed since they would be used for emergency trips. When a large number of miles was already available, the values of succeeding miles should become very small. Thus we expect the marginal value of additional miles to decrease continually.

The marginal value curve depends on the automotive alternative primarily through what we could summarize as "performance." For example, the marginal value of high annual mileage in a standard car could be quite high because the driver contemplates using it for vacations. A small electric car could have a low marginal value under these conditions if it were not suitable for vacation travel. Similarly, a car with very poor acceleration might make almost all driving miles less valuable and consequently lower the curve.

Of course, it is practically almost impossible to develop the information of Fig. 3 for every driver and every alternative. However, once the concept is clear, only a few inputs may be necessary to determine the level of detail necessary for the social decision to be made.

The Net Value Calculation: Now we are in a position to determine the driver's decision. First we integrate (or sum) the marginal value curves of Fig. 3 with respect to m to obtain the driver's total value of driving m annual miles using alternative a, $v(m,a)$. This value is shown for a given alternative a in Fig. 4. The total value of driving m miles always increases with m, but with a steadily decreasing slope.

Also shown in Fig. 4 is the cost of driving m annual miles with alternative a, $c(m,a)$. For most alternatives, $c(m,a)$ will have a nonzero intercept $c(0,a)$ caused by the fixed costs of ownership such as capital costs, property and license taxes, and insurance. The annual cost $c(m,a)$ should be almost linear in m, but the figure allows for a lower marginal cost at high mileages.

The difference between the value and cost curves is the net value to the driver $\pi_d(m,a)$. For each a he should select the mileage $m^*(a)$ that maximizes this value,

$$m^*(a) = \max_m{}^{-1}\ \pi_d(m,a).$$

Once he knows the mileage he would drive given each automotive alternative, he would select the alternative a^* with the highest net value at its optimum mileage,

$$a^* = \max_a{}^{-1}\ \pi_d(m^*(a),a).$$

Some of the automotive alternatives may have negative net values, others positive. However, our range of alternatives can be quite broad from "call a taxi" to "rent a sports car." We would expect almost every member of the driver class to find an alternative of positive net value. The net value to each driver is then

$$\pi_d^* = \pi_d(m^*(a^*),a^*).$$

When this net value is summed over all drivers, we obtain the total net value of the driver class π_d.

The Breather's Consequences: The drivers as a result of their decisions have collectively produced emissions for the breathers. Since presumably no breather likes the emissions, we would expect that we would have to pay each breather a sum of money to achieve a situation where he was as happy to have the emissions and the money as he would be to have completely clean air. If we determine this sum for each level of emissions and sum the results over all breathers, we might obtain the net value to breathers curve of Fig. 5. The net value $\pi_b(e)$ is, of course, negative. We have chosen to make the curve almost linear, but with a slightly increasing slope to

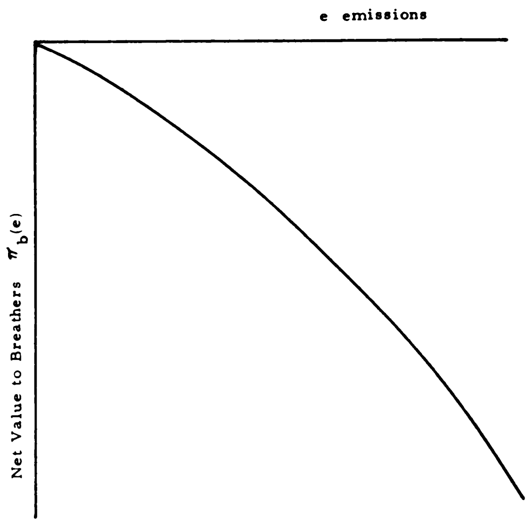

Fig. 5. Net value to breathers.

indicate that high levels of emissions may be considered to be even more undesirable than the linear model would specify.

Net Social Value: We can now evaluate any social alternative A_i for the treatment of the emission externality problem using the model of Fig. 1. We determine how drivers would make their decisions under this alternative and find the associated net drivers' value $\pi_d(A_i)$ by summing the optimum net values of each driver. Of course, it is possible that certain alternatives would cause nondrivers to become drivers or vice versa. If this effect is important, it can be treated by considering the driver class to be the union of driver classes for each alternative.

Then we would find the emissions $e(A_i)$ generated by the driver class under this alternative. The net value to breathers of this level of emissions is then $\pi_b(e(A_i)) = \pi_b(A_i)$, as computed from the net value to breathers curve. Finally the net value to society of alternative A_i, $\pi(A_i)$, is given by

$$\pi(A_i) = \pi_d(A_i) + \pi_b(A_i).$$

The alternative with the highest net value to society would be the best alternative for implementation, if alternative A_i did not in itself have associated administrative costs.

Administrative Costs: Most social alternatives A_i will, in fact, impose additional costs to the society as a result of their implementation. If these costs are $k(A_i)$, then the net social value will be

$$\pi(A_i) = \pi_d(A_i) + \pi_b(A_i) - k(A_i).$$

Change in Net Social Value: We shall now find the best social alternative by determining which alternative has the largest net social value including implementation costs. However, it serves as well to designate some particular alternative as a base alternative and then to determine how the net social value will change if any other alternative is used. Thus if we use A_0 to designate the base alternative, conveniently a continuation of present policy, then $\Delta\pi(A_i)$, the increase in social value caused by switching to social alternative A_i, is given by the following

$$\Delta\pi(A_i) = \pi(A_i) - \pi(A_0).$$

The best social alternative is then the one with the largest $\Delta\pi(A_i)$. Finally, using a parallel notation, we can write $\Delta\pi(A_i)$ in terms of the components of social value as

$$\Delta\pi(A_i) = \Delta\pi_d(A_i) + \Delta\pi_b(A_i) - \Delta k(A_i).$$

The advantage of this formulation is that we need deal only with the changes in driver net value, breather net value, and implementation cost when considering a new alternative rather than measuring the magnitude of any of these quantities. This approach simplifies considerably the input questions we discussed previously.

D. Social Alternatives

Now let us consider various social alternatives that could be employed in treating the emission externality.

Creating an Emission Market: Perhaps the simplest conceptual approach to the emission externality is to create a real emission market. Consider an agency that would charge drivers for the emissions they create and then use the money to compensate the breathers. If there were some way to discriminate among drivers so that those to whom the rights of emission were most valuable paid the highest tax per unit or if there were a way to discriminate among breathers so that those most affected by emissions received the highest compensation, then it could be possible for the revenue received by the agency in the form of taxes to be different from the compensation costs it incurs. It is likely that the agency could run at a profit even after administration costs were paid.

A Practical Pollution Market: However, for simplicity and practicality, we shall assume that the market functions by charging drivers r for each unit of emission and then compensating the breathers equally according to the total amount of emissions produced. For example, if each breather would pay \$$b$ for clean air, if there were n_b breathers, and if the total emissions were e kilograms, then the charge per kilogram (r) for emissions would be $r = bn_b/e$.

The Proper Signal: The importance of this approach is that it would reflect to the creators of the externality the costs that they are imposing on those they affect. Instituting such a charge would not drive emissions to zero, but would assure that each unit of emissions created by drivers was worth at least as much to them as the cost they were imposing on the breathers. Thus a proper "signal" is created to illustrate to those who emit just what is in the social interest. Note that this effect is preserved even if the compensation is not actually paid directly to the breathers but simply to the general funds of the community. Of course, if the emission level falls, the amount b the breathers are willing to pay for the clean air will probably fall even more so that we would expect the emission charge to fall as the air contains less emissions.

System Costs: We have not yet faced up to the question of how to implement even this simplified market. What would be required is a device to measure emissions directly, for example, a device that could be attached to the tailpipe to experience chemical changes as emissions are created. Perhaps the device could change color when a certain quantity of emissions had passed, thus requiring the owner to buy a new one at a service station. Perhaps he would have to exchange it periodically for a new one and send the old one to a center to determine the emissions he had created and consequently the tax he owed. The point is that such a method would make it the owner's responsibility to determine what kind of car he drove, how he drove it, and how he maintained it. "The more you emit, the more you pay."

TABLE I
DRIVER CHARACTERISTICS

MILEAGE CLASS (THOUSAND MILES PER YEAR)	%	VARIABLE COST PER MILE (DOLLARS)
> 30	.03	.040
20 - 29	.06	.048
15 - 20	.12	.056
10 - 15	.25	.064
5 - 10	.34	.072
< 5	.20	.080

TABLE II
ANTIPOLLUTION DEVICE CHARACTERISTICS

Abbreviation	Descriptions	Exhaust Emission Rate (g/mi) HC	NO_x	Sum	Initial Cost (dollars)	Variable Cost (dollars/1000 mi)
None	Uncontrolled[2]	11	4	15	0	0.00
CAP 68	1968 clear air package[1] Changes in the settings of carburetor and distribution parameters and some engine design modifications	4.4	6	10.4	21	0.00
LLO 71	1971 low lead, low octane engine[1] Engine modified to run well on low lead fuel	3.5	4.5	8	35	1.40
TUNE	Minor engine tune-up[1] Adjustment of idle speed and mixture	4	6.4	10.4	0	3.00
KIT	Used car kit[1] A General Motors package of parts and instructions for older cars with ignition and carburetion changes	5.1	2.8	7.9	35	0.00
CAT	Catalytic converter[1] A double catalyst bed with both oxidizing and reducing agents (low-lead gas required)	0.9	2	2.9	135	4.40
EXR	Exhaust gas recirculation[1] Some exhaust routed back to the carburetor with engine modified to run on unleaded fuel	3.4	2.1	5.5	60	1.40
VSAD	Vacuum spark advance disconnect[2] Restrictions placed on the spark advance in some driving modes	5.5	1.6	7.1	35	3.00
FLAME	Flame-afterburner[2] A system designed to complete the burning of exhaust gasses by means of a sustained flame in the exhaust system	0.3	4	4.3	200	3.00

[1] D. N. Dewees, "Automobile air pollution: An economic analysis," Ph.D. dissertation, Dep. Economics, Harvard University, Cambridge, Mass., Sept. 1971.
[2] P. B. Downing and L. Stoddard, "Benefit/cost analysis of air pollution control devices for used cars," Project Clear Air, Univ. California, Los Angeles, Sept. 1970.

Such a system could be quite expensive. The capital and service charges for the device would have to be paid. An enforcement mechanism would be necessary to assure that emissions were being propertly recorded; however, this enforcement mechanism might be self-supporting. Finally, the effectiveness of the system would be diminished if it is subverted by drivers; technology and enforcement will affect rates of subversion. Thus, when we consider the emission market solution, we must include in the net social value what could possibly be very high system costs of administration, enforcement, and subversion. It is possible that these system costs could make the emission market have a negative net social value compared with the present social policy.

Other Social Alternatives: Other social alternatives are easy to evaluate within the driver–breather model. For example, instituting a gas tax on gasoline sold, or having the tax depend on the class of car driven and on its time of last emission checkup. Other alternatives could be a mileage tax for each type of car collected according to odometer readings reported regularly by drivers. All of these are surrogates for the direct measurement of emissions. They lack the very attractive signaling effects of the emission tax, but their system costs could be sufficiently low to make them a better social policy.

The Present Alternative: Finally, we come to the present social alternative: requiring vehicles to be factory-modified with emission-control equipment. We observe first that this alternative has extremely remote signaling effects. Given that you have purchased one of these cars, there is little incentive to reduce emissions by curtailing driving or installing additional equipment. Nor is there any evidence that the present alternative is best when the effects on drivers, breathers, and

TABLE III
DRIVERS' DECISIONS AS A FUNCTION OF TAX RATE

TAX RATE	DEVICE SELECTED BY DRIVER CLASS > 30K	20-29K	15-20K	10-15K	5-10K	< 5K	
0.0000	PAYTAX	PAYTAX	PAYTAX	PAYTAX	PAYTAX	PAYTAX	
0.1000	CAP 68	CAP 68	CAP 68	CAP 68	CAP 68	PAYTAX	
0.2000	CAP 68	CAP 68	CAP 68	CAP 68	CAP 68	PAYTAX	
0.3000	KIT	KIT	KIT	CAP 68	CAP 68	PAYTAX	
0.4000	KIT	KIT	KIT	KIT	KIT	CAP 68	BASE
0.5000	KIT	KIT	KIT	KIT	KIT	CAP 68	CASE
0.6000	KIT	KIT	KIT	KIT	KIT	CAP 68	
0.7000	KIT	KIT	KIT	KIT	KIT	KIT	
0.8000	KIT	KIT	KIT	KIT	KIT	KIT	
0.9000	KIT	KIT	KIT	KIT	KIT	KIT	
1.0000	KIT	KIT	KIT	KIT	KIT	KIT	
1.1000	KIT	KIT	KIT	KIT	KIT	KIT	
1.2000	KIT	KIT	KIT	KIT	KIT	KIT	
1.3000	FLAME	KIT	KIT	KIT	KIT	KIT	
1.4000	FLAME	FLAME	FLAME	KIT	KIT	KIT	
1.5000	FLAME	FLAME	FLAME	FLAME	KIT	KIT	
1.6000	FLAME	FLAME	FLAME	FLAME	KIT	KIT	
1.7000	FLAME	FLAME	FLAME	FLAME	CAT	KIT	
1.8000	FLAME	CAT	CAT	CAT	CAT	KIT	
1.9000	CAT	CAT	CAT	CAT	CAT	KIT	
2.0000	CAT	CAT	CAT	CAT	CAT	KIT	
2.1000	CAT	CAT	CAT	CAT	CAT	KIT	
2.2000	CAT	CAT	CAT	CAT	CAT	KIT	
2.3000	CAT	CAT	CAT	CAT	CAT	KIT	
2.4000	CAT	CAT	CAT	CAT	CAT	KIT	
2.5000	CAT	CAT	CAT	CAT	CAT	KIT	
2.6000	CAT	CAT	CAT	CAT	CAT	KIT	
2.7000	CAT	CAT	CAT	CAT	CAT	KIT	
2.8000	CAT	CAT	CAT	CAT	CAT	KIT	
2.9000	CAT	CAT	CAT	CAT	CAT	CAT	
3.0000	CAT	CAT	CAT	CAT	CAT	CAT	
$/KG							

TABLE IV
BASE CASE RESULTS

TAX RATE R: $0.50/kg

DRIVER CLASS (Thousands of Miles Driven)	DEVICE	ANNUALIZED FIXED COST	VARIABLE COST ($/yr)	TAX ($/yr)	CHANGE IN ANNUAL MILES DRIVEN
> 30	KIT	9	29	150	-2095
20-29	KIT	7	17	94	-1091
15-20	KIT	5	12	67	- 655
10-15	KIT	4	9	48	- 409
5-10	KIT	4	5	29	- 218
< 5	CAP 68	4	0	13	- 74

ALL DRIVERS:

Average mileage	9962
Average emissions	8.53 g/Mi
Average tax/year	$43

system costs are considered. In fact, the wisdom of the present alternative can be evaluated only within the type of structure we have developed.

E. Numerical Results

Some preliminary insights into wise social choices for emission control can be gained from a numerical model developed by Judd to investigate the usefulness of the evaluation structure we have presented [4]. The model is applied to Los Angeles County in 1965, a reference year before extensive automotive modification and for which data are readily available. The 2.5 million households are assumed to be willing to pay $100 each to eliminate the 500 million kg of emissions generated annually by 3.3 million autos driving 10 000 mi/year and averaging 15 g of emissions per mile. From these data we find r = $0.50/kg is the nominal charge that should be placed on emissions.

Drivers were divided into six classes according to the number of miles they would normally drive by the partition of 5, 10, 15, 20, and 30 thousand miles per year. Each driver faced economics according to his class. Table I shows the class, the percentage of drivers in each class, and the variable cost per mile. Table II shows the various modifications to a standard car considered to be available to each driver. Table III displays how the drivers would respond in terms of these modifications if they faced various levels of emission tax. If the tax is small, up to and including $0.30/kg, low mileage drivers simply pay the tax, while high mileage drivers install CAP 68. As the tax increases beyond $0.30/kg, installation of KIT begins to drive out CAP 68 for all but the lowest mileage class. The nominal charge of $0.50 finds KIT used by all drivers except those driving less than 5000 mi/year; they use CAP 68. Increasing the charge above $0.60 drives everyone to KIT. When the charge reaches $1.30, FLAME becomes increasingly desirable for high mileage drivers. The charge has to reach 1.70 before CAT becomes desirable and it does not drive out KIT entirely until the charge reaches $2.90.

Table IV shows the detailed results of the nominal case where the emission charge is $0.50. We observe here how the drivers' choices have affected their costs, taxes, and mileage driven. We also observe that the effect of the $0.50 charge is to cut automotive emissions approximately in half.

TABLE V
SOCIAL VALUE INFERRED FROM REQUIRED DEVICES

DEVICE	SOCIAL VALUE RANGE FOR WHICH DEVICE IS OPTIMAL $/KG	IMPLIED ANNUAL VALUE OF ELIMINATING AIR POLLUTION (PER HOUSEHOLD) $/YEAR
None	0 - 0.05	0 - 10
CAP 68	.05 - 0.30	10 - 60
KIT	.30 - 1.35	60 - 270
FLAME	1.35 - 1.75	270 - 350
CAT	> 1.75	> 350
OTHERS	NEVER OPTIMAL	

Table V shows how the emission charge would have to be changed to make the various devices optimal, and then illustrates what this would mean in terms of the value to each household of eliminating automotive air pollution. The catalytic converter CAT is being installed on most 1975 cars in this region. Each household in Los Angeles County would have to be willing to pay at least $350/year, in 1965 dollars, if this is to be a wise choice for society. Since this figure represents about 5 percent of 1965 household annual income there is reason to reexamine the advisability of requiring the catalytic converter.

F. Concluding Observations

We should not take the results of the numerical model too seriously for it is only a pilot analysis of the emission situation in one region. Yet it could be expanded to comprise as many additional aspects of the problem as we wish to include, from nonautomotive pollution to the use of public transit. Our goal should be a balanced analysis comprehensive enough to investigate most concerns or suggestions regarding emissions.

It should provide a framework large enough to allow varying views to be assessed against a common background, so that information and reason will be more valuable tools of advocacy than rhetoric and histrionics.

Within the model several specific issues can be examined, such as uncertainty in the effectiveness or cost of any alternative, the choice of criteria for evaluating social desirability, or the value of emission measurement devices. The construction of a realistic model with this capability would require expertise in law, engineering, economics, psychology, and environmental medicine, to name but a few. The cost might be several hundred thousand dollars. But even this cost is a fraction of 1 percent of the total annual cost of automotive emissions estimated for Los Angeles County alone. The question is not when we can afford such models but rather whether we can afford to be without them.

VI. Hurricane Seeding—Uncertainty and Responsibility

As a further example of social decision analysis we can cite a recent study on hurricane seeding [5]. Seeding hurricanes with silver iodide crystals in an attempt to mitigate their destructive effects was a promising idea experimentally tested during the 1960's. The first experiments on Hurricanes Esther (1961) and Beulah (1963) were encouraging, but the experiment on Hurricane Debbie (1969) was the most impressive experimental result. Massive seedings on August 18 and 20, 1969, were followed by reductions of 31 and 15 percent in peak wind speed.

Hurricanes are costly to the nation. Annual property damage averages over $400 million dollars; individual Hurricanes Betsy (1965) and Camille (1969) each cost about 1.5 billion dollars of property damage. Seeding is cheap: about $0.25 million dollars. If the results of the Debbie experiment are applicable to even a few hurricanes, the savings could be very significant.

A. The Decision

The U.S. Hurricane Modification Program at the time of the study was exclusively scientific in its mission. Any seeding of hurricanes that could threaten coastal areas was prohibited. The decision that had to be faced was whether this prohibition should be removed. If it were, then study could begin on whether any specific hurricane threatening the U.S. should be seeded.

B. The Effect of Seeding

The characteristic of hurricanes of greatest economic consequence is their maximum sustained surface wind speed, since this is the characteristic that is primarily responsible for property damage. Happily, this is just the characteristic that seeding promises to reduce. However, even if seeding were certain to reduce wind speed, it still would not be clear that it should be performed. For by the very act of seeding, the Government stamps its name on the subsequent damage caused by the hurricane—"acts of God" become "acts of Uncle Sam." Even though scientists may state that the damage would probably have been worse without the seeding, the government would be in a difficult legal position.

The position is aggravated because the natural behavior of hurricanes is notoriously unpredictable. They can spontaneously intensify or diminish in their destructive power. If a hurricane were seeded when its peak winds were 100 miles per hour and it crossed the coast with peak winds of 125 miles per hour there will be few who pause to think that the winds might have been 150 miles per hour at land encounter if the seeding had not been performed. The public may well note only that a government-seeded hurricane intensified.

C. Characterizing Information on Seeding Effects

Three hypotheses were used to describe possible beliefs about the effect of seeding:

1) H_1, the "*beneficial*" hypothesis: the average effect of seeding is to reduce the maximum sustained wind speeds;
2) H_2, the "*null*" hypothesis: seeding has no effect on hurricanes;
3) H_3, the "*detrimental*" hypothesis: the average effect of seeding is to increase the maximum sustained wind speed.

After discussions with the scientists concerned, hypothesis H_1 was taken to mean a reduction of 15 percent in wind speed over a 12-h period; hypothesis H_3, a 10-percent increase. The actual wind speed after the 12-h period as a percentage of the initial wind speed is assumed to have a normal distribution in all cases. If seeding has no effect, natural variability is described by a standard deviation of 15.6 percent. If seeding has an effect, the standard deviation of its average effectiveness (7 percent) and the standard deviation of its effectiveness in any particular case (7 percent) combine to produce a 10-percent standard deviation for the effect of seeding. When this is superimposed on the natural variability, we obtain a standard deviation of 18.6 percent for the percentage change in wind speed of a seeded hurricane. Thus the 12-h percentage wind change distributions are

$$P(w \mid H_1) = f_n(85 \text{ percent}, 18.6 \text{ percent})$$
$$P(w \mid H_2) = f_n(100 \text{ percent}, 15.6 \text{ percent})$$
$$P(w \mid H_3) = f_n(110 \text{ percent}, 18.6 \text{ percent}).$$

It was now necessary to assign probabilities to the hypotheses based on meteorological judgment. The Debbie experiment produced what were considered to be two independent observations of 31- and 15-percent reductions in wind speed. In other words, on one occasion, the wind was 69 percent of its original speed; in the other, 85 percent. The relative likelihood of these observations under each hypothesis can be computed by multiplying together the ordinates of the corresponding normal density function at the values 69 percent and 85 percent. If we let D be the Debbie observations, the result is:

$$P(D \mid H_1) \sim 3.21$$
$$P(D \mid H_2) \sim 0.61$$
$$P(D \mid H_3) \sim 0.173.$$

The pre-Debbie hypothesis probabilities $P(H_1)$ must be related to the post-Debbie probabilities $P(H_i D)$ by Bayes' equation,

$$P(H_i \mid D) = \frac{P(D \mid H_i) P(H_i)}{\sum_{i=1}^{3} P(D \mid H_i) P(H_i)}.$$

Of course, since no one recalled the probabilities before the experiment, they had to be reconstructed. Meteorologists believed that before Debbie, H_1 was more likely than H_3 if seeding had any effect on hurricanes. They also believed that after Debbie they saw H_1 and H_2 as equally likely. Prior and posterior probabilities that meet these conditions and also satisfy Bayes' equation are

$$P(H_1) = 0.15 \qquad P(H_2) = 0.75 \qquad P(H_3) = 0.10$$

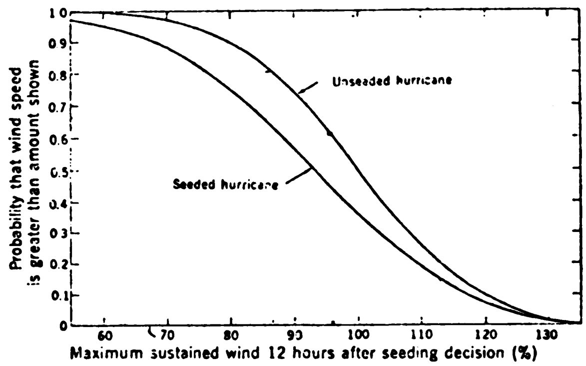

Fig. 6. Probability distribution on 12-h wind changes for the seeded and unseeded hurricane.

and

$$P(H_1 | D) = 0.49 \qquad P(H_2 | D) = 0.49 \qquad P(H_3 | D) = 0.02.$$

Note that before Debbie these assignments revealed three to one odds that seeding would have no effect, and that there was one chance in ten that it would have a detrimental effect. The Debbie experiment reduced the probability of a detrimental effect to only one in fifty, and made a beneficial and no effect equally likely. Sensitivity analysis showed there was little incentive in refining the state of knowledge beyond this specification.

D. The Wind Speed Probability Distribution

Having the posterior probabilities on the hypotheses and the percentage change in wind speed distributions given the hypotheses, we can multiply and sum to obtain the probability distribution on percentage change in wind if a hurricane is seeded. This distribution is plotted in complementary cumulative form in Fig. 6. It is also shown for a nonseeded hurricane; this is simply the complementary cumulative of $P_n(w | H_2)$. Note that for any wind speed you choose, the probability that a hurricane will exceed that wind speed is always greater for an unseeded than for a seeded hurricane. We say that the seeding alternative stochastically dominates the alternative of leaving the hurricane alone. As long as higher winds are worse, the logical man would prefer seeding to not seeding.

We must note, however, that seeded hurricanes can still intensify. For example, the probability that a seeded hurricane will intensify by 10 percent or more is 0.18, almost one chance in five. We choose seeding because an unseeded hurricane has an even higher probability of such intensification: 0.26.

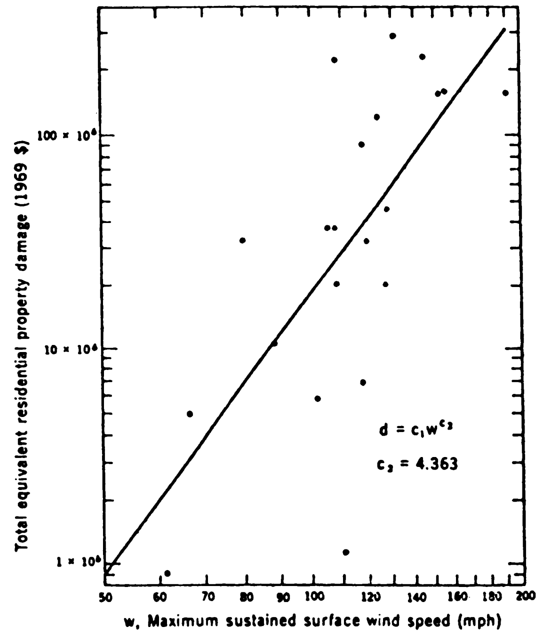

Fig. 7. Property damage plotted against maximum sustained wind speed.

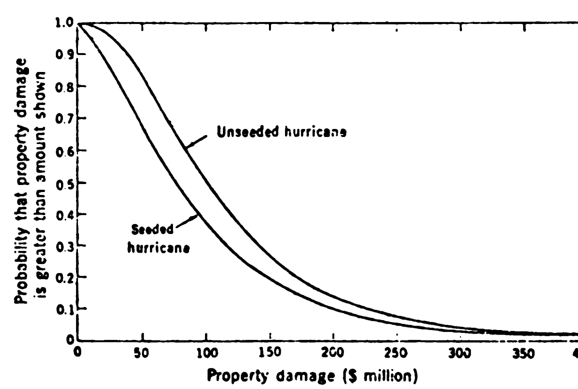

Fig. 8. Probability distributions on property damage for the seeded and unseeded hurricane.

E. Property Damage

To make the case complete we must examine property damage to find the economic consequences of wind speed reduction. Fig. 7 shows 21 instances of residential property damage caused by hurricanes with the associated wind speed. A least-squares fit produced an exponent of 4.363 in relating damage to wind speed. This implies that a 15-percent reduction in wind speed will cause a 51-percent reduction in damage. If seeding is beneficial, it can have a major economic impact.

Using the wind-damage function of Fig. 7, we can transform the distribution on wind speed of Fig. 6 into the distribution on property damage shown in Fig. 8. Once again we see that the seeding alternative stochastically dominates the alternative of not seeding. The probability of property damage exceeding any given level is always higher for the unseeded hurricane. If our objective is to minimize property damage, then seeding is the better alternative.

F. The Seeding Decision

A detailed description of the seeding decision for a nominal hurricane appears in Fig. 9. Here the two alternatives of seeding and not seeding lead to five-branch chance nodes with probabilities on change in maximum sustained wind obtained by discretizing the distributions of Fig. 6. The property damage associated with each change in maximum wind speed is shown at the tips of the tree. The expected loss for the seeding alternative is \$94.33 million including the \$0.25 million cost of seeding. The expected loss of not seeding is \$116 million, or \$21.67 million higher. Not only is seeding the better alternative, it is in an absolute sense significantly better.

But perhaps this result is very sensitive to the assumptions made. It does not appear to be. The probability of H_3 must

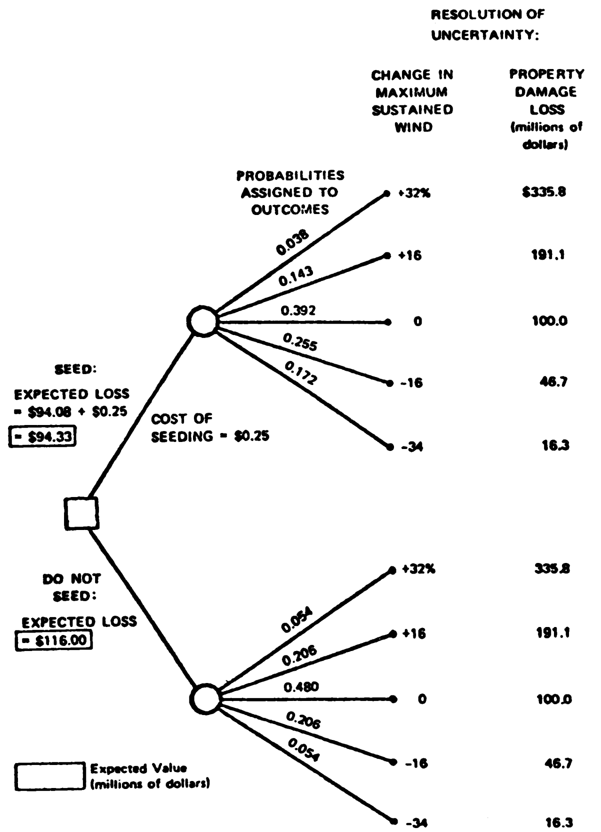

Fig. 9. The seeding decision for the nominal hurricane.

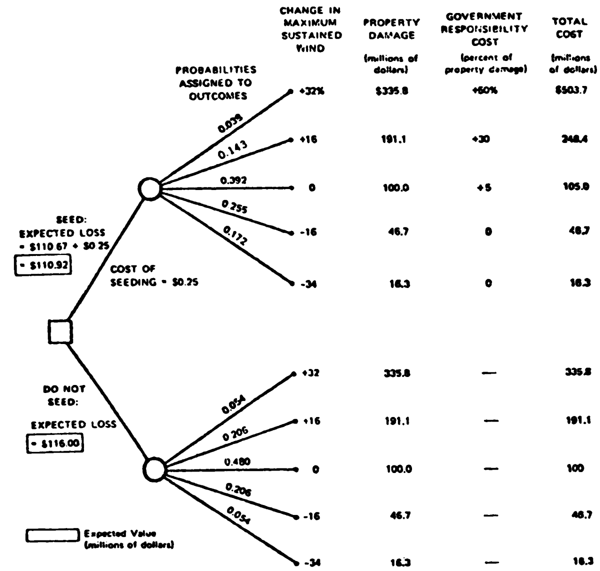

Fig. 10. The seeding decision for the nominal hurricane (government responsibility cost included).

be raised to 0.07 before stochastic dominance of seeding no longer holds. Even if it is raised to 0.20 seeding still has an expected loss 7 percent less than the not-seeding alternative. As stated in [5]: "the results of extensive sensitivity analyses may be summarized as follows: The expected loss in terms of property damage appears to be about 20 percent less if the hurricane is seeded. Varying the assumptions of the analysis causes this reduction to vary between 10 and 30 percent but does not change the preferred alternative."

G. Government Responsibility

If seeding appears to be so economically desirable, why is it not done? The answer must lie in the reluctance of government officials to expose the government and themselves to charges of tampering detrimentally with the weather. It would be surprising if government officials did not have such concern; if the concerns are large in comparison with the objective benefits to society, perhaps it is wise to forgo the benefits using the argument that the gain is not worth the trouble.

To see how large such concerns would have to be before hurricane seeding should be avoided, we introduced the concept of government responsibility cost. This cost is added to the economic cost to represent the trouble that the government would have in explaining apparent unfortunate consequences of seeding. We assigned it by asking, in each such case, what percentage increase in economic losses the official would be willing to incur without seeding before he would prefer to have the original losses and the necessity of explaining them if the storm had intensified by a given amount after seeding. For example, an official might say that he was just indifferent between $150 million in damage from a nonseeded storm and $100 million in damage from a seeded storm that had intensified by 32 percent after seeding. In this case, we would say that the government responsibility cost of seeding a storm that then intensified by 32 percent and produced $150 million of damage was 50 percent.

Fig. 10 shows the seeding decision with government responsibility cost included. A wind increase of 32 percent in a seeded hurricane is associated with a 50-percent increase in property damage in accordance with our discussion. An increase of 16 percent in wind speed for a seeded hurricane is assigned a 30-percent government responsibility cost. Even if seeding is followed by no wind change, a 5-percent government responsibility cost is assigned to account for charges that the hurricane would have diminished if the government had left it alone. Cases where the wind actually decreases after seeding are assigned a zero government responsibility cost, since we would expect few complaints about the decision in this case.

No government responsibility costs are assigned if the storm is not seeded, although there is an argument for doing so. Given that seeding is promising but uncertain, is it not just as irresponsible to refrain from seeding a storm that subsequently intensifies as it is to seed it? Nevertheless, since errors of omission seem to attract less notice than errors of commission, no government responsibility costs are associated with a refusal to seed.

Analysis of the tree in Fig. 10 shows that in spite of the relatively high values for government responsibility costs, the seeding alternative has an expected loss less than that of not seeding. The political costs must be considerably higher than those assumed if the model is to recommend against seeding. In fact, for a hurricane in the $1 billion class, like Betsy (1965) and Camille (1969), government responsibility costs would have to be equivalent to $200 million in property damage before seeding should be avoided.

TABLE VI
SUMMARY OF VALUE OF ADDITIONAL INFORMATION ON EFFECT OF SEEDING[a]

Item	Nominal hurricane used in analysis	%	Single hurricane Season	All future hurricane seasons, discounted
Expected Property Damage without Seeding	116.0	100	220.0	3142
Expected Value of Perfect Information	13.6	11.8	26.0	370
Expected Value of a Field Experiment Consisting of Two Experimental Seedings	5.4	4.7	10.2	146

[a]Only the 50 percent of hurricanes that are assumed to be possible candidates for seeding on the basis of tactical considerations are considered. If all hurricanes are assumed to be candidates for operational seeding, the figures of the last two columns should be doubled. All figures are in millions of dollars.

H. Value of Information and Experimentation

One of the most valuable aspects of decision analysis is its ability to place a monetary value on acquiring information on the uncertain aspects of the decision problem [2]. Ideally, the information would be perfect, producing complete resolution of the uncertainty. Practically, information is provided by tests or experiments that only partially resolve uncertainty.

A value of information analysis was performed on the hurricane seeding decision with the results shown in Table VI. The expected value of perfect information on the effectiveness of seeding is $13.6 million when considering whether to seed one nominal hurricane, $26 million for each hurricane season, and $370 million for all future seasons when discounted at a 7-percent rate. The value of another Debbie-size experiment, which would provide partial information, is $5.4 million for the nominal hurricane, $10.2 million for the season, and $146 million in present value. If the question were whether or not to seed a hurricane in the billion-dollar class, the experiment would be worth about 10 times as much as for the nominal hurricane, or about $50 million for that decision alone. Clearly there is a high value in conducting the hurricane seeding experiments.

I. Conclusion

This analysis served to illuminate the hurricane-seeding decision. The decision cannot be avoided: either the government must assume responsibility for seeding and hence for what appear to be detrimental effects of seeding or it must assume responsibility for not seeding and thus exposing the public to greater losses from hurricanes.

VII. NUCLEAR POWER PLANT SAFETY

The question of nuclear power plant safety offers an important illustration of the use of decision analysis in considering the general issue of social safety.

A. Safety

The question of safety is an aspect of virtually all large human systems. Whether in energy production and distribution, transportation, manufacturing, or space exploration, everyone prefers more safety to less. The issue is not whether any system can be made more safe—the answer is almost always "yes." The question is how to find a proper balance between the benefits and costs of the entire system, including the costs associated with safety.

To carry the discussion one level further, we might try to divide the safety costs into the costs of building and maintaining the safety-ensuring systems and the costs incurred from accidents. However, it is not always easy to identify safety system costs per se as separate from total system costs. For example, while the addition of seatbelts in an automobile could be considered as primarily a safety investment, an improved braking system may enhance both safety and performance: its costs cannot reasonably be considered solely an investment in safety.

A more logical way to proceed is to compare total system benefits less total system costs for each alternative system design. This comparison, while simple in principle, is challenging in practice because many system effects are particularly difficult to detect and evaluate. For example, if we consider the alternative of raising the gasoline tax as an energy conservation measure, we find many potential safety effects. If the increased tax causes less traveling, that would presumably increase safety. There would likely be less transport of dangerous gasoline by tank trucks—another diminished hazard. But would people switch to smaller, more vulnerable cars? As you can see, the net effect on safety of such an alternative poses difficult analysis problems.

Now let us return to the accident cost side of the safety evaluation. When, in spite of safety devices, an accident occurs, how should the social cost be measured? The accident could cause death, injury, sickness, property destruction or loss of use, aesthetic impairment, etc. One approach to evaluation is to assess a monetary value to each unit of these consequences, multiply by the number of units, and sum to obtain the total social costs of the accident. While some might see this procedure as callous, it is important to observe that we are not attempting to value a human life or injury in a moral sense, we are rather assigning values to life and the other loss categories as a means of making social decisions. The values are not measurements, but the results of a decision made by society concerning the balancing of social benefit and social risk. By assessing these values in monetary terms, we are able to determine the monetary amount that should be spent on safety systems, and at a higher level, to choose between entire systems.

To demonstrate, consider two possibilities for our society. In the first we assign a $1 thousand value to the consequence of a human death resulting from an involuntarily assumed risk. In the second we assign a $10 million value. The first assignment would lead to decisions regarding aircraft, roads, fire pro-

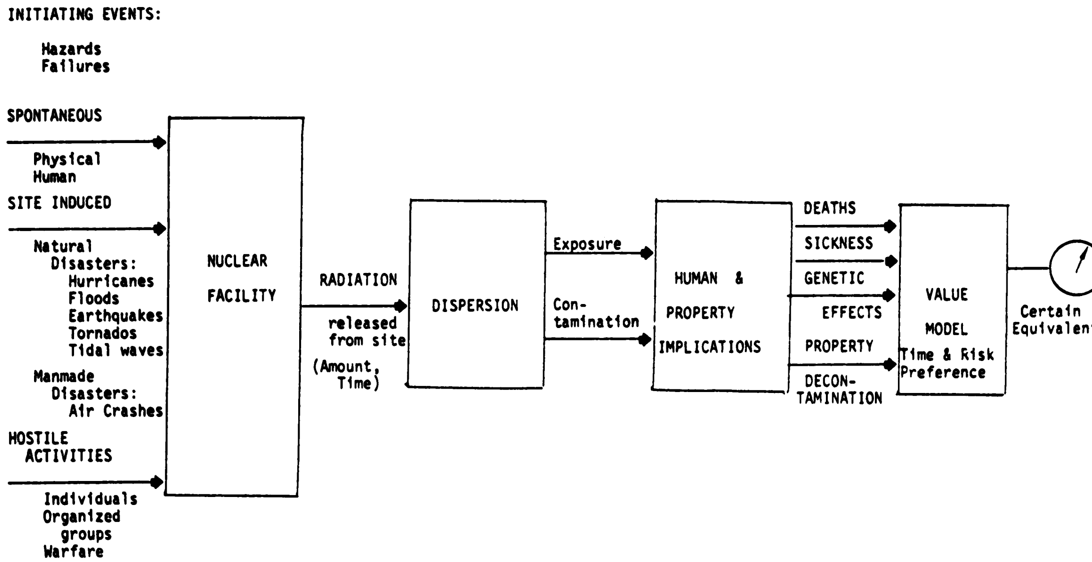

Fig. 11. Nuclear power plant safety—accident model logical structure.

tection, etc., that would soon leave us a nation of the afraid. A popular clamor would arise for "safer" social alternatives. On the other hand, the $10 million assignment would lead to decisions requiring huge sums to be spent on safety equipment. The price of airplane tickets, cars, and electricity would grow until the population again complained about the "waste" of money on "unnecessary" safeguards. Thus the problem is to assign a value to human life (and to injury, etc.) that reflects what we can afford to pay as a society to avoid a death from involuntary risk. Various analyses may be helpful in assessing this value, but in the end, it is a policy decision for society.

Therefore, a logical way to look at the safety problem in context is to select the alternative with the greatest excess of benefits over costs, where the costs due to accidents are evaluated using value assignments made by a social decision process. While the safest system is the one with the lowest accident cost, the most desirable system for society will seldom be the safest.

We shall now present a procedure for evaluating the accident cost of a nuclear power plant. As we have shown, safety is an inseparable part of the process for choosing between system alternatives, such as choosing between fossil and nuclear power generation. Consequently, a safety analysis cannot, by itself, logically determine whether nuclear power generation is a desirable alternative for our society, or even if it is "safe enough" compared to other energy conversion systems.

B. Methodology–Nuclear Power Plant Safety

Assessing nuclear power plant safety may be discussed conveniently with the aid of Fig. 11. Any particular facility is exposed to various hazards. These hazards could be spontaneous failures, such as physical failure of a piece of equipment, or a human error caused by incompetence, inattention, or incapacity. There could also be failure because the reactor is at a particular site. These site-included failures, in turn, could be caused by either natural disasters like windstorms, floods, or earthquakes or by man-made disasters, like airplane crashes. Finally, the reactor might be the target of hostile activities by individuals or organized groups.

It is possible that as a result of those hazards, the nuclear power plant will suffer damage. The question regarding public health and safety is whether a radioactive release from the site occurs, for such a release could be the cause of public loss. The release from the site could be described by the amount of radioactivity of each type emitted in each future time interval.

For a given release, the magnitude of the loss will be governed by the water, by the subsequent spatial and temporal distribution of population in the surrounding area, and by the land use patterns in the area. The dispersion model thus utilizes meteorological, geographic, and demographic information to predict the human exposure and property contamination produced by a given release.

Since humans suffer from differing types and amounts of radiation in different ways, a human implications model is necessary to assess deaths, sickness and genetic damage likely as the consequence of a given exposure. Similarly, the cost of decontaminating or burying property infested with radioactivity is assessed in a property implications model.

Finally, the value model reduces to a single value measure the combined societal evaluation of both human and property losses, in accordance with the discussion of the last section.

C. Model Implementation

The logic of Fig. 11 can be conveniently incorporated in a probability tree. Successive chance nodes of the tree would characterize such uncertain possibilities as the magnitude of an earthquake, the amount of release, the extent of contamination, and the number of deaths. Experts on the various stages could assign probabilities either directly, on the basis of the analysis of data (e.g., records of pipe breakage), or by exercising subsidiary models (e.g., meteorological dispersion simulations). The net result would be a probability distribution on the magnitude of the annual loss that could be produced by a single reactor, or, for certain purposes, by all reactors.

Studies of this kind have been performed by various agencies, but always with serious limitations. They have tended to emphasize the physical causes of accidents and to stop short of considering the value implications of the accident consequences. Studies of the effects of hostile activity ranging from disgruntled employees to vulnerability in time of war have been conspicuously rare. Yet there is no reason why such hazards cannot be analyzed within the same framework used for physical malfunctions. In fact, as improved engineering reduces physical hazards, it becomes increasingly likely that if a serious power plant accident should occur, the cause will be found in human intervention.

D. Summary

Decision analysis models provide a comprehensive and practical framework for investigating safety effects. In choosing between two forms of power generation, safety system costs and uncertainties can be balanced against economic and social benefits. The value of further experiments and the costs of delay can be computed and compared to evaluate a proposed course of action. The controversies that permeate licensing hearings could be reduced to their simplest issues, thus offering the hope that energy now devoted to argument would be applied to resolution.

VIII. Conclusion

We have now seen how the use of extrapersonal models provides a logical and open procedure for making social decisions. Various stakeholders can offer their views and then observe how their contributions would affect the choice of alternatives. Anyone could suggest an alternative, all citizens, either individually or in groups, could influence the values used, and experts would be able to supply information within the realm of their expertise.

The net result would be a more distributive and more effective form of government. We would see more debate on the fundamental issues of the society, rather than on individual laws. For example, the question of whether society has the right to protect mature competent citizens from the consequences of their own actions would have to be faced forthrightly rather than skirted.

A government of law that is not logically based on the information and values of the society is just as arbitrary in its effect on individuals as a government of men.

References

[1] R. A. Howard, "Decision analysis: Applied decision theory," in *The Proceedings of the Fourth International Conference on Operations Research.* New York: Wiley, 1966, pp. 55-71.

[2] —, "The foundations of decision analysis," *IEEE Trans. Syst., Sci., Cybern.*, vol. SSC-4, pp. 211-219, Sept. 1968.

[3] —, "Decision analysis in systems engineering," in *Systems Concepts, Lectures on Contemporary Approaches to Systems.* New York: Wiley, 1973, pp. 51-85.

[4] R. A. Howard and B. R. Judd, "A pilot model for evaluation of economic incentive and direct control policies for limiting automobile air pollution," report for the National Science Foundation under Grant GK-36491, Dec. 1973.

[5] R. A. Howard, J. E. Matheson, and D. W. North, "The decision to seed hurricanes," *Science*, vol. 176, pp. 1191-1202, June 1972.

23

THE DECISION TO SEED HURRICANES

R. A. Howard and J. E. Matheson
Strategic Decisions Group

D. W. North
Decision Focus Incorporated

Article reprinted with permission from Science, 16 June 1972, Vol. 176, pp. 1191-1202. Letters appeared in Science, 23 February 1973, Vol. 179, pp. 744-747; and 14 September 1973, Vol. 181, pp. 1072-1073.

The Decision to Seed Hurricanes

On the basis of present information, the probability of severe damage is less if a hurricane is seeded.

R. A. Howard, J. E. Matheson, D. W. North

The possibility of mitigating the destructive force of hurricanes by seeding them with silver iodide was suggested by R. H. Simpson in 1961. Early experiments on hurricanes Esther (1961) and Beulah (1963) were encouraging (*1*), but strong evidence for the effectiveness of seeding was not obtained until the 1969 experiments on Hurricane Debbie (*2*). Debbie was seeded with massive amounts of silver iodide on 18 and 20 August 1969. Reductions of 31 and 15 percent in peak wind speed were observed after the seedings.

Over the last 10 years property damage caused by hurricanes has averaged \$440 million annually. Hurricane Betsy (1965) and Hurricane Camille (1969) each caused property damage of approximately \$1.5 billion. Any means of reducing the destructive force of hurricanes would therefore have great economic implications.

Decision to Permit Operational Seeding

In the spring of 1970 Stanford Research Institute began a small study for the Environmental Science Service Administration (ESSA) (*3*) to explore areas in which decision analysis (*4, 5*) might make significant contributions to ESSA, both in its technical operations and in its management and planning function. At the suggestion of Myron Tribus, Assistant Secretary of Commerce for Science and Technology, we decided to focus the study on the decision problems inherent in hurricane modification (*6*).

The objective of the present U.S. government program in hurricane modification, Project Stormfury, is strictly scientific: to add to man's knowledge about hurricanes. Any seeding of hurricanes that threaten inhabited coastal areas is prohibited. According to the policy currently in force, seeding will be carried out only if there is less than a 10 percent chance of the hurricane center coming within 50 miles of a populated land area within 18 hours after seeding.

If the seeding of hurricanes threatening inhabited coastal areas is to be undertaken, it will be necessary to modify the existing policies. The purpose of our analysis is to examine the circumstances that bear on the decision to change or not to change these existing policies.

The decision to seed a hurricane threatening a coastal area should therefore be viewed as a two-stage process: (i) a decision is taken to lift the present prohibition against seeding threatening hurricanes and (ii) a decision is taken to seed a particular hurricane a few hours before that hurricane is expected to strike the coast. Our study is concentrated on the policy decision rather than on the tactical decision to seed a particular hurricane at a particular time. It is also addressed to the experimental question: What would be the value of expanding research in hurricane modification, and, specifically, what would be the value of conducting additional field experiments such as the seedings of Hurricane Debbie in 1969?

Our approach was to consider a representative severe hurricane bearing down on a coastal area and to analyze the decision to seed or not to seed this "nominal" hurricane. The level of the analysis was relatively coarse, because for the policy decision we did not have to consider many geographical and meteorological details that might influence the tactical decision to seed. We described the hurricane by a single measure of intensity, its maximum sustained surface wind speed, since it is this characteristic that seeding is expected to influence (*7*). The surface winds, directly and indirectly (through the storm tide), are the primary cause of the destruction wrought by most hurricanes (*8*). The direct consequence of a decision for or against seeding a hurricane is considered to be the property damage caused by that hurricane. (Injuries and loss of life are often dependent on the issuance and effectiveness of storm warnings; they were not explicitly included in our analysis.)

However, property damage alone is not sufficient to describe the consequence of the decision. There are indirect legal and social effects that arise from the fact that the hurricane is known to have been seeded. For example, the government might have some legal responsibility for the damage caused by a seeded hurricane (*9*). Even if legal action against the government were not possible, a strong public outcry might result if a seeded hurricane caused an unusual amount of damage. Nearly all the government hurricane meteorologists that we questioned said they would seed a hurricane threatening their homes and families—if they could be freed from professional liability.

The importance of the indirect effects stems in large part from uncertainty about the consequences of taking either decision. A hurricane is complex and highly variable, and at present meteor-

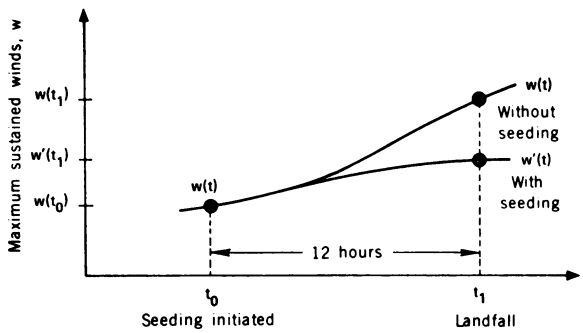

Fig. 1. Maximum sustained winds over time.

ologists cannot predict accurately how the behavior of a hurricane will evolve over time. The effect of seeding is uncertain also; consequently, the behavior of a hurricane that is seeded will be a combination of two uncertain effects: natural changes and the changes induced by seeding.

The seeding decision would remain difficult even if the uncertainty were removed. Suppose that, if the hurricane is not seeded, the surface wind intensifies as shown by the curve $w(t)$ in Fig. 1 and that, if the hurricane is seeded, the behavior of the wind is that shown by the curve $w'(t)$. The effect of the seeding has been to diminish the wind, thus reducing property damage, yet the wind speed $w'(t_1)$ when the hurricane strikes land at time t_1 is higher than the wind speed when the seeding was initiated at time t_0. Even if the decision-maker were certain of $w(t_1)$ and $w'(t_1)$, he would still have a difficult choice. If he chooses not to seed, the citizens may have more property damage. On the other hand, if he chooses to seed, the citizens may not perceive themselves as better off because of his decision. Instead, they may perceive only that the storm became worse after the seeding and they may blame the decision-maker for his choice. The trade off between accepting the responsibility for seeding and accepting higher probabilities of severe property damage is the crucial issue in the decision to seed hurricanes.

Decision under Uncertainty

The decision to seed a threatening hurricane would be taken about 12 hours before the hurricane is predicted to strike the coast. At this time the consequences are uncertain for both alternatives; the decision-maker does not know what amount of property damage will be sustained if the hurricane is seeded or is not seeded. We may illustrate the situation facing him in the form of a decision tree, as shown in Fig. 2. The decision-maker must select one of the two alternatives, seeding or not seeding. The decision cannot be avoided for inaction is equivalent to selecting the alternative of not seeding. Each alternative leads to a set of possible consequences: property damage caused by the hurricane and the responsibility incurred by the government. These consequences are, in turn, related to the intensity of the hurricane and whether or not it was seeded. The consequences for each alternative are uncertain at the time the decision is made; the uncertainty will be resolved after the decision-maker selects his choice. This decision under uncertainty may be examined according to the usual procedures of a decision analysis. We use the information that is currently available to develop a probability distribution over changes in the intensity of the hurricane as measured by its maximum sustained surface wind speed for each of the two decision alternatives. Then we use data from past hurricanes to infer a relation between wind speed and property damage. By assessing the consequences in property damage and government responsibility and the probability that these consequences will be achieved, we are able to determine which of the decision alternatives is the preferred choice.

Uncertainty in Hurricane Wind Changes

We began our analysis by considering the change in maximum sustained surface winds over a 12-hour period for a hurricane that is not seeded. If enough data had been available on the changes in hurricane wind speeds with time, a probability distribution for wind changes could have been based largely on these past data. Wind-change data were not available, but data were available for changes over time in the central pressure of hurricanes. The central pressure and the maximum wind speed of a hurricane are closely related; Holliday has shown that the available data can be summarized fairly well by a linear relation (*10*). We combined this relation with observations of the change in central pressure over a 12-hour period, using the assumption that the discrepancies from the Holliday relation are independent over a 12-hour period and independent of the change in central pressure. These assumptions imply a probability distribution on wind changes over a 12-hour period that is normal with a mean of zero and a standard deviation of 15.6 percent (*11*).

Therefore, present information is consistent with rather large natural changes in hurricane intensity over a 12-hour period. There is about one chance in six that a hurricane whose maximum sustained wind speed is 100 miles per hour will intensify over a 12-hour period to a maximum wind speed of over 115 miles per hour; there is also about one chance in six that the winds would naturally diminish to less than 85 miles per hour. In assessing these probabilities only general historical and meteorological information has been used. In a specific hurricane situation additional meteorological information might indicate that the hurricane would be more likely to intensify or more likely to diminish.

Effect of Seeding

The next step is to develop a probability distribution for the wind speed if the hurricane is seeded. The change in wind speed over 12 hours would then be a combination of the natural change occurring in the hurricane and the change caused by seeding. With the limited data available it is reasonable to assume that the two effects would be independent of each other and act in an additive fashion; for example, if the natural change is an intensification such that the maximum sustained wind speed is increased from 100 to $(100+x)$ percent, and if the effect of seeding is to diminish the maximum sustained wind speed from 100 to $(100-y)$ percent, the net observed change over 12 hours is from 100 to $(100+x-y)$ percent. A probability distribution has already been assigned for natural changes; we need to assign a probability distribution for the change caused by seeding. In developing this probability distribution it is necessary to distinguish between the effect of seeding on one hurricane and the average effect of seeding on many hurricanes. The effect of seeding on a particular hurricane might be quite different from its average effect.

After discussion with meteorologists associated with Project Stormfury, we concluded that the major uncertainty about the effect of seeding would be resolved if we knew which of the following mutually exclusive and collectively exhaustive hypotheses described the effect of seeding:

1) H_1, the "beneficial" hypothesis. The average effect of seeding is to reduce the maximum sustained wind speed.

2) H_2, the "null" hypothesis. Seeding has no effect on hurricanes. No change is induced in maximum sustained wind speed.

3) H_3, the "detrimental" hypothesis. The average effect of seeding is to increase the maximum sustained wind speed.

The scientific basis for the "beneficial" hypothesis, H_1, had its origins in the original Simpson theory (*1*). It has been modified and strengthened by Project Stormfury studies involving a computer model of hurricane dynamics (*1, 12*). This hypothesis, in fact, motivated the formation of the Project Stormfury research program. A possible basis for the "null" hypothesis, H_2, is that seeding does not release enough latent heat to affect the dynamics of the hurricane. The "detrimental" hypothesis, H_3, has been added to complete the set. Meteorologists do not have a basis in physical theory for H_3 that is comparable to that for H_1 or H_2.

Even if we know which of the hypotheses is true, there remain uncertainties about the effects of seeding. We now describe the approach we followed in creating a model to formalize existing knowledge about these uncertainties. Then we shall return to the hypotheses.

Let us suppose we have access to a clairvoyant who can tell us which hypothesis, H_1, H_2, or H_3, represents the actual effect of seeding on hurricanes. What probability would we assign to the 12-hour change in the maximum sustained winds of a seeded hurricane for each of his three possible answers? If the clairvoyant says H_2 is true, the assignment process is simple. Seeding has no effect, and the same probabilities are assigned to the wind speed w' if the hurricane is seeded as to the wind speed w if the hurricane is not seeded (*13*).

$$P(w'|H_2) = P(w) = f_N(100\%, 15.6\%) \quad (1)$$

If H_1 is the clairvoyant's answer, the process is more difficult. The average effect is known to be a reduction in storm intensity, but the amount of this average reduction is uncertain. The Simpson theory and the computer studies indicate that a reduction of 10 to 20 percent in wind speed should be expected, with 15 percent as the most likely value. This information was summarized by assigning to the change in wind speed a normal probability distribution with a mean of -15 percent and a standard deviation of 7 percent. An average reduction greater than 15 percent is considered as likely as an average reduction less than 15 percent, and the odds are about 2 to 1 that the average reduction will lie between 22 and 8 percent rather than outside this interval.

The effect of seeding on an individual hurricane would be uncertain even if the average effect of seeding were known. Odds of about 2 to 1 were considered appropriate that the effect of seeding would not differ from the average effect by more than about 7 percent; thus, a normal distribution centered at the average value with a standard deviation of 7 percent was judged an adequate summary of the information available on fluctuations in seeding effects. Combining the uncertainty about fluctuations with the uncertainty about the average effect leads to a probability distribution for the effect of seeding a specific hurricane that is normal with a mean equal to -15 percent and a standard deviation of 10 percent (*14*).

Adding the natural change in the hurricane over a 12-hour period to the change resulting from seeding gives the total 12-hour change occurring in a seeded hurricane if hypothesis H_1 is true. The probability distribution assigned to w' is then normal with a mean of 85 percent and a standard deviation of 18.6 percent (*15*):

$$P(w'|H_1) = f_N(85\%, 18.6\%) \quad (2)$$

The development of a probability distribution for w', if it is considered that H_3 is true, proceeds in a similar way. The average change effected by seeding is described by a normal probability distribution with a mean of $+10$ percent and a standard deviation of 7 percent. The fluctuations expected when an individual hurricane is seeded are normally distributed around the average with a standard deviation of 7 percent. Combining these uncertainties

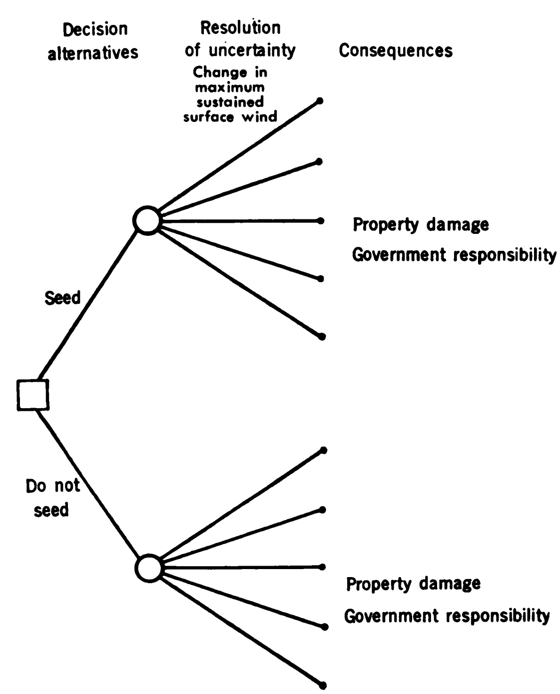

Fig. 2. The seeding decision: decision tree.

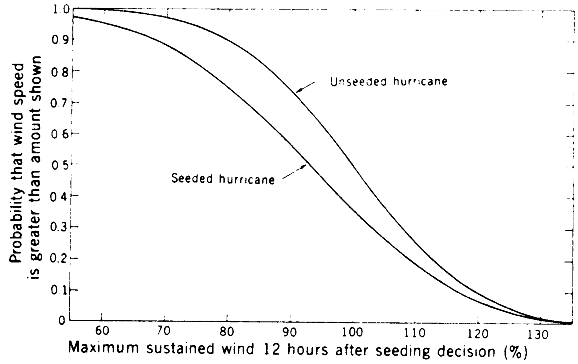

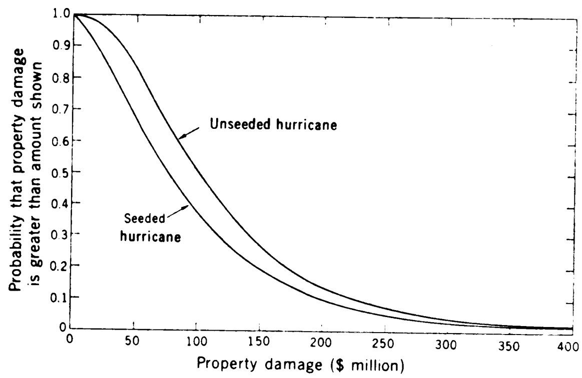

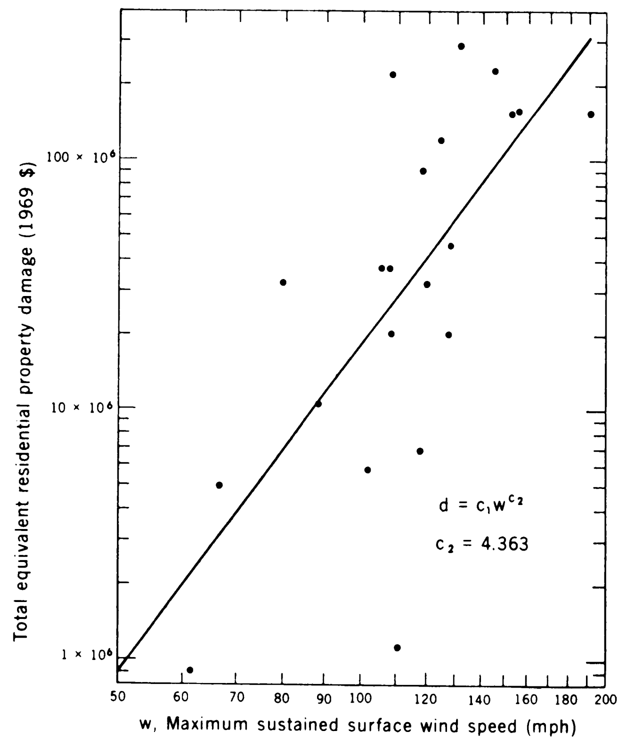

Fig. 3 (upper left). Probability distributions on 12-hour wind changes for the seeded and unseeded hurricane. Fig. 4 (right). Property damage plotted against maximum sustained wind speed. Fig. 5 (lower left). Probability distributions on property damage for the seeded and unseeded hurricane.

with the uncertainty about the natural change in the hurricane over a 12-hour period, we obtain a probability distribution for w' that is normal with a mean of 110 percent and a standard deviation of 18.6 percent:

$$P(w'|\mathrm{H}_3) = f_N(110\%, 18.6\%) \quad (3)$$

We have now developed probability distributions for the wind speed w' over a 12-hour period following the initiation of seeding for each of the three hypotheses. To obtain the probability distribution for w' that represents present information about the change in a seeded hurricane, we multiply each of the above distributions by the probability that is presently assigned to each of the hypotheses being true and sum over the three hypotheses:

$$P(w') = \sum_{i=1}^{3} P(w'|\mathrm{H}_i)P(\mathrm{H}_i) \quad (4)$$

Assigning Probabilities to the Hypotheses

The last element in developing a probability distribution for w' is to assign the probabilities $P(\mathrm{H}_1)$, $P(\mathrm{H}_2)$, and $P(\mathrm{H}_3)$. These probabilities should take into account both present meteorological information and meteorological information before the results of the 1969 Debbie experiments. The models we have just constructed allow us to examine the effect of experimental observations, such as the Debbie results, in revising the probabilities assigned to the three hypotheses. If a wind speed $w' = u$ has been observed after a seeding experiment, the posterior probabilities $P(\mathrm{H}_i|u)$ are related to the probabilities $P(\mathrm{H}_i)$ assigned before the experiment by Bayes' equation (*5, 16, 17*):

$$P(\mathrm{H}_i|u) = \frac{P(u|\mathrm{H}_i)P(\mathrm{H}_i)}{P(u)} \quad (5)$$

where the denominator is

$$P(u) = P(w' = u) = \sum_{i=1}^{3} P(w' = u|\mathrm{H}_i)P(\mathrm{H}_i) \quad (6)$$

The extension to several independent experiments is straightforward. The Debbie results are considered as two independent experiments in which reductions of 31 and 15 percent in wind speed were observed over a 12-hour period. The posterior probabilities assigned to the hypotheses are computed by multiplying together the appropriate values of two normal probability density functions. The probability density function for the Debbie results if hypothesis H_i is true, $P(u_1 = 69 \text{ percent}, u_2 = 85 \text{ percent}|\mathrm{H}_i)$, is

$$P(69\%, 85\%|\mathrm{H}_1) = 1.50 \times 2.14 = 3.21$$
$$P(69\%, 85\%|\mathrm{H}_2) = 0.372 \times 1.64 = 0.61$$
$$P(69\%, 85\%|\mathrm{H}_3) = 0.195 \times 0.886 = 0.173 \quad (7)$$

These numbers can be used to compute the posterior probabilities appropriate after the Debbie results from any set of probabilities assigned to the hypotheses before the Debbie results were known. For example, suppose that before the Debbie experiments the three hypotheses H_1, H_2, and H_3 were considered to be equally likely, that is, each had a probability of 1/3. Then, after the Debbie results are incorporated through Bayes' equation, the corresponding posterior probabilities assigned to the hypotheses are

$$P(\mathrm{H}_1|\text{Debbie}) = \frac{3.21 \times 1/3}{3.21 \times 1/3 + 0.61 \times 1/3 + 0.173 \times 1/3} = .81$$
$$P(\mathrm{H}_2|\text{Debbie}) = .15$$
$$P(\mathrm{H}_3|\text{Debbie}) = .04 \quad (8)$$

However, meteorologists did not believe that H_1, H_2, and H_3 were equally likely before the Debbie experiments. They thought that seeding was unlikely to have any effect but that, if seeding did have an effect, it was more likely to be a reduction in wind speed than an increase, because a reduction was expected from both the Simpson theory and the computer model studies. Further, the four field experiments that were conducted before Debbie all led to no change or to reductions in the maximum wind speeds (*1*).

We determined probability assignments for the three hypotheses to reflect present information by two conditions: (i) Before Debbie, meteorologists believed that H_1 was more likely than H_3 if seeding had any effect on a hurricane. (ii) Since Debbie, meteorologists believe that H_1 and H_2 are equally likely.

These conditions led us to use the probabilities

$$\begin{aligned} P(H_1) &= .49 \\ P(H_2) &= .49 \\ P(H_3) &= .02 \end{aligned} \qquad (9)$$

in our analysis. These posterior probabilities correspond to the pre-Debbie probabilities

$$\begin{aligned} P(H_1) &= .15 \\ P(H_2) &= .75 \\ P(H_3) &= .10 \end{aligned} \qquad (10)$$

This set of probability assignments implies that prior to Debbie the odds were 3 to 1 that seeding would have no effect but that, if seeding did have an effect, the odds were 3 to 2 for wind reduction rather than wind intensification. Since the Debbie results, the chance of seeding causing an average intensification of hurricanes is assessed at 1 in 50, and the "null" hypothesis, H_2, of no effect and the "beneficial" hypothesis, H_1, of an average reduction are judged equally likely.

The probability assignments (Eq. 9) representing present information were reviewed with Project Stormfury officials before being used in the analysis. However, the results of the analysis are not particularly sensitive to the specific numbers, as we discuss below.

Probability Distributions on Wind Speed

We now can compute the probability distributions on wind speed for the seeding and not-seeding alternatives (from Eqs. 1–4 and Eq. 9). These distributions are plotted in Fig. 3 as complementary cumulative distribution functions. By reading the ordinate values corresponding to an initial wind intensity of 100 percent, we find that the probability assigned to intensification if a hurricane is seeded is .36; if the hurricane is not seeded, the probability is .50. The probability of intensification by 10 percent or more is .18 if a hurricane is seeded and .26 if it is unseeded. For any particular wind speed, the probability that this speed will be exceeded is always greater if the hurricane is unseeded than if it is seeded because the complementary cumulative distribution function for the not-seeding alternative is always above the curve for the seeding alternative. This result is called stochastic dominance of the seeding alternative.

We have now specified the uncertainties about the outcome of the decision to seed. The same methods could be applied if the outcome were specified by several variables rather than simply by the relative change in maximum sustained wind speed. Much of the uncertainty in the outcome is the result of uncertainty about the natural change in hurricane behavior, not about the effect of seeding. This characteristic holds even more strongly if other aspects of hurricane behavior are examined, such as the trajectory of a hurricane or the precipitation it generates. Although it is considered unlikely that seeding would have a significant effect on these features of hurricanes, substantial variations may occur from natural causes.

The uncertainty about the natural behavior of a hurricane makes the issue of government responsibility of paramount importance. The intensification after seeding illustrated in Fig. 1 is a distinct possibility. Even if further experiments confirm that the "beneficial" hypothesis, H_1, is true, there would still be about one chance in ten that a seeded hurricane will intensify by 10 percent or more. Meteorological advances and improved computer models may eventually allow many of the natural changes in a hurricane to be predicted accurately, but this capability may require many years to achieve.

Wind Changes and Property Damage

The winds of a hurricane cause property damage directly and indirectly, the latter by creating a high storm tide that can flood low-lying coastal areas. The data available for past hurricanes do not distinguish wind and storm-tide damage; consequently, a detailed basis is lacking for a causal model relating wind and property damage. In our analysis, we assumed a general power law of the form

$$d = c_1 w^{c_2} \qquad (11)$$

where d is property damage in millions of dollars, w is the maximum sustained wind speed in miles per hour, and c_1 and c_2 are empirical constants to be determined from historical data on hurricanes. We estimated c_2 from data obtained from the American Red Cross on residential damage from 21 hurricanes. Since the Red Cross data were available for counties, we could isolate the damage caused by precipitation-induced flooding rather than by the wind or the storm tide by assuming that such damage occurred well inland. (The Red Cross data are the only statistics available that permit even this crude distinction between causes of damage.) Corrections for construction cost inflation and population growth were included, and c_2 was determined as 4.36 by a linear least-squares fit of the logarithms (Fig. 4). Thus, a change in the wind speed by a factor x implies a change in property damage by the factor x to the power 4.36. If x is 0.85, corresponding to a 15 percent reduction in maximum wind speed, the corresponding reduction in property damage is 51 percent (*18*).

The approximations of this method and the limited data indicate that broad limits are appropriate in a sensitivity analysis. If c_2 is 3, the reduction in damage corresponding to a 15 percent reduction in wind speed is 39 percent; if c_2 is 6, the corresponding damage reduction is 62 percent.

Since the probability assignments to wind changes were made on relative rather than absolute changes in maximum sustained wind speeds, the scaling factor c_1 can be assigned as the last step in the analysis. We assume a nominal hurricane whose maximum wind speed at the time of the seeding decision is such that, if no change occurs in the 12 hours before landfall, the property damage will be \$100 million. The analysis for a more or a less severe hurricane can be obtained by a suitable change in scale factor (*19*).

Using this relationship between property damage and maximum wind speed, we can develop the probability distributions for property damage for the nominal hurricane, whether seeded or unseeded. Figure 5 shows that the seeding alternative stochastically dominates the not-seeding alternative: the probability of exceeding a particular amount of property damage is always greater if the hurricane is not seeded than if it is seeded. Hence, if property damage is the criterion, the better alternative is to seed.

Further Analysis of the Decision to Seed

The decision to seed is shown in the form of a decision tree in Fig. 6. The decision to seed or not to seed is shown at the decision node denoted by the small square box; the consequent resolution of the uncertainty about wind change is indicated at the chance nodes denoted by open circles. For expository clarity and convenience, especially in the later stages of the analysis, it is convenient to use discrete approximations to the probability distributions for wind change (*20*) (Table 1).

As a measure of the worth of each alternative we can compute the expected loss for each alternative by multiplying the property damage for each of the five possible outcomes by the probability that the outcome will be achieved and summing over the possible consequences. The expected loss for the seeding alternative is \$94.33 million (including a cost of \$0.25 million to carry out the seeding); the expected loss for the not-seeding alternative is \$116 million; the difference is \$21.67 million or 18.7 percent.

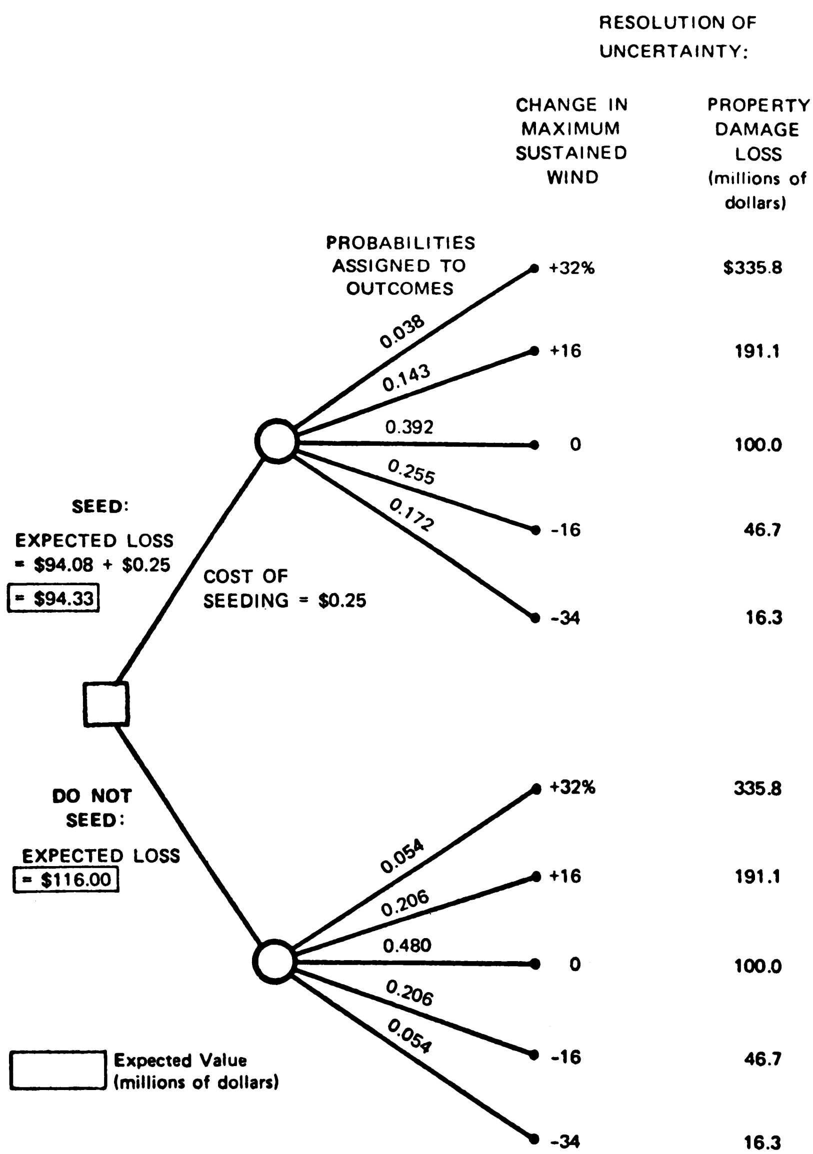

Fig. 6. The seeding decision for the nominal hurricane.

These results should be examined to see how much they depend on the specific assumptions in the model. Stochastic dominance is a general result that does not depend on the specific form of the relationship between property damage and maximum wind speed (see Eq. 11); rather, it depends on the probabilities assigned to hypotheses H_1, H_2, and H_3. The probability of H_3 must be raised to .07 before stochastic dominance no longer holds. Even if the probability of H_3 is raised much higher, seeding still results in the least expected property damage. If $P(H_1)$ is .40, $P(H_2)$ is .40, and $P(H_3)$ is .20, the expected loss for the seeding alternative is \$107.8 million—7 percent less than for the not-seeding alternative. Variation of the exponent c_2 from 3 to 6 does not change the decision: if c_2 is 3, the expected property damage with seeding is 14 percent less; if c_2 is 6, the expected reduction in damage is 22 percent. If the criterion of expected cost is replaced by a nonlinear utility function reflecting aversion to risk, the relative advantage of the seeding alternative is even greater (*21*). The results of extensive sensitivity analysis may be summarized as follows: The expected loss in terms of property damage appears to be about 20 percent less if the hurricane is seeded. Varying the assumptions of the analysis causes this reduction to vary between 10 and 30 percent but does not change the preferred alternative.

Government Responsibility

The analysis in the section above indicates that, if minimizing the expected loss in terms of property damage (and the cost of seeding) is the only criterion, then seeding is preferred. However, an important aspect of the decision—the matter of government responsibility—has not yet been included in the analysis. We have calculated a probability of .36 that a seeded hurricane will intensify between seeding and landfall and a probability of .18 that this intensification will be at least 10 percent. This high probability is largely the result of the great natural variability in hurricane intensity. It is advisable to consider both the legal and the social consequences that might occur if a seeded hurricane intensified.

The crucial issue in the decision to seed a hurricane threatening a coastal area is the relative desirability of reducing the expected property damage and assuming the responsibility for a dangerous and erratic natural phenomenon. This is difficult to assess, and to have a simple way of regarding it we use the concept of a government responsibility cost, defined as follows. The government is faced with a choice between assuming the responsibility for a hurricane and accepting higher probabilities of property damage. This situation is comparable to one of haggling over price: What increment of property-damage reduction justifies the assumption of responsibility entailed by seeding a hurricane? This increment of property damage is defined as the government responsibility cost. The government responsibility cost is a means of quantifying the indirect social, legal, and political factors related to seeding a hurricane. It is distinguished from the direct measure—property damage—that is assumed to be the same for both modified and natural hurricanes with the same maximum sustained wind speed.

We define the government responsibility cost so that it is incurred only if the hurricane is seeded. It is conceivable that the public may hold the government responsible for not seeding a severe hurricane, which implies that a responsibility cost should also be attached to the alternative of not seeding. Such a cost would strengthen the implication of the analysis in favor of permitting seeding.

The assessment of government responsibility cost is made by considering the seeding decision in a hypothetical situation in which no uncertainty is present. Suppose the government must choose between two outcomes:

1) A seeded hurricane that intensifies 16 percent between the time of seeding and landfall.

2) An unseeded hurricane that intensifies more than 16 percent between the time of seeding and landfall. The property damage from outcome 2 is x percent more than the property damage from outcome 1.

If x is near zero, the government will choose outcome 2. If x is large, the government will prefer outcome 1. We then adjust x until the choice becomes very difficult; that is, the government is indifferent to which outcome it receives. For example, the indifference

Table 1. Probabilities assigned to wind changes occurring in the 12 hours before hurricane landfall. Discrete approximation for five outcomes.

Interval of changes in maximum sustained wind	Representative value in discrete approximation (%)	Probability that wind change will be within interval: If seeded	If not seeded
Increase of 25% or more	+32	.038	.054
Increase of 10 to 25%	+16	.143	.206
Little change, +10 to −10%	0	.392	.480
Reduction of 10 to 25%	−16	.255	.206
Reduction of 25% or more	−34	.172	.054

point might occur when x is 30 percent. An increase of 16 percent in the intensity of the nominal hurricane corresponds to property damage of $191 million, so that the corresponding responsibility cost defined by the indifference point at 30 percent is (.30) ($191 million), or $57.3 million. The responsibility cost is then assessed for other possible changes in hurricane intensity.

The assessment of government responsibility costs entails considerable introspective effort on the part of the decision-maker who represents the government. The difficulty of determining the numbers does not provide an excuse to avoid the issue. Any decision or policy prohibiting seeding implicitly determines a set of government responsibility costs. As shown in the last section, seeding is the preferred decision unless the government responsibility costs are high.

Let us consider an illustrative set of responsibility costs. The government is indifferent, if the choice is between:

1) A seeded hurricane that intensifies 32 percent and an unseeded hurricane that intensifies even more, causing 50 percent more property damage.

2) A seeded hurricane that intensifies 16 percent and an unseeded hurricane that causes 30 percent more property damage.

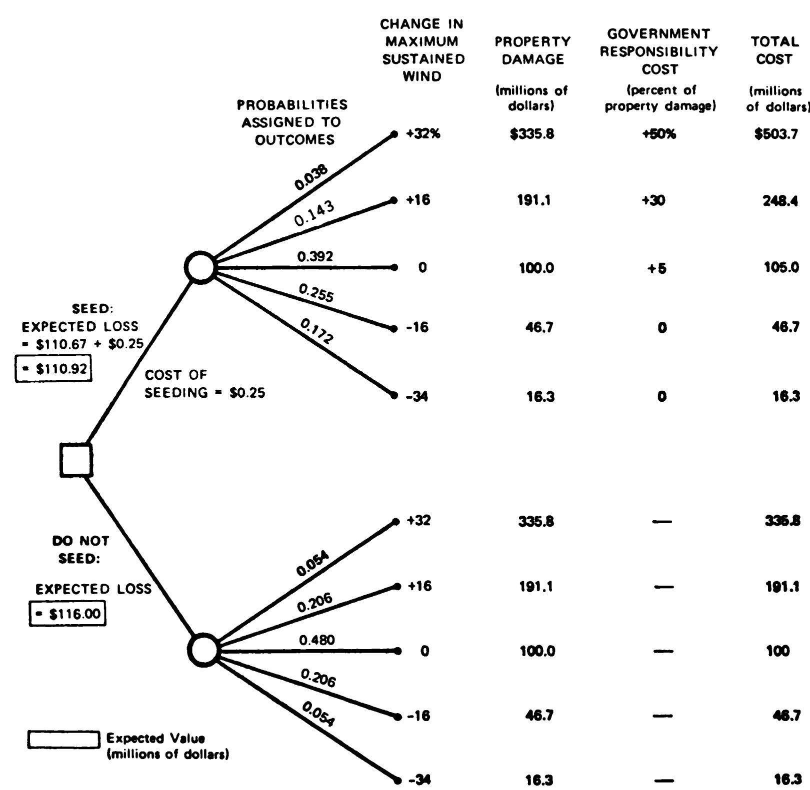

Fig. 7. The seeding decision for the nominal hurricane (government responsibility cost included).

3) A seeded hurricane that neither intensifies nor diminishes (0 percent change in the maximum sustained wind speed after the seeding) and an unseeded hurricane that intensifies slightly, causing 5 percent more property damage.

4) A seeded hurricane that diminishes by more than 10 percent and an unseeded hurricane that diminishes by the same amount. (If the hurricane diminishes after seeding, everyone agrees that the government acted wisely; thus, responsibility costs are set at zero.)

The analysis of the seeding decision with these government responsibility costs included is diagramed in Fig. 7. Even with these large responsibility costs, the preferred decision is still to seed.

The responsibility costs needed to change the decision are a substantial fraction of the property damage caused by the hurricane. For the $100-million hurricane chosen as the example for this section, the average responsibility cost must be about $22 million to change the decision. If the hurricane were in the $1-billion class, as Camille (1969) and Betsy (1965) were, an average responsibility cost of $200 million would be needed. In other words, an expected reduction of $200 million in property damage would be foregone if the government decided not to accept the responsibility of seeding the hurricane.

The importance of the responsibility issue led us to investigate the legal basis for hurricane seeding in some detail. These investigations were carried out by Gary Widman, Hastings College of the Law, University of California. A firm legal basis for operational seeding apparently does not now exist. The doctrine of sovereign immunity provides the government only partial and unpredictable protection against lawsuits, and substantial grounds for bringing such lawsuits might exist (*22*). A better legal basis for government seeding activities is needed before hurricane seeding could be considered other than as an extraordinary emergency action. Specific congressional legislation may be the best means of investing a government agency with the authority to seed hurricanes threatening the coast of the United States.

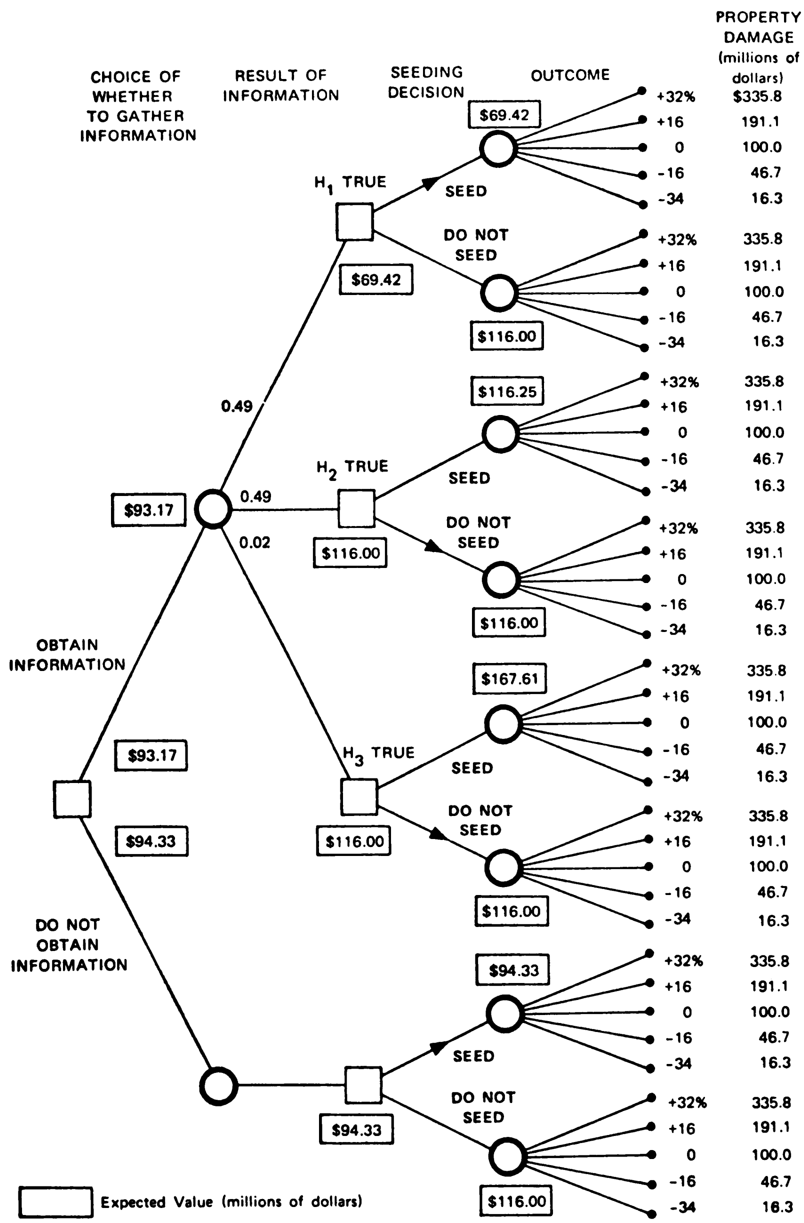

Fig. 8. Expected value of the clairvoyant's information—which hypothesis describes the effect of seeding? (There is no government responsibility cost.)

Value of Information

One of the most important concepts in decision analysis is the value of information: How much it would be worth to make the decision after rather than before uncertainty is resolved? In the case of hurricane modification, how much should be the government pay to learn which of the three hypotheses, H_1, H_2, or H_3, is actually true (*23*)? We imagine that the government has access to a clairvoyant who has this information and is willing to sell it to the government, if he is paid before he makes the information available. It is easiest to understand the calculation in terms of the decision to seed one hurricane threatening a coastal area.

Let us consider the choice between the two decision situations shown in Fig. 8. The government can choose to buy the information and make the decision after it has learned which hypothesis is true, or it can choose not to buy the information and can make the seeding decision on the basis of the present uncertainty.

Let us, for the moment, consider only property damage and the cost of seeding and disregard government responsibility costs. If H_1 is true, the preferred decision is to seed because the expected loss is \$69.42 million compared with a loss of \$116.00 million for the alternative of not seeding. If H_2 is true, then by choosing not to seed, the government saves the \$0.25-million cost of seeding; the loss expected from property damage is the same for both alternatives: \$116.00 million. If H_3 is true, seeding is a poor choice; the expected loss from property damage is \$167.61 million, \$51.61 million more than for the alternative of not seeding. At the present time, the government does not know what the clairvoyant will say, but probabilities have been assigned to his answers:

$$P(H_1) = .49 \quad P(H_2) = .49 \quad P(H_3) = .02 \tag{12}$$

The expected loss corresponding to the decision situation in Fig. 8 is then computed by multiplying the probability of each of the clairvoyant's answers by the expected loss associated with that answer and summing over the three possible answers:

$$(.49)(\$69.42) + (.49)(\$116.00) + (.02)(\$116.00) = \$93.17 \text{ million} \tag{13}$$

Comparing this with the expected loss for the best alternative (seed) without the clairvoyant's information, which was \$94.33 million, we see that it is \$1.16 million less. This difference represents the expected value of the clairvoyant's information in allowing the government to make a better decision. It is a relatively small number compared with the expected losses because the information is not expected to be of much value—the probability assignments indicate that seeding is already a good idea. Without the clairvoyant's information the government should seed; with the clairvoyant's information, with probability .49, the government will save the cost of seeding (\$0.25 million), and with the low probability .02 it will avert the potentially disastrous intensification expected from H_3, saving \$167.61 million − \$116 million = \$51.61 million. By this reasoning we get the same answer as before for the value of information

$$(.49)(\$0.25) + (.02)(\$51.61) = \$1.16 \text{ million} \tag{14}$$

and we can see that the value is very sensitive to the small probability assigned to H_3.

Now suppose that the government responsibility costs assumed previously are included. The expected value of perfect information is then much higher because, if H_2 is true, the government responsibility costs can be saved by not seeding. If the decision without perfect information is to seed, the expected saving from engaging the clairvoyant is

$$(.49)(\$0.25 + \$23.28) + (.02)(\$51.61 + \$53.57) = \$13.63 \text{ million} \tag{15}$$

This figure represents 11.75 percent of the expected property damage if the alternative of not seeding is taken for the nominal hurricane.

The value of information largely derives from the fact that it allows the government to avoid the responsibility for seeding if seeding turns out to have no effect. The large increase over the value computed in Eq. 13 is due to the contribution of the government responsibility costs. Most of the increase of \$12.47 million, namely \$11.41 million, comes from the first or H_2 term.

The value of information depends on the extent to which the government is willing to assume responsibility for seeding a hurricane. If responsibility were not an issue, the government would seed operationally now, and information would have a comparatively low value in the context of this decision. The value of information is greatest when the government responsibility costs are large enough to make the decision essentially even between seeding and not seeding. Still higher responsibility costs cause the value of information to decrease (*24*).

Value of Further Seeding Experiments

The analysis of the value of a seeding experiment is similar to the determination of the value of the clairvoyant's information. The difference is that the resolution of uncertainty is only partial. The information obtained in the experiment is used in Bayes' equation (Eq. 5) to revise the prior probability assignments to the hypotheses. The original decision is then reevaluated with the posterior probabilities (Fig. 9). The result of the experiment is uncertain when the decision to experiment is made; consequently, the value of experimentation must be computed as an expectation over the possible posterior decision situations. The situation can be diagramed in tree form as shown in Fig. 9.

The analysis for two experimental seedings is given in Table 2 (*25*). The values assumed above for the govern-

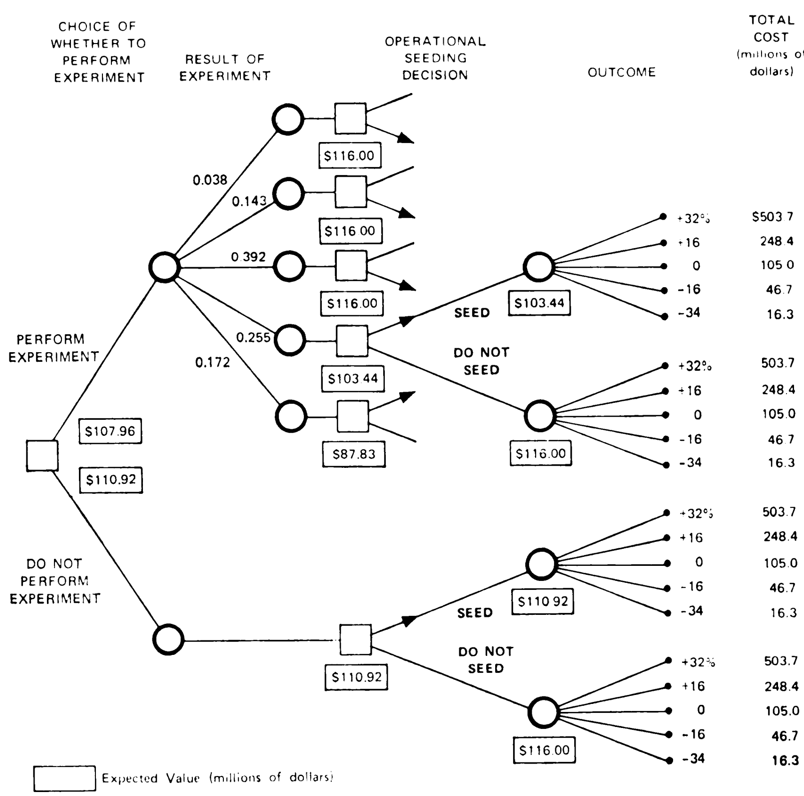

Fig. 9. Value of a seeding experiment (government responsibility cost included).

ment responsibility costs have been used. The expected value of the experiment in improving one operational seeding decision is \$5.39 million, slightly less than twice the value of a single experimental seeding and more than ten times the assumed experimental cost, \$0.50 million. This value represents 4.7 percent of the expected property damage if the alternative of not seeding is taken. In the discrete version used in the analysis, one of five possible values (see Table 1) is taken as representative of the observed change in hurricane intensity over a 12-hour period following seeding: − 34, − 16, 0, + 16, and + 32 percent. The order in which the results are obtained is not significant, and a total of 15 pairs of results could be obtained with two experiments (Table 2). These pairs might be placed in three groups: favorable, unfavorable, and mixed results. The probability of obtaining a pair of favorable results (− 34, − 34; − 34, − 16; − 34, 0; and − 16, − 16 percent) (*26*) in the two experimental seedings is .327; a pair of results in this group would provide substantial confirmation of hypothesis H_1.

For example, a repetition of the pair of results obtained with Debbie in 1969 (− 34, − 16 percent in the discrete approximation) would lead to posterior probabilities of .89 for H_1, .11 for H_2, and less than .005 for H_3. A probability of .075 is computed for a pair of strongly unfavorable results (0, + 32; + 16, + 16; + 16, + 32; + 32, + 32 percent); in this case the probability assigned to H_1 would be revised strongly downward. The remaining mixed pairs of results do not significantly confirm or deny H_1, and these results have a total probability of .595. Within this group a small probability (.055) is accorded to conflicting results in the two experiments (− 34, + 16; − 34, + 32; − 16, + 16; − 16, + 32 percent).

Another Approach to Determining the Value of Seeding Experiments

The preceding discussion indicates that the value of experiments is sensitive to the government responsibility costs that are assumed in the analysis. We may wish to determine the value of experiments in a different manner in which the issue of government responsibility is treated implicitly.

Table 2. Evaluation of a future experiment with two (independent) experimental seedings. Government responsibility cost is included.

Observed change in wind speed u_1	u_2	Prior probability of observation	Posterior probability of hypotheses H_1	H_2	H_3	Subsequent operational seeding decision expected values (million dollars): Loss with seeding alternative	Loss with the better alternative	Posterior value of perfect information
− 34	− 34	.0441	.97	.03	< .005	79.87	79.87	0.80
− 34	− 16	.1009	.89	.11	< .005	84.67	84.67	2.68
− 34	0	.1140	.77	.22	< .005	92.11	92.11	5.64
− 16	− 16	.0684	.69	.30	< .005	97.08	97.08	7.53
		.3274						
− 34	+ 16	.0324	.65	.34	.01	100.16	100.16	9.06
− 34	+ 32	.0078	.60	.37	.03	105.27	105.27	12.10
− 16	0	.1915	.49	.51	.01	110.25	110.25	12.78
− 16	+ 16	.0651	.34	.64	.02	120.07	116.00	13.05
0	0	.1610	.28	.70	.02	123.37	116.00	10.81
− 16	+ 32	.0167	.29	.65	.06	126.05	116.00	11.15
0	+ 16	.1229	.18	.79	.03	131.35	116.00	6.78
		.5974						
0	+ 32	.0332	.14	.77	.09	138.02	116.00	5.51
+ 16	+ 16	.0251	.10	.83	.07	138.62	116.00	3.98
+ 16	+ 32	.0145	.08	.75	.17	148.37	116.00	3.02
+ 32	+ 32	.0024	.05	.59	.36	165.72	116.00	1.98
		.0752						

Value of seeding decision with prior information	110.92
Expected value of seeding decision with seeding experiments	105.53
Value of experiment	5.39
Cost of experiment	0.50
Net expected value of experiment	4.89

Suppose that operational seeding will be permitted only after another successful result is obtained in a pair of experiments of the Debbie type. This approximation gives a lower bound to the value of experiments because only a successful experimental result is regarded as valuable. Even if wind reductions are not observed, knowledge gained about the effects of seeding may have implications for future successful operational seeding.

The probability of a favorable pair of results in two experimental seedings of a hurricane was computed as .327. If favorable experimental results are obtained and a subsequent hurricane is seeded operationally, the expected reduction in property-damage losses is \$37.88 million. Even if government responsibility costs are included, the reduction in expected losses is \$26.80 million. Since these reductions occur with a probability of .327, the expected value of the experiment in improving one operational decision is \$12.40 million if only the property damage is considered and \$8.77 million if the decrease in property damage is partially offset by the government responsibility costs. The figures \$8.77 million and \$12.40 million represent 7.6 and 10.7 percent, respectively, of the \$116-million property damage expected from the not-seeding alternative in the seeding decision for the nominal hurricane.

We see that the value of experiments is considerably higher than the values computed earlier. This difference results from the high responsibility costs implicit in the decision not to seed on the basis of present information. It may be a reasonable assumption that a bad outcome for the first seeding of a hurricane threatening a coastal area would have much less severe legal and social consequences if it were preceded by another successful experiment. Therefore, lowering the government responsibility costs may be appropriate after another successful field experiment.

Generalizing the Value of Additional Information

The preceding discussions are directed specifically toward updating our information about which hypothesis, H_1, H_2, or H_3, describes the effect of seeding on the maximum sustained wind speed of a hurricane. The analysis has been done for a single seeding decision for a moderately intense hurricane threatening a coastal area. Perfect information applies not only to a single hurricane but to all hurricanes that might be seeded operationally. The numerical results for the single nominal hurricane are summarized in the extreme left column of Table 3 and are extended to multiple hurricanes in the remaining columns.

Even if only half the hurricanes could be seeded because of tactical considerations having to do with precipitation, hurricane trajectory, and so on, the expected annual benefit from perfect information is $26 million. If we assume that only half the hurricanes could be seeded, and discount the expected benefits of perfect information for all future hurricane seasons at a discount rate of 7 percent, we arrive at $370 million. This figure represents the value of a "perfect" experiment that would determine whether H_1 is true.

A single repetition of the 1969 Hurricane Debbie experiment has an expected value of $5.39 million in the context of the nominal hurricane, or about 4.7 percent of expected property damage. For the decision to seed a single hurricane in the billion-dollar range, the expected value of the experiment is ten times as high, about $50 million. For one hurricane season the value is 4.7 percent of $220 million, or $10.2 million (it is assumed again that various tactical considerations might preclude seeding in half of the cases). For all future hurricane seasons, with a discount rate of 7 percent, the value is $146 million compared with an experimental cost of about $500,000. The benefit to cost ratio is therefore about 300. Even if only a single hurricane season is considered, the expected benefits are 20 times greater than the cost of the experiment and ten times the present annual budget for Project Stormfury.

Table 3. Summary of the value of additional information on the effect of seeding. Only the 50 percent of hurricanes that are assumed to be possible candidates for seeding on the basis of tactical considerations are considered. If all hurricanes are assumed to be candidates for operational seeding, the figures of the last two columns should be doubled.

Item	Nominal hurricane used in analysis		Single hurricane season (million dollars)	All future hurricane seasons, discounted at 7% (million dollars)
	Million dollars	Percentage		
Expected property damage without seeding	116.0	100	220.0	3142
Expected value of perfect information	13.6	11.8	26.0	370
Expected value of a field experiment consisting of two experimental seedings	5.4	4.7	10.2	146
Expected value of field experiments:*				
With government responsibility costs	8.8	7.6	16.6	238
Government responsibility costs = 0	12.4	10.7	23.5	335

* If it is assumed that prior operational seeding is not permitted.

Experimental Capability Decision

The occurrence of hurricanes is a random phenomenon. Therefore, it is uncertain whether there will be an opportunity for an experimental seeding before the arrival of a threatening storm that might be operationally seeded. Opportunities for experimental seeding have been scarce. In the last few years there have been only six experimental seedings, and these have been conducted on three hurricanes, Esther (1961), Beulah (1963), and Debbie (1969) (*7*). Experimental seedings have been limited to a small region of the Atlantic Ocean accessible to aircraft based in Puerto Rico, and few hurricanes have passed through this region.

There are many other regions of the ocean where hurricanes might be found that satisfy the present criterion for experimental seeding—that is, the hurricane will be seeded only if the probability is less than .10 that it will come within 50 miles of a populated land area within 18 hours after seeding. However, a decision to expand the present experimental capability of Project Stormfury would need to be made well before the experiment itself. Whereas the seeding itself requires only that an aircraft be fitted with silver iodide pyrotechnic generators, the monitoring of the subsequent development of the hurricane requires other aircraft fitted with the appropriate instrumentation. The requirements in equipment, crew training, and communications and support facilities are substantial. In addition, permission may be needed from nations whose shores might be threatened by the seeded hurricane. The experimental decision, then, involves an investment in the capability to perform an experimental seeding. Whether an experiment is performed depends on the uncertain occurrences of hurricanes in the experimental areas.

The expected time before another experimental opportunity for Project Stormfury's present capability is about one full hurricane season. There was no opportunity during 1970. Preliminary estimates of the cost of a capability to seed hurricanes in the Pacific are about $1 million (*27*). The incidence of experimentally seedable hurricanes in the Pacific appears to be more than twice that in the Atlantic (*28*). Therefore, it appears advisable to develop a capability to conduct experimental hurricane seeding in the Pacific Ocean since the benefits expected from this capability outweigh the costs by a factor of at least 5 (*29*).

Conclusions from the Analysis

The decision to seed a hurricane imposes a great responsibility on public officials. This decision cannot be avoided because inaction is equivalent to a decision not to permit seeding. Either the government must accept the responsibility of a seeding that may be perceived by the public as deleterious, or it must accept the responsibility for not seeding and thereby exposing the public to higher probabilities of severe storm damage.

Our report to the National Oceanic and Atmospheric Administration recommended that seeding be permitted on an emergency basis. We hope that further experimental results and a formal analysis of the tactical decision to seed a particular hurricane will precede the emergency. However, a decision may be required before additional experimental or analytical results are available. A hurricane with the intensity of Camille threatening a populous coastal area of the United States would confront public officials with an agonizing but unavoidable choice.

The decision to seed hurricanes can not be resolved on strictly scientific grounds. It is a complex decision whose uncertain consequences affect many people. Appropriate legal and political institutions should be designated for making the hurricane-seeding decision, and further analysis should be conducted to support these institutions in carrying out their work.

Role of Decision Analysis

The results of a decision analysis depend on the information available at the time it is performed. Decision analysis should not be used to arrive at a static recommendation to be verified by further research, rather it should be used as a dynamic tool for making necessary decisions at any time. Various sensitivity analyses included here indicate how new information might be expected to influence policy recommendations. However, the advent of a severe hurricane will necessitate a decision on the basis of the information then available.

The analysis of hurricane modification points up a difficulty that is common in public decision-making on complex technological issues. When the consequences of deploying new technology are uncertain, who will make the choice? While many individuals or groups may share responsibility, decision analysis conceptually separates the roles of the executive decision-maker, the expert, and the analyst. The analyst's role is to structure a complex problem in a tractable manner so that the uncertain consequences of the alternative actions may be assessed. Various experts provide the technical information from which the analysis is fashioned. The decision-maker acts for society in providing the basis for choosing among the alternatives. The analysis provides a mechanism for integration and communication so that the technical judgments of the experts and the value judgments of the decision-maker may be seen in relation to each other, examined, and debated. Decision analysis makes not only the decision but the decision process a matter of formal record. For any complex decision that may affect the lives of millions, a decision analysis showing explicitly the uncertainties and decision criteria can and should be carried out.

References and Notes

1. R. H. Simpson and J. S. Malkus, *Sci. Amer.* **211**, 27 (Dec. 1964).
2. R. C. Gentry, *Science* **168**, 473 (1970).
3. Now incorporated in the National Oceanic and Atmospheric Administration.
4. R. A. Howard, *Proceedings of the Fourth International Conference on Operational Research* (Wiley, New York, 1966).
5. ———, *IEEE Trans. Syst. Sci. Cybern.* *4* (1968), p. 211.
6. A detailed discussion of the research is to be found in the project's final report [D. W. Boyd, R. A. Howard, J. E. Matheson, D. W. North, *Decision Analysis of Hurricane Modification* (Project 8503, Stanford Research Institute, Menlo Park, Calif., 1971)]. This report is available through the National Technical Information Service, U.S. Department of Commerce, Washington, D.C., accession number COM-71-00784.
7. The meteorological information leading to this approximation is discussed in detail in the SRI project final report (*6*), especially appendix B. Meteorologists connected with Project Stormfury believe it highly improbable that seeding will cause any substantial change in the course of the hurricane, and other important consequences of seeding are not foreseen at this time. We wish to stress that our role in the decision analysis of hurricane modification has been to provide the methodology for analyzing a complex decision with uncertain consequences. The specific assumptions have been provided by the hurricane meteorologists associated with Project Stormfury and by other experts in relevant fields. Because of space limitations these assumptions cannot be discussed in detail in this article; the interested reader is advised to consult the project's final report or communicate directly with the authors. The type of seeding is assumed to be the same as that used in the Hurricane Debbie experiments: massive multiple seeding of the clouds in the outer eyewall region with silver iodide. During September 1971, Project Stormfury conducted seeding experiments of a different type on Hurricane Ginger (R. C. Gentry, internal communication, National Oceanic and Atmospheric Administration, October 1971). The Ginger experiment involved the seeding of clouds in the rain bands well outside the eyewall region. This "rainsector" experiment was selected because Ginger had a large and poorly formed eyewall and was judged not to be a good subject for eyewall-region seeding. Although some changes in cloud structure and wind field occurred at a time when they might have been caused by seeding, these changes were minor compared with the dramatic changes that occurred in Hurricane Debbie after seeding. Because of the difference in type of seeding, the Ginger results do not imply a need for revision of the analysis or data presented in this article.
8. In some hurricanes, such as Diane (1955) and Camille (1969), precipitation-induced inland flooding has also been an important cause of property damage. Seeding might cause some increase in precipitation. In considering the policy decision to permit seeding we ignored these precipitation effects, but they might sometimes be important in the decision to seed a specific hurricane.
9. Throughout the analysis it is assumed that seeding would be authorized and carried out by some agency of the federal government.
10. C. Holliday, *Technical Memorandum WBTM SR-45* (Environmental Science Service Administration, Washington, D.C., 1969).
11. The details of the derivation of this probability distribution are given in appendix B of (*6*). The indirect approach of using the Holliday relation combined with pressure-change observations was first suggested by R. C. Sheets of the National Hurricane Research Laboratory.
12. S. L. Rosenthal, *Technical Memorandum ERLTM-NHRL 88* (Environmental Science Service Administration, Washington, D.C., 1970). See also *Project Stormfury Annual Report 1970* (National Hurricane Research Laboratory, Miami, 1971).
13. A probability distribution on an uncertain quantity x will be denoted $P(x)$ whether x takes on discrete or continuous values. If x is discrete, $P(x)$ will be the probability mass function; if x is continuous, $P(x)$ will be the probability density function. A probability distribution of the normal or Gaussian family specified by its mean m and standard deviation σ will be denoted $f_N(m, \sigma)$.
14. The average effect of seeding and the fluctuation from the average may be regarded as (independent) normal random variables whose sum represents the effect of seeding on a specific hurricane. According to well-known results in probability theory, this sum will be normally distributed with a mean equal to the sum of the two means and a standard deviation equal to the square root of the sum of the squares of the two standard deviations.
15. The effect of seeding and the natural change in the storm are described as independent normal random variables and the total change is their sum. The independence assumption is judged an appropriate summary of present knowledge; sensitivity to this assumption is examined in (*6*). Important assumptions such as this one were reviewed with Project Stormfury meteorologists. A letter to us from R. C. Gentry (October 1970) stated, "while seeding may affect different hurricanes by different amounts, we are not yet prepared to predict these differences." The assumption of independence does not deny that there may be a relationship between the natural change occurring in a hurricane and the effect of modification. When information about this relationship becomes available, it should be incorporated into the analysis and the independence assumption should be withdrawn.
16. H. Raiffa, *Decision Analysis: Introductory Lectures on Choices Under Uncertainty* (Addison-Wesley, Reading, Mass., 1968); M. Tribus, *Rational Descriptions, Decisions, and Designs* (Pergamon, New York, 1969).
17. D. W. North, *IEEE Trans. Syst. Sci. Cybern.* *4* (1968), p. 200.
18. The details of the calculation of c_2 are given in (*6*). Similar relationships between maximum sustained wind speed and property damage have been stated by other investigators [R. L. Hendrick and D. G. Friedman, in *Human Dimensions in Weather Modification*, W. R. Derrick Sewell, Ed. (Univ. of Chicago Press, Chicago, 1966), pp. 227–246]. In November 1971, D. G. Friedman communicated to us some results from analyzing insurance claim data. He finds an exponent of 6.7; this value would lead to much larger reductions in property damage than were assumed in our analysis. Other investigators have suggested an equation of the form $d = c_1 (w - w_0)^2$, where c_1 and w_0 are empirical constants (R. C. Gentry, private communication). This equation would give results essentially equivalent to ours.
19. This procedure is an approximation, which depends on the fact that seeding costs are small compared with costs of property damage.
20. It is shown in (*6*) that the results are not sensitive to the discrete approximation.
21. For example, if an exponential utility function with a risk aversion coefficient of $\gamma = 0.001$ is used, the difference between the certain equivalents for the two alternatives increases from \$21.67 million to \$24.2 million. Because of stochastic dominance, any risk attitude will always leave seeding as the preferred alternative. Further discussion on risk preference may be found in (*5*) and (*17*).
22. These issues are discussed in detail in appendixes E and F of (*6*).
23. In answering this question we assume that the government is willing to pay up to \$1 to avoid \$1 of property damage.
24. It is possible for the responsibility costs to be so high that a hurricane would not be seeded even if it were certain that H_1 is true. This amount of responsibility cost implies that the government would prefer an unseeded hurricane to a seeded hurricane that caused only half as much property damage.
25. For these calculations a system of computer programs for evaluating large decision trees, developed by W. Rousseau of Stanford Research Institute, was used.
26. These discrete outcomes correspond to a reduction of 10 percent or more. The discrete approximation simplifies the analysis by restricting the number of possible experimental results. Earlier we considered the revision of probabilities based on the results of the 1969 Hurricane Debbie experiments. There the discrete approximation was not used, but it would have given equivalent results.
27. R. C. Gentry, personal communication. In arriving at this figure it was assumed that military aircraft based in the Pacific could be used in the seeding.
28. *Project Stormfury Annual Report 1968* (National Hurricane Research Laboratory, Miami, 1969).
29. The details of this calculation are given in (*6*).

30. This article summarizes research performed for the National Oceanic and Atmospheric Administration, U.S. Department of Commerce, contract 0-35172; the project leader is D. W. North. The authors acknowledge the substantial contribution of Dr. Dean W. Boyd. The authors also wish to acknowledge Professor Gary Widman of Hastings College of the Law, University of California, San Francisco, for legal research supporting the project and Dr. Cecil Gentry, Dr. Robert Simpson, Dr. Joanne Simpson, and many others who have been associated with Project Stormfury for their assistance and cooperation. The findings and conclusions presented are the sole responsibility of the authors and do not necessarily reflect the views of the U.S. government or any of the individuals mentioned above.

Letters

Seeding Hurricanes

Howard, Matheson, and North, in their article "The decision to seed hurricanes" (16 June, p. 1191), provide a good framework for an initial consideration of this important subject. They do not, however, include the effect of seeding on the hurricane rainfall rate—only the effects on the maximum sustained wind and on the wind-related storm tide.

While these latter effects may be paramount for coastal areas, in the light of the floods that accompanied hurricane Agnes, the storm rainfall should also be considered a decision factor when further studies are made of the seeding of hurricanes that threaten coastal areas. As seeding could conceivably increase the storm rainfall, both at the coastline and inland, the increased damage from flooding would then have to be balanced against the hoped-for reduction in damage from wind and storm tide.

Research on the control of hurricane direction, as well as on the reduction of wind intensity, appears indicated. If directional control were feasible—and, for example, some control of the rate of release of latent heat in different storm sectors is technically available now—this would be an attractive option in the case of storms approaching coastal areas.

The question of loss of life in seeded hurricanes, not covered in the article, must eventually be faced. The parallel question of seeding in war (News and Comment, 16 June, p. 1216), could also benefit from the same kind of rational and orderly analysis as that begun by Howard *et al.* To be fully useful, a study should attempt to separate the military from the civil effects, and the value judgments should be founded on an analysis of war as a moral problem.

BERNARD A. POWER
255 Touzin Avenue,
Dorval, Quebec, Canada

The article by Howard, Matheson, and North is an elegant decision-making analysis (within a Bayesian framework) that considers the consequences of both property damage and government responsibility of seeding versus nonseeding of hurricanes in terms of change in maximum sustained surface wind. Within this somewhat constrained analysis (surface wind as a surrogate for a complex physical phenomenon, property damage, and government responsibility are the only effects considered), a thorough range of possible outcomes is examined, including the three key hypotheses that seeding is beneficial, ineffective, or detrimental to the goal of reducing the social cost of hurricanes. The central conclusion is that "On the basis of present information, the probability of severe damage is less if a hurricane is seeded" and that seeding should be permitted on an emergency basis and encouraged on an experimental basis. But beyond this recommendation, the analysis itself is suggested as a model for "any complex decision that may affect the lives of millions, a decision analysis showing explicitly the uncertainties and decision criteria [that] can and should be carried out."

Among social scientists working at the boundaries of atmospheric science, the Howard *et al.* analysis has been received with critical enthusiasm. Some 6 years ago, along with Sewell, I suggested a process of analyzing social impacts akin to this analysis (*1*) and undertook with Julian and Sewell (*2*) a modest field survey to determine the expectations of leading atmospheric scientists about the viability of a range of weather modification technologies.

In the spirit of decision analysis, I question the use of the Howard *et al.* analysis and offer three alternative hypotheses to be used in the context of the current social, political, and scientific milieu when such questions as the decision to seed hurricanes are being dealt with.

1) Hypothesis H_1. Decision analysis is a rational method of analysis that systematically precludes in nonrandom fashion significant aspects of the problem, because these aspects are either not known, poorly understood, have low a priori estimates of probability, or seem inappropriate to the terms of reference.

2) Hypothesis H_2. Decision analysis is a rational method of analysis which will be used in an "arational" way.

3) Hypothesis H_3. Decision analysis is a rational method of analysis employed rationally for amoral purposes.

The first hypothesis emphasizes the problem of where to make the cut in systems analysis. Howard *et al.* have so constrained their analysis as to ignore the beneficial and detrimental effects of hurricanes on the water balance of the areas affected (*3*). They also seem unaware of the counterintuitive effects, well documented from other forms of hazard control (*4*), in which the knowledge of seeding may increase the damage toll by influencing negatively other human responses, such as evacuation, preventive measures, and so forth. And there is no mention of the low-probability outcomes, for example, the potentially negative environmental impacts of large-scale injection of silver iodide particles into the atmosphere. Such analyses are always constrained by time, effort, and imagination and must systematically exclude many considerations. And indeed many are missing from the article.

Under the second hypothesis, the use of the analysis serves as justification for decisions made on other more trans-scientific grounds. Thus if a decision is taken on the basis of considerations extraneous to the analysis (for example, the bureaucratic ambition of an organization for its own growth), will "arational" analysis be used to buttress the decision and give an unwarranted gloss of respectability? How often are even negative results ignored in such cases, with the comforting statement, "Oh, we had Stanford Research Institute carefully study the question." The precedents for this misuse are ample. The most extensive use of rational analysis to date, benefit-cost analysis in water resource development (less elegant than decision analysis, but relevant nonetheless) has served for 35 years to justify a program of water resource development that many feel has served the public less well than it could have if such analysis had been absent (*5*). In another instance of rational analysis, the results of cloud-seeding experiments in Texas, Arizona, and Florida were quickly used to justify operational cloud-seeding programs before adequate control experiments were made in dry periods.

As for the final hypothesis, one need only follow the recent reports in *Science* (News and Comment, 16 June,

p. 1216; 21 July, p. 239; 1 Sept., p. 776; 13 Oct., p. 145) concerning the massive use of environmental modification, including weather control, in Southeast Asia to consider that the experience gained in peaceful geophysical modification can be quickly turned to other purposes less helpful to mankind.

To the extent that any one of these hypotheses is valid, the social scientist committed both to rational analysis and to responsibility for his or her actions is in a dilemma. If the limits of the analysis or its possible misuse are great, would society be better off without it? I think in some cases the answer must be yes, as much in social science as in new technology. Indeed to the extent that social science becomes important (that is, people really take it seriously) social scientists must be as self-critical and responsible about their methods and their possible abuse and misuse as technologists should be about their inventions. In some cases where uncertainty is very great, it may be as irresponsible to advocate a decision-making methodology that does nothing to really reduce the uncertainty or to control its use as it is to build an SST. At the very least, until we can take into account both the limits and unintended use of decision analysis, we should be cautious in its advocacy. And in areas of great scientific unknowns, such as weather modification, where heavy pressure exists for its "arational" use and some pressure for its amoral use, extreme caution is indicated.

ROBERT W. KATES
Graduate School of Geography,
Clark University,
Worcester, Massachusetts 01610

References

1. R. W. Kates and W. R. D. Sewell, *Human Dimensions of Weather Modification* (Department of Geography Research Paper No. 105, University of Chicago, Chicago, 1966), pp. 347–362.
2. P. R. Julian, R. W. Kates, W. R. D. Sewell, *Water Resour. Res.* **5**, 215 (1969).
3. L. M. Hartman, D. Hilland, M. Giddings, *ibid.*, p. 555.
4. Taskforce on Federal Flood Control Policy, *A Unified National Program for Managing Flood Losses*, House Document 465, 89th Congr., 2nd sess. (Government Printing Office, Washington, D.C., 1966).
5. G. F. White, in *Water Research*, A. V. Kneese and S. C. Smith, Eds. (Johns Hopkins Press, Baltimore, Md., 1966), pp. 251–273.

Power is correct in suggesting that rainfall and steering effects are important issues in hurricane seeding. Another important factor is storm tide, which can be affected significantly by coastal geography. These effects might be of critical importance in the tactical decision to seed a particular hurricane. As the full report referenced in our article shows, present knowledge concerning these factors is consistent with our strategic recommendation to permit, as an emergency measure, the seeding of some hurricanes threatening a coastal area.

It is possible to conduct a decision analysis to determine the value of research on hurricane steering. However, our discussions with meteorologists have indicated that while the ability to steer hurricanes would be valuable, this ability is unlikely to result from a research program. Consequently, it is not clear that the decision analysis of steering research would demonstrate that the research has a high value.

On the question of loss of life, we found that, given the effective hurricane warnings provided by the U.S. Weather Service, the expected number of lives lost in a present-day hurricane is relatively small. If these lives are valued for decision-making purposes in a range from \$100,000 to \$300,000 each, they constitute an expected loss of only about one-tenth the expected property damage for the hurricane. Furthermore, since storms that damage less property also tend to kill fewer people, the case for removing the prohibition against seeding is only strengthened by including human loss.

We direct our commentary on Kates's letter to the three hypotheses he suggests for the nature of decision analysis.

Hypothesis H_1 is that decision analysis systematically excludes significant aspects of the problem because they are uncertain or improbable. Anyone familiar with decision analysis knows that its procedures involve not excluding, but discovering and emphasizing, significant aspects of the problem. In fact, decision analysis is uniquely concerned with assessing probabilities and their implications. Kates presents no evidence that our recommendations would be changed by additional analysis of any of the factors he mentions.

Hypothesis H_2 is that decision analysis might be misused. We agree that anything from hammers to medicine may be misused, but we find no logical argument that they should be unused. Moreover, Kates presents no evidence that our hurricane analysis has been or will be misused.

Hypothesis H_3 is that decision analysis might be used for amoral purposes. Presuming that amoral means immoral, we can only reiterate that the fact that hammers and medicine can be instruments of crime is no argument for discontinuing their production. Kates presents no evidence that our analysis has been or will be used for immoral purposes.

But Kates's hypotheses do not form a collectively exhaustive set. We would like to include a fourth hypothesis, H_4: Decision analysis is a rational method for displaying and balancing the important uncertain, complex, and dynamic factors that surround a decision. We leave it to others to judge whether this hypothesis is supported by our work.

RONALD A. HOWARD
Department of Engineering–Economic Systems, Stanford University, Stanford, California 94301
JAMES E. MATHESON
D. WARNER NORTH
Decision Analysis Group, Stanford Research Institute, Menlo Park, California 94025

Hurricane Seeding Analysis

In the article "The decision to seed hurricanes" by Howard *et al.* (*1*) it is stated in the subtitle that "On the basis of present information, the probability of severe damage is less if a hurricane is seeded." In my opinion present knowledge does not support such a statement because the results of studies of this problem do not provide a unique answer. Consequently, no conclusions can presently be made about the economic effects resulting from seeding hurricanes.

The data available from seeding experiments (such as those from Hurricane Debbie and possibly Hurricane Ginger) are too few for a statistical analysis to yield confident conclusions. Furthermore, the results of the numerical model studies referred to by Howard *et al.* conflict with results which I reported (*2*). In fact, if the method of Howard *et al.* is applied to my results, the conclusion reached is the opposite of that reached by Howard *et al.*, as I show below.

The standard deviations adopted here are the same as those in Howard *et al.* for all three hypotheses concerning the effect of seeding (H_1, reduction of the maximum wind; H_2, no effect; H_3, increase of the maximum wind). The probability distribution for the wind speed if the hurricane is seeded, w', if H_2 is true, is the same as that of Howard *et al.* (*3*):

$$P'(w'|H_2) = P(w'|H_2) = P(w) = f_N(100\%,\ 15.6\%) \quad (1)$$

where w is the wind speed of the unseeded hurricane (*4*).

Using the results of the numerical experiments presented in (*2*) I assign the following probability distribution to w' for the case that H_3 is true

$$P'(w'|H_3) = f_N{}' \ (107\%,\ 18.6\%) \quad (2)$$

The probability distribution employed for w', if it is considered that H_1 is true, is

$$P'(w'|H_1) = f_N{}' \ (95\%,\ 18.6\%) \quad (3)$$

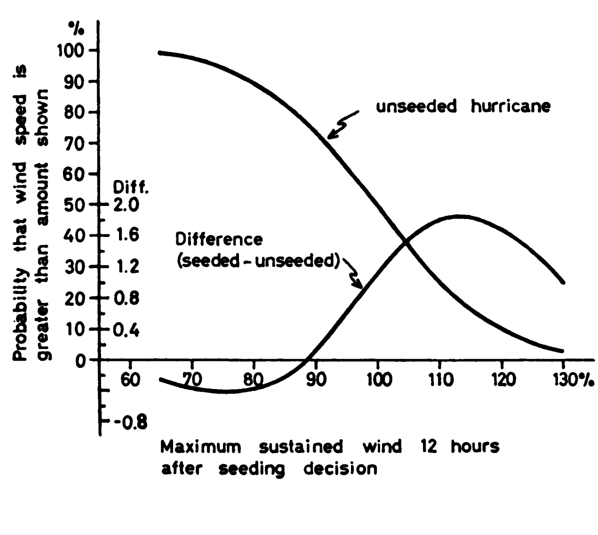

Fig. 1. Probability distribution on 12-hour wind changes for the unseeded hurricane, and the difference in probability distributions between the seeded and the unseeded hurricane.

Therefore the probability density function for the Debbie results, if hypothesis H_i is true, now becomes

$$P'(69\%, 85\%|H_i) = \begin{cases} 1.4996 & i=1 \\ 0.5716 & i=2 \\ 0.2827 & i=3 \end{cases} \quad (4)$$

Now, considering the deductions made in (*2*)—on the basis of physical reasoning and the results of numerical model experiments which definitely indicate an effect of intensification by seeding—I assign the pre-Debbie probabilities

$$\begin{aligned} P'(H_1) &= .0227 \\ P'(H_2) &= .7500 \\ P'(H_3) &= .2273 \end{aligned} \quad (5)$$

whereas in Howard *et al.* the corresponding set is

$$\begin{aligned} P(H_1) &= .15 \\ P(H_2) &= .75 \\ P(H_3) &= .10 \end{aligned}$$

Hence, the pre-Debbie odds that seeding has no effect are the same in set 5 as in Howard *et al.* However, $P'(H_3)$ is taken to be one order of magnitude larger than $P'(H_1)$ to reflect that, if seeding affects the intensity at all, an increase of the maximum wind is expected.

When sets 4 and 5 are introduced in Bayes' equation the posterior probabilities become

$$\begin{aligned} P'(H_1) &= .0647 \\ P'(H_2) &= .8131 \\ P'(H_3) &= .1222 \end{aligned} \quad (6)$$

whereas in Howard *et al.*

$$\begin{aligned} P(H_1) &= .49 \\ P(H_2) &= .49 \\ P(H_3) &= .02 \end{aligned}$$

Set 6 implies that, since the Debbie results, the odds are about 4 to 1 that seeding has no effect, and if seeding does have an effect the odds are 2 to 1 for wind intensification rather than wind reduction.

Finally, I can compute the probability distribution on wind speed [from Eqs. 1, 2, 3, and 6 above and equation 4 in (*1*)]. The difference in probability between the seeding and not-seeding alternatives is so small that it is hard to show it in a plot of the complementary cumulative distribution functions of those two alternatives. Instead, I plot this function for the not-seeding alternative and the difference (the function for seeding minus the function for not-seeding) in Fig. 1. I find that the probability for intensification (wind speed more than 100 percent of the initial wind speed) if a hurricane is seeded is .511; if the hurricane is not seeded the probability is .500 [in (*1*) these values are .36 and .50, respectively]. The probability of intensification by 10 percent or more is .278 if a hurricane is seeded and .261 if it is not seeded [.18 and .26, respectively, in (*1*)].

Furthermore, for any particular wind speed larger than 88 percent of its initial value, the probability that this speed will be exceeded is greater if the hurricane is seeded than if it is not seeded. For wind speeds less than 88 percent of their initial values the situation is reversed; however, the difference in this interval is much smaller in magnitude than it is in the former interval.

Since the analysis given above may be considered to be as soundly based as that in (*1*), it shows that the available data are too sparse to yield a statistical basis for conclusive statements. I suggest that the method of statistical analysis (possibly somewhat modified) should be used to investigate the requirements on reliability and volume of results from model studies and field experiments in order to permit confident conclusions and recommendations.

HILDING SUNDQVIST
Institute of Meteorology, University of Stockholm, Stockholm 19, Sweden

References and Notes

1. R. A. Howard, J. E. Matheson, D. W. North, *Science* **176**, 1191 (1972).
2. H. Sundqvist, *Tellus* **24** (No. 1), 6 (1972).
3. The same notation is used here as in (*1*). For convenience, some of the data given in (*1*) are repeated; the probabilities from that article are designated by P and those of the treatment given here by P'.
4. As in Howard *et al.* [reference *13* in (*1*)] a probability distribution on a quantity x is denoted $P(x)$, and a probability distribution of the normal or Gaussian family specified by its mean m and standard deviation σ is denoted $f_N(m, \sigma)$.

31 October 1972; revised 1 June 1973

In the concluding section of our article we stated: "The results of a decision analysis depend on the information available at the time it is performed. Decision analysis should not be used to arrive at a static recommendation to be verified by further research, rather it should be used as a dynamic tool for making necessary decisions at any time." We are pleased that Sundqvist finds our analysis a useful format in which to present his views regarding the results of hurricane modification. He has succinctly summarized his opinion in the form of a prior probability distribution and then used the Debbie experimental results to develop consequent probability distributions for the wind speed, both with and without seeding. His pre-Debbie probability assignment was that there was a 75 percent chance of no seeding effect, and that if there were an effect, the odds were 10 to 1 that it would be deleterious. The Debbie experiment is not sufficient to overcome this pessimistic prior probability distribution: a decision-maker who subscribed to Sundqvist's view would not wish to attempt operational hurricane seeding at this time.

Our analysis was based on the best information we could obtain from U.S. hurricane modification experts. As decision analysts we cannot comment on Sundqvist's differing opinion, except to say that our information sources were aware of his work and did not subscribe to his views. Further dialogue between Sundqvist and the community of U.S. hurricane modification experts would be appropriate to determine whether the latter see any new reason to modify their judgments.

RONALD A. HOWARD
Department of Engineering-Economic Systems, Stanford University, Stanford, California 94305
JAMES E. MATHESON
D. WARNER NORTH
Decision Analysis Group, Stanford Research Institute, Menlo Park, California 94025

24 April 1973; revised 18 June 1973

24

DECISION ANALYSIS OF THE SYNTHETIC FUELS COMMERCIALIZATION PROGRAM

Steven N. Tani
Investment Intelligence Systems Corporation

Reprinted with permission of American Federation of Information Processing Societies, Inc. The paper appears in 1978 National Computer Conference Proceedings, AFIPS Volume 47, pp. 23-29.

Decision analysis of the synthetic fuels commercialization program

by STEVEN N. TANI
SRI International

INTRODUCTION

In his State of the Union Message in January 1975, President Gerald Ford called for the accelerated development of U.S. energy technology and resources and proposed a comprehensive set of energy supply and conservation measures to reduce U.S. requirements for imported oil. As one of these measures, the President proposed a federal incentive program whose goal would be the commercial production of one million barrels per day of synthetic fuels by 1985. In such a program, the Federal Government would provide suitable financial and regulatory incentives to stimulate private sector investment in commercial-scale plants to convert coal, oil shale, and other relatively abundant domestic resources into clean liquid and gaseous fuels. It was generally believed that without such incentives, industry would be unlikely to undertake the large risks of synthetic fuel plant investments.

The benefits to be achieved by the synthetic fuels program would be the following:

1. An accelerated accumulation of experience and information on the technical, environmental, economic, and institutional aspects of commercial-scale synthetic fuel production for better-informed private sector investment decisions.
2. The development of an industry infrastructure to support subsequent expansion of the synthetic fuels industry.
3. Insurance against high world oil prices and against early depletion of domestic sources of conventional fuels.
4. Protection against the losses of an oil embargo.
5. Improvement in the U.S. international bargaining position.

These benefits, however, would be counterbalanced by the possible costs of subsidizing synthetic fuels relative to less expensive energy sources such as imported oil and other domestic resources and by the environmental and socioeconomic costs associated with rapid development of coal and oil shale reserves.

In response to the President's message, the Interagency Task Force on Synthetic Fuel Commercialization was formed to evaluate the economic and environmental costs and benefits of the program and to recommend to the President a program of appropriate size and scope. The Task Force was chaired by the Office of Management and Budget and included members from the Federal Energy Administration, the Environmental Protection Agency, the Departments of State, Commerce, and Treasury, the Council on Environmental Quality, and the National Science Foundation. We, in the Decision Analysis Group at Stanford Research Institute (now SRI International), were engaged to assist the Task Force in conducting an analysis of the program.

STRUCTURE OF THE ANALYSIS

The fundamental question addressed by this analysis was whether the U.S. should have a synthetic fuels commercialization program and, if so, how large the program should be. The Task Force defined four distinct program alternatives to be evaluated:

1. No Program—No federal funding of synthetic fuels commercialization but continuation of research and development.
2. Informational Program—A minimal program designed primarily to generate technical, environmental, and economic data on various resource-to-fuel conversion processes, with synthetic fuel production of about 350,000 barrels per day by 1985.
3. Medium Program—A program designed to generate more complete information on a wider range of processes and to meet the President's goal of 1,000,000 barrels per day by 1985.
4. Maximum Program—A program designed to achieve the greatest amount of synthetic fuel production in 1985 possible without causing major dislocations in the economy: 1,700,000 barrels per day.

The object of the analysis was to determine which of these alternatives would be of greatest net benefit to the nation as a whole, where net benefit includes both economic and non-economic impacts.

We defined four components of net national benefit for the evaluation of the program alternatives: economic impact on consumers, economic impact on producers, embargo protection, and environmental and socio-economic impacts.

To measure the economic impact on consumers, we utilized the concept of consumer surplus. Consumer surplus is the difference between the value of a good to consumers and the amount of money they must pay for it. This is shown graphically in Figure 1. The demand curve, by definition, is the most consumers would pay for each unit of the good, which is the value of that unit. If the market price is p, then q units will be purchased. For every unit except the last one, the value of the good exceeds the price paid for it. The shaded area between the price line and the demand curve represents the total excess value the consumers receive from this good; this is called the consumer surplus.

In the case of the synthetic fuels program, it was felt that a demonstration that synthetic fuels could be produced cheaply, if achieved, would have the effect of holding down the price of imported oil. The resulting increase in consumer surplus would then be credited as a positive benefit of the program.

To measure the impact on producers, we used a concept analogous to that of consumer surplus—producer surplus. This is the difference between the amount producers receive for a good and their marginal cost of producing it. Clearly, producer surplus is directly related to the idea of profitability. Figure 2 shows producer surplus graphically. The supply curve represents the marginal cost of producing each unit of the good, which is the least amount of money the producers would accept for it. The shaded area between the price line and the supply curve is equal to the total producer surplus for that good.

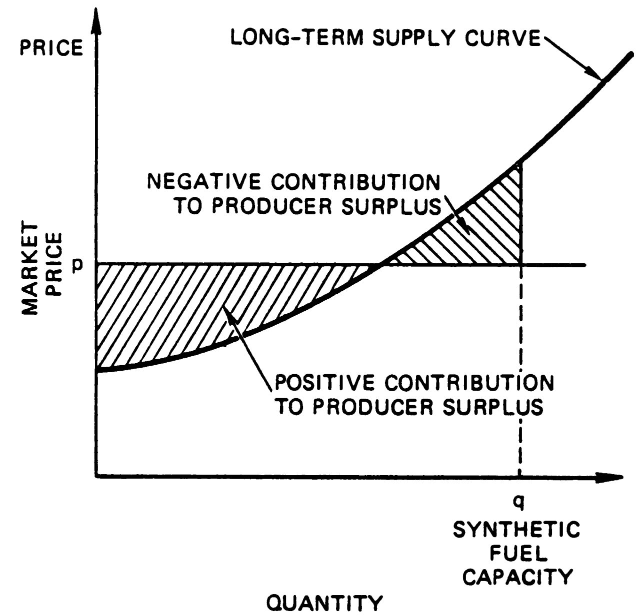

Figure 2—Producer surplus

We assumed in this analysis that synthetic fuel would be a substitute for imported oil. Therefore, if the cost of the synthetic fuels turned out to be less than the cost of imported oil, the industry would accrue positive producer surplus, which would be credited to the program as a benefit. How-

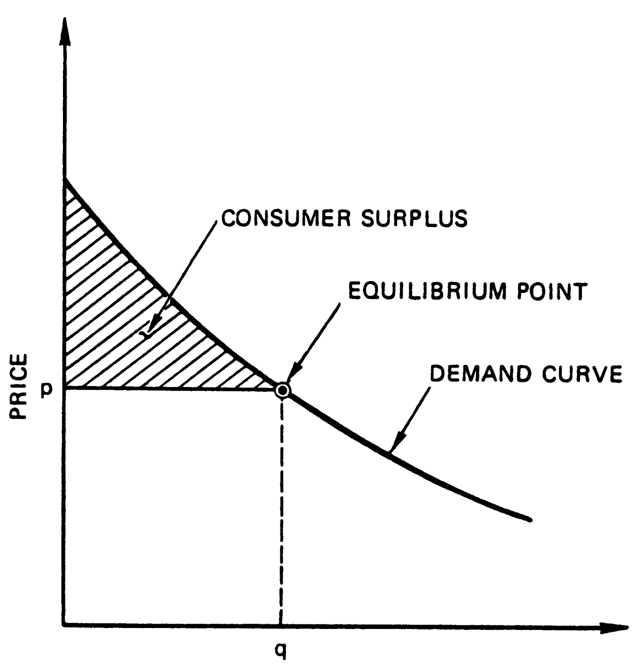

Figure 1—Consumer surplus

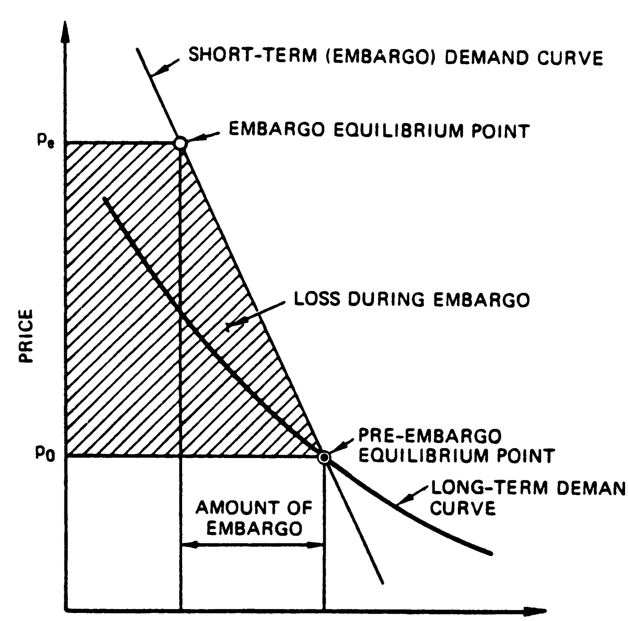

Figure 3—Embargo loss

ever, if synthetic fuels turned out to be costlier than imported oil, producer surplus would be negative and industry would require a subsidy from the government to cover its losses. The amount of this negative producer surplus would be charged as a cost of the program.

The algebraic sum of consumer and producer surplus is a measure of the total economic impact of the program on the nation assuming normal market conditions. However, it does not include the impact of the program in the event of an oil embargo. The situation during an oil embargo is illustrated in Figure 3. The pre-embargo price and quantity of oil are established on the long-term demand curve. If an embargo occurs, the quantity of oil available for consumption decreases abruptly. Because of short-term inflexibilities in consumption patterns, the marginal value (or shadow price) of oil is much higher than the long-term demand curve indicates. Here we use a linear short-term demand curve to show this effect. The economic cost of the embargo is the loss of consumer surplus during the embargo and is represented by the shaded trapezoidal area.

The synthetic fuels program, by providing a substitute for some of the imported oil, would reduce this embargo loss by increasing the amount of fuel available for consumption during the embargo. This reduction of embargo loss, weighted by the probability of occurrence of an embargo, is credited to the program as a benefit.

Finally, the synthetic fuels program would result in noneconomic costs in the form of environmental damage (e.g., increased air pollution) and socio-economic disruption (e.g., "boom towns" near mining and conversion facilities). These costs, to the extent that they are not internalized in the producers' costs (e.g., pollution control costs), are charged to the program.

The sum of these four components of program impact is the measure of net national benefit we used in the analysis to evaluate the alternatives.

Clearly, to determine the net benefit of each program alternative, we need to know something about the energy supply and demand situation in the future. There was, of course, considerable uncertainty about the future energy picture, so we used probabilistic modeling techniques to quantify and incorporate the uncertainty in the analysis.

Figure 4 shows the decision tree structure that we ultimately developed for this analysis. We treated the dynamics in a simple manner by looking at three discrete time periods. In 1975, the government would make its program decision,

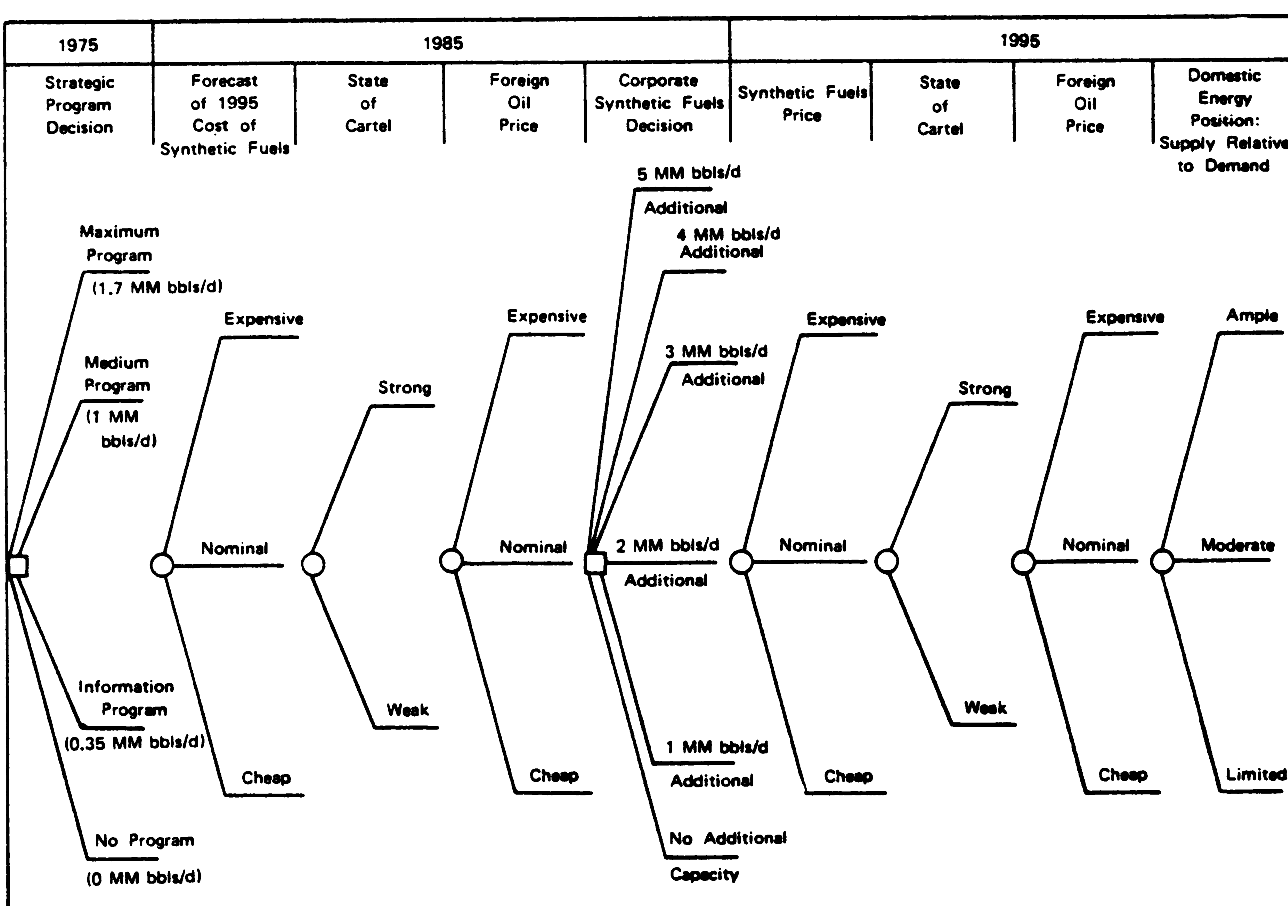

Figure 4—Decision tree

choosing one of the four alternatives. Then, in the mid-1980's, the program would result in information about the ultimate cost of synthetic fuels production. Based on this information and on the prevailing and projected price of imported oil, the industry would make its decision on further investment in synthetic fuel plants. The price of imported oil would, of course, depend on whether or not the oil producers' cartel remained effective in controlling prices. Finally, in the mid-1990's, when the new synthetic fuel plants are on-stream, the program impacts can be determined by looking at the cost of synthetic fuels, the price of imported oil, which again depends on the current state of the cartel, and the U.S. energy supply and demand balance.

The year 1985 was used to typify the decade of the 1980's and the year 1995 to typify the decade of the 1990's. Program cost and benefits were measured in constant 1975 dollars and were discounted in 1975 using a discount rate of ten percent.

The decision tree in Figure 4 shows how uncertainty was explicitly incorporated in the analysis. Uncertainty about each of the factors shown was quantified in the form of probabilities. Then, for each combination of factors, which defined a unique scenario of the future, both the probability of occurrence for that scenario and the discounted net national benefit associated with it were calculated. Finally, for each alternative, the *expected* net benefit was calculated by weighting the outcome of each scenario by its probability and summing. Note that the decision tree in Figure 4 defines 5,832 different scenarios for each of the four program alternatives.

The industry decision in 1985 of how much further investment to make in synthetic fuels plants required special treatment. While the government decision would be made on the basis of overall national benefits, the private sector decision would be made on the basis of corporate profits only. Therefore, in the analysis, the level of corporate investment that maximized expected future producers surplus was selected.

Figure 5 illustrates the techniques we used to quickly encode uncertainty in the various factors. The probability distribution shown represents uncertainty in the 1985 imported oil price given that the cartel is strong. According to this distribution, it is equally likely that the price will be above or below $15 per barrel (the median value). Also, there is a ten percent chance that the price will be below

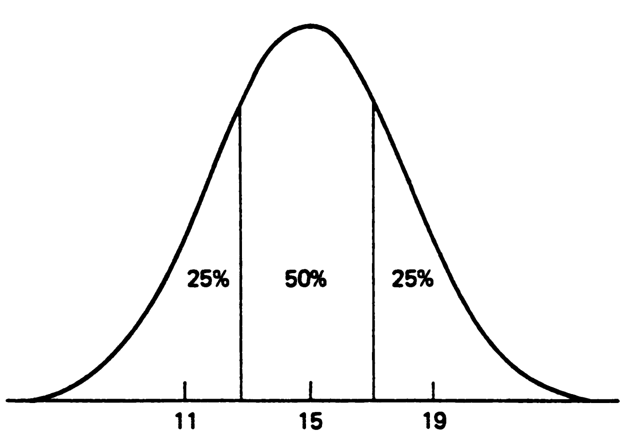

Figure 5—Simple encoding technique

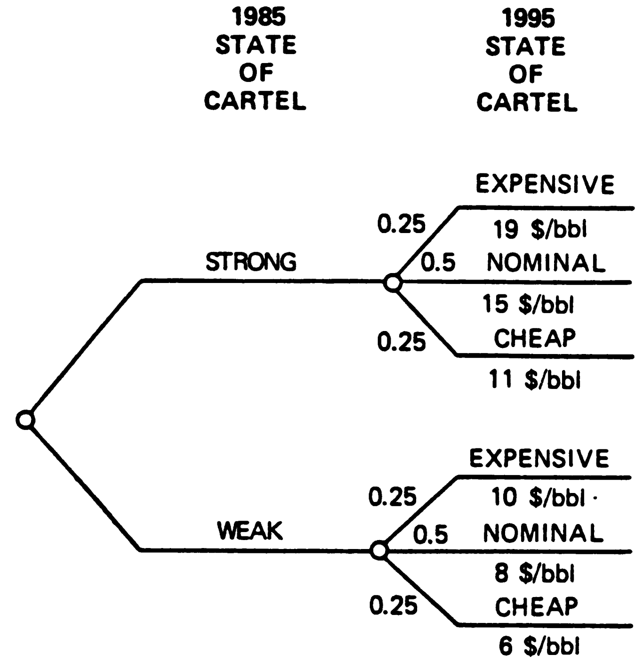

Figure 6—Probabilities of 1985 imported oil prices

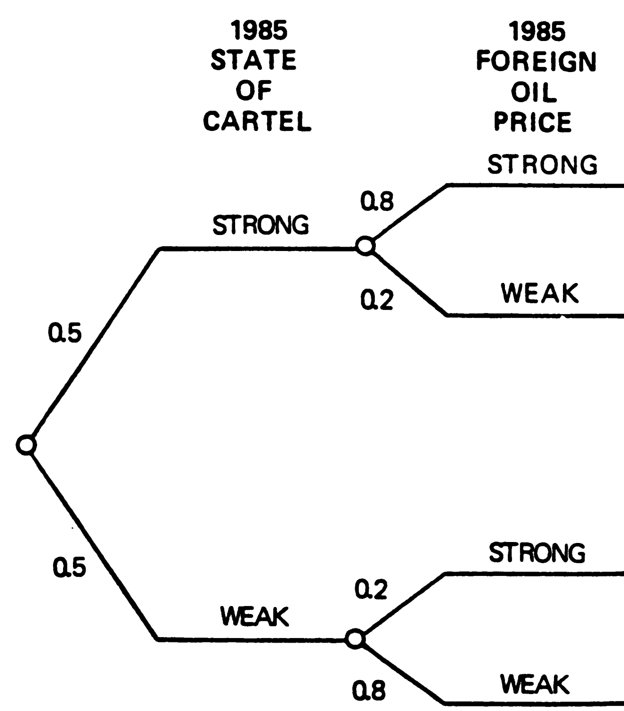

Figure 7—Probabilities of state of cartel

$11 per barrel (the 10 percent fractile) and a ten percent chance that it will be above $19 per barrel (the 90 percent fractile). We divided the distribution into three sections having areas of 25 percent, 50 percent, and 25 percent and used the median value to represent the middle section and the 10 percent and 90 percent fractiles to represent the two tails. Thus, as shown in Figure 6, we said in the analysis that there is a 25 percent chance that the 1985 imported oil price would be $19 per barrel, a 50 percent chance that it would be $15 per barrel, and 25 percent that it would be $11 per barrel, given that the cartel is strong. The imported oil price given that the cartel is weak is assessed to be much lower—$10, $8 and $6 per barrel, respectively.

Using this simple technique, we encoded the uncertainty of the Task Force members about each of the factors shown in the decision tree. Of particular interest are the assessments of the future state of the oil producers' cartel. As shown in Figure 7, the chances of the cartel remaining strong through 1985 were assessed by the Task Force to be 50-50. Given that it is strong in 1985, the probability that it would remain strong through 1995 was assessed to be 80 percent, while if it is weak in 1985, the chance that it would become strong by 1995 was assessed to be only 20 percent.

RESULTS OF THE ANALYSIS

After we had structured the problem with the decision tree, constructed a computer model to calculate the net national benefit for each of the thousands of scenarios in the tree, and encoded the uncertainties of the Task Force, we were ready to compute the analytic results.

Figure 8 summarizes these results. The total expected discounted net benefit (in billions of 1975 dollars) is shown, along with its components, for each of the three synthetic fuels program levels relative to having no program at all. These results indicated that, on balance, the synthetic fuels commercialization program was not in the best national interest and that the bigger the program, the greater the national loss. The small informational program had an expected impact on minus $1.65 billion. The larger program had expected impacts of minus $5.41 billion and minus $10.98 billion, respectively.

We can get more insight by looking at the components of total net benefit. While the synthetic fuels program is expected to have positive impacts on consumer surplus

	Expected Discounted Net Benefit (billions of 1975 dollars)				
Program Alternative	Consumer Surplus	Producer Surplus	Embargo Protection	Environmental and Socioeconomic	Total
No Program	0	0	0	0	0
Information Program (0.35 mm bbl/day)	1.07	-2.71	0.43	0.44	-1.65
Medium Program (1 mm bbl/day)	3.29	-8.74	1.18	-1.14	-5.41
Maximum Program (1.7 mm bbl/day)	4.55	-15.77	2.23	-1.99	-10.98

Figure 8—Expected program impacts

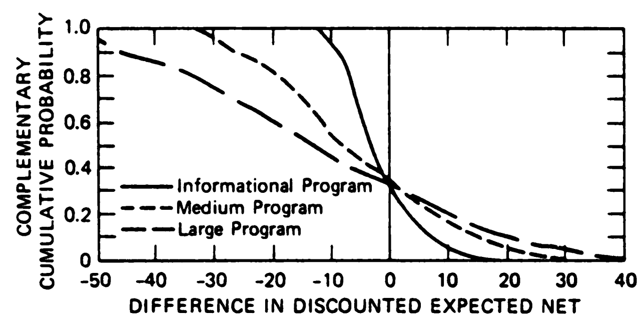

Figure 9—Uncertainty in program impacts

through the possible moderation of future imported oil prices and on embargo losses through a slight reduction in oil imports, these benefits are far outweighed by the negative impact on producer surplus. Basically, it was far more likely than not that synthetic fuels would be more expensive than imported oil and therefore need a subsidy. The negative impact of environmental and socio-economic costs is relatively minor.

The results shown in Figure 8 are the expected values of program impacts. There is, of course, considerable uncertainty about the impact of the program, as shown in Figure 9. While the expected impact of the informational program is $1.65 billion, there is a 30 percent chance that the net impact will be positive and a 10 percent chance that it will be as much as +$7 billion. On the other hand, there is a 10 percent chance that it will be as negative as −$9 billion. It is equally likely that the impact will be worse than or better than −$4 billion. The uncertainty in the impact of the larger program is even greater.

Figure 10 is a partial expansion of the decision tree that shows how two of the factors affect the results of the analysis. The −$1.65 billion expected impact of the information program consists of a 50 percent chance of −$4.86 billion if the cartel in 1985 is weak and a 50 percent chance of +$1.55 billion if it is strong. Note that a weak cartel, which leads to generally lower imported oil prices, is bad for the syn-

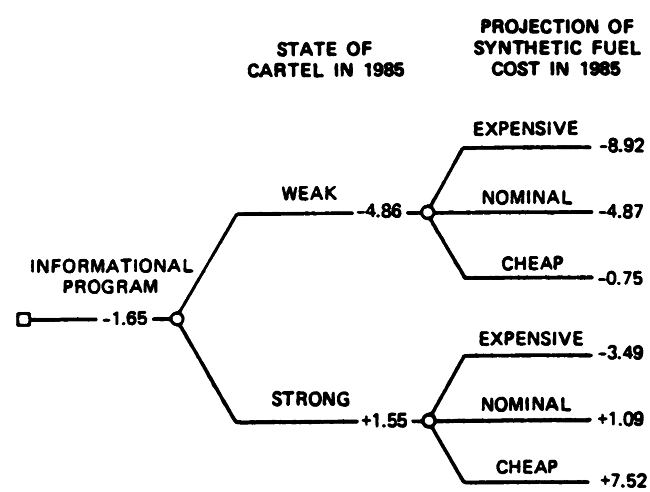

Figure 10—Partial expansion of tree

thetic fuels program while presumably very good for the nation as a whole. Conversely, a strong cartel, with higher imported oil prices, makes the program look good while being bad for the nation. This emphasizes that the synthetic fuels program is a hedging strategy—it pays off when other things are going badly. Note also that if the cartel is weak, the program looks bad even if synthetic fuels turn out to be cheap to produce. On the other hand, if synthetic fuels turn out to be expensive, the program looks bad even if the cartel is strong. That is why, on balance, the program looks bad.

The assessment of a 50-50 chance that the cartel would remain strong through 1985 turned out to be pivotal and more than a little controversial within the Task Force. To show the implications of different probabilities for this factor, we performed a sensitivity analysis, which is shown in Figure 11. This gives the expected net impact of each program level relative to no program as a function of the probability of a strong cartel in 1985. It assumes that with 80 percent probability, the cartel will remain in the same state from 1985 to 1995. Figure 11 shows that only if the probability of a strong cartel in 1985 exceeds 75 percent does the information program look better than no program and that the probability must exceed 82 percent for the medium size program to be the best alternative. An interesting result is that the maximum size program is never optimal for any value of this probability.

So far, the analytic results have been presented only in terms of expected values. It might be argued that the decision should not be made on the basis of expected values but rather on the basis of values that are adjusted for risk. To show how various levels of risk aversion would affect the

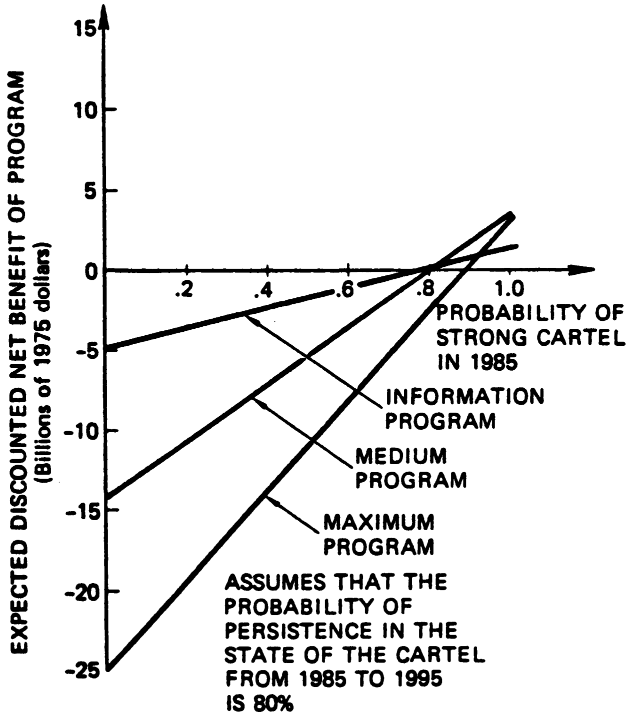

Figure 11—Sensitivity to cartel probabilities

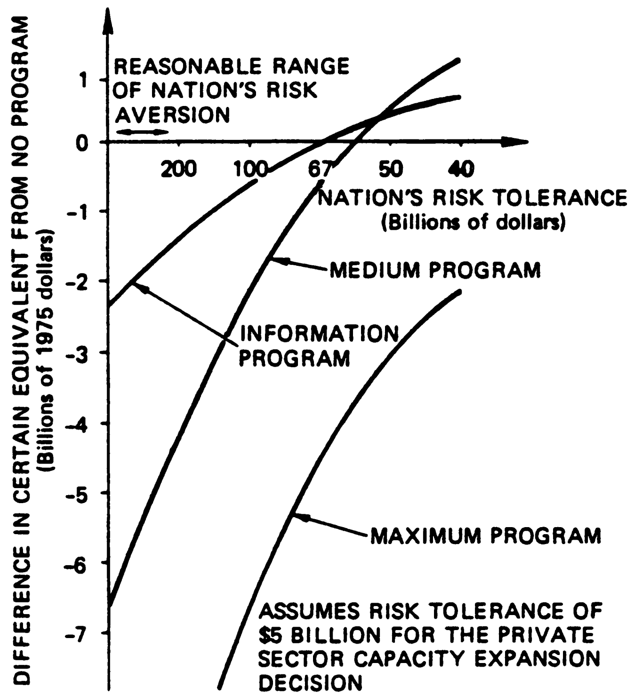

Figure 12—Sensitivity to nation's risk aversion

results, we prepared the sensitivity profile shown in Figure 12. Here, we have assumed that the nation's risk attitude is expressed by one of the family of exponential utility curves. The degree of risk aversion expressed by this curve is given by its one parameter, called the risk tolerance; the smaller the risk tolerance, the greater the degree of risk aversion. In personal terms, an individual's risk tolerance is the largest amount of money he would willingly risk in a gamble that is equally likely to halve or double that amount.

Figure 12 shows the value to the nation of each program level relative to no program as a function of the nation's risk tolerance assuming an industry risk tolerance of $5 billion for the private sector capacity expansion decision. Note first that the value of the program increases as the nation's risk aversion increases. This is characteristic of a hedging strategy, since it reduces overall uncertainty. However, the nation's risk tolerance must be less than $67 billion for the information program to be better than no program and it must be less than $56 billion for the medium size program to be the best alternative. We believe that a reasonable range for the nation's risk tolerance is from one-fourth to one-half of annual GNP, or about $300 billion to $600 billion. As Figure 12 shows, for any risk tolerance in this range, the ranking of program alternatives is the same as in the expected value case, with the best alternative being no program at all.

EPILOGUE

As far as we can determine, this is the first decision analysis to be presented in the White House. The chairman of the Task Force presented it to the President's Energy Resources Council in July 1975. Citing benefits of the program that were not quantified in the analysis, such as the international leverage gained by the U.S. in asserting positive leadership in developing alternate fuel sources, as well as the "relatively small risk and expected cost" of the small program, the Task Force recommended that the government undertake the informational program alternative with a possibility that it could switch to the medium size program pending additional information on crucial factors. The Administration's bill incorporating this recommendation ultimately failed to pass through Congress.

COMPUTER UTILIZATION

For this analysis, we wrote a FORTRAN program for use on a commercial time sharing system. The program, which we created from scratch rather than use off-the-shelf routines, required approximately 400 lines of code and cost roughly fifteen dollars to run a complete evaluation of the decision tree.

One feature of the time sharing service that we found to be especially useful was the accessibility from different locations. We did most of the model development at SRI headquarters in Menlo Park, California, but we used the program extensively while working with the Task Force in Washington D.C. Indeed, one of the most valuable aspects of our assistance to the Task Force was our ability to answer almost immediately their many questions about the effect of changes on the assumptions and assessments in the analysis.

REFERENCE

Recommendations for a Synthetic Fuels Commercialization Program. report submitted by Synfuels Interagency Task Force to The President's Energy Resources Council, November 1975. Volume I: Overview Report. Volume II: Cost/Benefit Analysis of Alternate Production Levels.

25

DECISION ANALYSIS OF SPACE PROJECTS: VOYAGER MARS

James E. Matheson
Strategic Decisions Group

William J. Roths
formerly with
Missile & Space Division
Voyager Spacecraft System Project
General Electric Company

Presented at the National Symposium, "Saturn/Apollo and Beyond," American Astronautical Society, 11 June 1967.

ABSTRACT

This paper presents the use of the concepts and techniques of decision analysis to select a space mission configuration, including both hardware and experiments, in the context of an overall space project. The techniques are applied to the selection of the Voyager Mars mission configuration of the 1970's.

The selection of each mission must be based on all of the available data from prior missions and on the policy or strategy for the selection of all future missions. First, in order to select the most economic configuration for a given mission, the most economic policy must be selected for the entire sequence of future missions that make up the space project. Then the given mission configuration is selected as the first mission of this sequence.

This sequence of mission configuration decisions and the possible outcomes is described by a decision tree in which the decision nodes represent points in the project at which the project manager will have to select a configuration. The branches following a decision node represent the configuration alternatives available at that opportunity. Chance nodes represent points in the project at which random outcomes occur. The branches following a chance node represent every major outcome that might occur at that particular chance node.

In addition to the decision tree structure for a space project, data are required on the costs of configurations, probabilities of outcomes, and values of outcomes. Thus, a configuration cost is attached to each alternative branch, and values and probabilities to each outcome branch. A method of determining values of the outcomes by means of a value tree for the space project is included in this paper. The most economic decision policy is determined by processes of expectation and maximization in the decision tree.

PREFACE

Work described in this paper was performed by General Electric and Stanford Research Institute during the Voyager Phase "A" or Phase "B" competition, which included preliminary Spacecraft design concepts only. The final Voyager Spacecraft design competition, entitled Phase "C" is scheduled to commence November 1, 1967.

The authors wish to acknowledge their many colleagues at the General Electric Missile and Space Division, Stanford Research Institute, and Stanford University, whose work has been drawn upon liberally in this paper.

I INTRODUCTION

Decision analysis, an applied extension of statistical decision theory, is a procedure for the logical and quantitative analysis of the factors that influence a decision. A decision, in this context, is an irrevocable allocation of resources as opposed to a generalized mental commitment to follow a given course of action or to pursue given objectives. A logical decision is one that selects the most economic alternative in terms of the preferences of the decision maker, the values and costs of the possible courses of action and their possible outcomes, and the probability of these outcomes as determined on the basis of the knowledge and experience available to the decision maker. A quantitative analysis is one that places the decision into an unambiguous structure in which numerical quantities encode the factors that influence the decision in order to quantitatively weigh these factors. Decision analysis incorporates the fundamentals of decision theory as a means of quantifying the decision process while maintaining the logical basis for a rational decision.

Decision making requires the study of uncertainty. Most decisions would be easy to make if there were no uncertainties in the outcomes resulting from a course of action. When the outcomes are uncertain even "simple" decisions, such as when to leave for the airport to catch a plane, become more difficult to analyze. The theory of probability provides the basis for the meaningful treatment of uncertainty. Probability is a state of mind and not of things. All prior experience must be used in measuring probabilities. If we have seen a lot of data, such as a million flips of a bent coin, the overwhelming data will be the predominate influence on the probability assignment to heads on the next flip. If we have never flipped a bent coin before, the probability measurement must depend on judgment and prior experience, possibly including physical models we might build to describe the dynamics of a flipping bent coin. If we have seen only a few flips of the bent coin, we must combine our prior experience with the limited new data. The inferential theory of probability, based on Bayes' interpretation, provides a means of logically combining the new data with the prior probability assessment.

In order to allocate resources logically, values must be placed on outcomes. Because the outcomes and hence the values are uncertain, a criteria must be established for choosing among various "value lotteries." The theory of "utility" provides a basis for encoding the risk aversion of the decision maker; that is, the desire to substitute an alternative with low expected value but low risk for an alternative with high expected value and high risk. In many cases the time preferences of the decision maker must also be incorporated.

The primary purpose of decision analysis is to increase the likelihood of good outcomes by making good decisions. A good outcome is one that we would like to occur. A good decision is one logically consistent with the information and preferences of the decision maker. Decision analysis provides a framework for making good decisions; only chance determines the ultimate outcome.

The principles upon which decision analysis is based were conceived in the days of Bernoulli, Bayes and Laplace, two to three hundred years ago. Since then these principles have been studied extensively and turned into elaborate theories. More recently operations research has applied these principles to operational problems, then management science brought them to repetitive management problems. Decision analysis is the natural next step into the one-of-a-kind major decisions. During the last ten years or so, decision analysis has been applied primarily to industrial and business decisions; decisions such as new product introduction, strategic planning of business operations and experimental program planning. Recently, applications have been made to governmental problems.

The decision analysis of space projects is a significantly new application. This paper outlines the conceptual structure by considering a simplified project for Voyager Mars. The heart of this analysis is the decision tree, a graphical method of representing the structure of a sequential decision process. Many of the important aspects of decision trees will be presented in this example.

II PROBLEM STRUCTURE

A. Nature of the Problem

Decisions must be made at many levels in conceiving or carrying out a space program. At the national level, for example, decisions are made regarding the total amount of funding to be allotted to the space program at the expense of other national goals. Within the National Aeronautics and Space Agency (NASA), this total space program funding is divided between manned and unmanned programs, and within unmanned programs the funding is still further divided into categories like planetary exploration and earth resources satellites.

Given a funding allocation for planetary exploration, decisions are then required concerning the most effective way of carrying out that exploration. Typical questions concern what launch vehicles should be employed and what type of spacecraft designs should be considered as a function of time? Still further down in the decision hierarchy: given a spacecraft concept, what design approaches for various hardware elements should be selected? For example, should power generation be based on solar or nuclear energy sources?

In all cases, the type of inputs required for making the decisions are essentially the same. The decision maker would like to know (1) what is the cost of the available alternatives, (2) what are the values of the outcomes produced by the various alternatives if they are successful, and (3) what are the relative probabilities of success of the various alternatives?

In this initial attempt to apply decision analysis to a space project, we have selected a relatively specific problem within the overall decision hierarchy previously mentioned. The situation postulated is:

> A Voyager project for the unmanned exploration of Mars has been approved with an initial launch scheduled for 1973. Preliminary studies have indicated that the Saturn V launch vehicle, which can put 40,000 to 70,000 pounds on a trajectory to Mars, is optimum for this project. It is desired to place orbiters about Mars as well as to land vehicles on Mars to collect the desired scientific data. In the preliminary studies which

showed the Saturn V to be the appropriate launch vehicle, general characteristics desired of the orbiter and landing vehicle were also determined.

While we have constrained the problem significantly by the above statement, the manager of such a Voyager project still has many decisions to make. Examples of these are:

1. Should one, two, or perhaps more Saturn V's be employed at each launch opportunity?
2. Assuming that the orbiter is to serve as a "bus" for delivering the landing vehicle to the vicinity of Mars, should each bus carry one or more landing vehicles?
3. Should the total capability of the Saturn V be used to carry a single, heavy Planetary vehicle (one orbiter with one or more landers) or should two lighter Planetary vehicles be carried to provide redundancy at the system level?
4. Given the desired characteristics of the orbiter and lander postulated previously, should the maximum desired capability be designed into the first vehicles or should a more evolutionary design approach be taken? If evolutionary, what steps in sophistication are logical?

In short, the Voyager project manager must define more precisely the mission configuration--that is, the number of launch vehicles, orbiters, and landers to be employed and the specific capabilities of these orbiters and landers. The mission configuration for the first launch must be selected, and a policy or strategy for determining what to do at subsequent launches must also be established. As we will discuss, the first choice cannot be made logically without considering the overall project objectives and the configuration sequences available to satisfy those objectives.

B. Approach to Solution

To develop the application of decision analysis to the problem posed, a two-phase program was adopted. The first or pilot phase consisted of defining a simplified version of the previously described required decision. To the maximum extent possible, however, the essential features of the problem were accurately represented and only the dimensionality was reduced. This smaller problem allowed easier development of the modeling approach, and exercising of the model provided insight into the level of detail required in structuring the inputs to the decision. The second phase consisted of incorporating all of the elements required to decide on a mission configuration for Voyager. The bulk of this report is devoted to discussion of the pilot phase, with some description of the additional factors being included in the second phase.

C. Pilot Problem Definition

For the development of the pilot model, it was postulated that the Voyager project manager had already answered many of the questions previously posed. He had decided that:

1. Only a single Saturn V will be launched at each opportunity.
2. Only one lander will be carried to Mars by each orbiter.
3. Two Planetary Vehicles will be carried on each Saturn V.
4. The design characteristics of the orbiter have been established and a single basic orbiter design will be used throughout the project.

The remaining decision, then, concerned the desired capability of the initial landing vehicle, and the desired steps in the evolution of the lander to the ultimate level of capability required. Should the project manager, for example, elect to provide the ultimate level of capability in the initial capsule in the face of uncertainties in the Martian environment and difficulties in developing complex equipment to survive the pre-launch sterilization environment? Or should he choose a much simpler capsule, which can obtain some information about the Martian environment to be factored into the design of subsequent, more complex vehicles? Which approach will yield the highest expected value from the project and what are the relative costs?

Four possible lander configurations have been postulated, which represent steps in sophistication from the simplest useful capsule to the most complex one capable of obtaining all the data ultimately desired.

The first configuration, C1, is a simple atmospheric probe, not intended to survive impact with the Martian surface. It would be separated from the orbiter as the Planetary Vehicle approaches Mars and be deflected to an impact trajectory by a small rocket. It would contain instruments that measure such parameters as density profile and composition of the Martian atmosphere during the entry phase. If successful, it would achieve an outcome that is denoted as level 1(L1).

The second configuration, C2, would be carried into Martian orbit with the orbiter and be subsequently de-orbited. It also contains the instruments for measuring atmospheric properties, and in addition has a TV camera for returning pictures of the Mars surface during the late descent phases. If totally successful, then, it would achieve L1, atmospheric measurements, plus L2, descent TV. It could, of course, achieve only L1 assuming a failure prevented returning the TV pictures. For the pilot model, it has been assumed that the C2 lander configuration could not achieve L2 without also achieving L1. For most cases this assumption is quite reasonable. There are, however, instances where it is not, and these are being handled differently in the phase 2 model.

The third configuration (C3) is the first capsule intended to survive impact and operate after landing on Mars. It also would enter from orbit, would contain the atmospheric experiments plus descent TV, and in addition would carry out relatively simple surface experiments and provide close-up TV after landing on Mars. If totally successful, then C3 can achieve L3, surface experiments plus landed TV, as well as L2 and L1. With partial success, it can achieve L2 and L1 or just L1.

The fourth configuration, C4, is the most sophisticated lander considered. It would contain all the experiments discussed above, plus the capability to carry out meaningful life-detection experiments on the surface of Mars. Total success of C4 then would lead to L4 + L3 + L2 + L1. Partial success would lead to lesser levels of achievement, of course.

The question is, again, what configuration should be selected for the first opportunity, and what sequence of configurations should be planned to follow the first choice?

D. Decision Tree

The heart of the decision model is a decision tree that represents the stucture of all possible sequences of decisions and outcomes, and contains slots into which cost, value, and probability inputs must be fed. The tree contains two types of nodes (decision nodes and chance nodes) and two types of branches (alternative branches and outcome branches), as illustrated in Figure 1. Emanating from each decision node is a set of alternative branches, each branch representing one of the configurations available for selection at that point of decision in the project. Each chance node is followed by a set of outcome branches, one branch for each outcome that may be achieved from the point in the

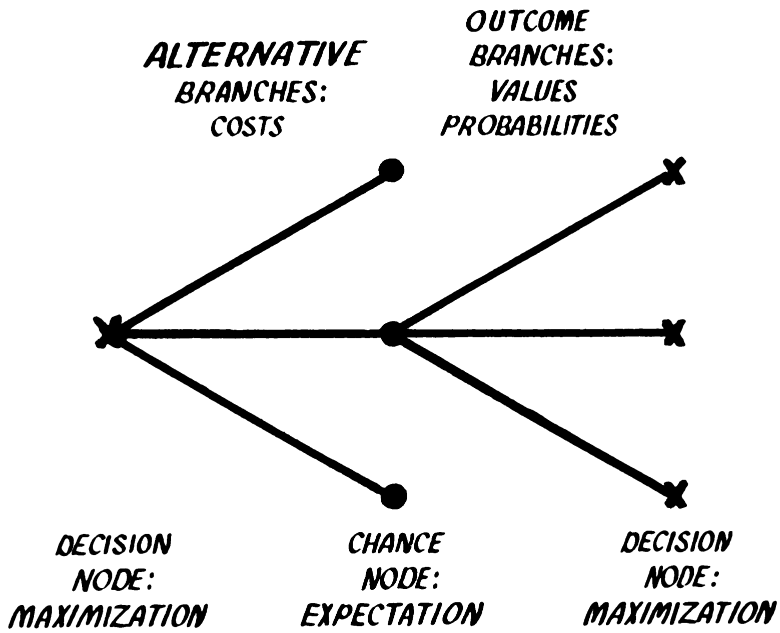

FIG. 1 TREE RELATIONS

project represented by that chance node. Probabilities of occurrence and values are assigned to each of these outcomes. Costs are assigned to each decision alternative.

Two fundamental operations, expectation and maximization, are used to determine the most economic decision from the tree. At each chance node the net expected (mean) value (NEV) is computed by summing the probabilities of each outcome, multiplied by the value of that outcome and the NEV of the node following that outcome. At each decision node the NEV of each alternative is calculated as the expected value of the following node ("successor node"), less the cost of the alternative. The optimum decision is found by maximization of these values over the set of possible alternatives, i.e., by selecting the alternative of highest NEV.

E. Order of Events

The particular sequencing of mission configuration decisions and outcomes is a significant feature of the pilot Voyager Project. As illustrated in Figure 2, the initial event of significance is the selection of the 1973 mission configuration. However, since lead time considerations require that the 1975 configuration decision be made

	1968	'69	'70	'71	'72	'73	'74	'75	'76	'77	'78	'79	...
1ST MISSION	SELECT					LAUNCH	OUTCOME						
2ND MISSION					SELECT			LAUNCH	OUTCOME				
3RD MISSION							SELECT			LAUNCH	OUTCOME		
4TH MISSION									SELECT			LAUNCH	OUTCOME
5TH MISSION											SELECT		...

FIG. 2 ORDER OF EVENTS

in 1972, we see that the second mission decision must be made prior to obtaining the first mission results. Similarly, the 1977 decision must be made before obtaining the results of the 1975 mission, although after the 1973 mission results. In general, then, a mission configuration decision must be made in ignorance of the results of the previous mission.

F. Constraining of Configuration Sequences

In defining what configurations can be selected at each decision node, some logic must be applied. It does not make sense to choose a C1 when the program has already reached outcome level 2. So this choice is not shown in the decision tree. Other restrictions can also be made that are not quite so obvious. For this pilot model, the following logic has been applied:

1. It has been assumed that the most complex lander capsule is not available in 1973 due to the development time required.

2. We have not considered sequences in which we follow a complex lander with a simpler one; that is, the complexity of the lander only increases with time.

3. Because of the order of events just discussed, where one vehicle would be in flight while another is being fabricated for the next opportunity, some logic is required to determine what will be done in the event of failure of the vehicle in flight. Failure is specifically defined as no improvement in the previous level of

success as a result of the current flight. We therefore have postulated that should failure occur, the next configuration chosen cannot be more complex than the one currently being constructed--that is, no advance is allowed in the face of failure.

4. The achievement of level 4 terminates the program. Additionally, we have postulated that two failures in succession, where again failure means no advance in level of achievement will terminate the program. These two factors constrain the overall tree size.

G. Tree Example

A completed decision tree for the simplified Voyager Project, with the additional assumption that L2 is the highest level of success, is presented in Figure 3. The model that produces the numerical probabilities, values, and costs used in the example will be discussed later. Node 1, at the left side of the tree, is the initial decision to select either a C1 or a C2 for the first launch opportunity. The box designated L0 above this node indicates that the state at this node is the current level of achievement. Suppose a C1 is selected. The cost of that C1 is $850 million, indicated by the "-850" that is written under that branch. As a result of this choice, the next node is decision node 2. The box designated L0, C1 above this node indicates that the state at this node is the current level of achievement and a C1 is being constructed for the first launch. Now either a C1 or C2 must be selected for the second launch. If a C1 is selected the cost is $575 million, and the next node is chance node 7. The two branches following this node represent the possible outcomes of the first launch. The L0' outcome, which would be failure to better L0 on the first try, occurs with probability 0.1 whereas the L1 outcome occurs with probability 0.9. The value of the L0' outcome is zero, whereas the value of the L0 outcome is 1224. Now follow the case of the L1 outcome to decision node 34. The state L1, C1 at this node, means that the highest level of success is L1 and that a C1 is being constructed for the next launch. Since L1 has already been achieved at this point in the tree, a C2 is the only configuration that may be launched in the third opportunity, at a cost of $740 million. This leads to decision node 35, where the state is L1, C2.

Node 35 in the example tree illustrates coalescence of nodes, a feature vital to maintaining a manageable tree size. Node 35 on the upper path through the tree can be reached from four other paths through the tree as indicated in the exhibit. If the coalescence did not occur, the portion of the tree following node 35 would have to be repeated four additional times. In the full pilot tree coalescence results in a reduction of the number of branches in the tree by a factor of 30.

Along the path 1-2-7-34-35, at decision node 35 a C2 must be selected for the fourth opportunity. At chance node 36 the outcome of the third launch is either an L1' (failure to better L1 with one attempt, which leads to node 38), or an L2 (which achieves a value of 1714 and successfully completes the program). These outcomes occur with probability 0.3 and 0.7, respectively. If L1' is the outcome, chance node 38 is reached where the outcome of the fourth launch is represented. The probability of L1" is 0.24, and the probability of L2 is 0.76. Note that the probability of L2 has increased over that of node 36 (0.7 to 0.76) because of the experience gained with the earlier attempt.

The reader can similarly follow and interpret many other paths through the tree. A decision policy is a complete selection of particular alternatives at all decision nodes. This limits the set of all possible paths to a smaller subset. (It is not possible, for example, to reach node 26 if a C1 is chosen at node 1.) The probabilities, values, and cost of these paths then determine the characteristics of the decision policy.

The most economic decision policy, given the input data specifications, is defined as the policy that maximizes the net expected value of the Project, i.e., expected value less expected cost. The technique illustrated here eliminates many of the nonoptimum policies from explicit consideration; it is the "roll back" technique that starts from the <u>right</u> side of the tree and progresses <u>left</u> to the beginning of the tree, making all decisions and calculations in <u>reverse</u> chronological order. Thus, when each decision is made, only policies that optimize decisions for the following decision nodes are considered.

Consider node 38 in Figure 3. At this chance node the probability of achieving L1", which is worth nothing, is 0.24, and the probability of achieving L2, which is worth 1714, is 0.76. Thus, the NEV of node 38 is: $0.24(0) + 0.76(1714) = 1303$. This number is written near node 38.

The calculations are carried out in this manner backwards through the tree. The first decision node with more than one choice is node 2. If a C1 is selected, it costs \$575 million (-575) and leads to node 7 with an NEV of 1408, which yields $-575 + 1408 = 833$. If a C2 is selected, it costs \$740 million (-740) and leads to node 12 with an NEV of 2106, which yields $-740 + 2106 = 1366$. Since 1366 is greater than 833, the most economic decision is to select a C2 at node 2, which results in NEV (2) = 1366.

Finally, the first decision is a choice between a C1 with an NEV of 516 or a C2 with an NEV of 832. Maximum NEV is achieved by the choice of a C2 resulting in NEV(1) = 832. This is the NEV of the entire project at the time the first decision is made.

The full pilot decision tree is shown in Figure 4.

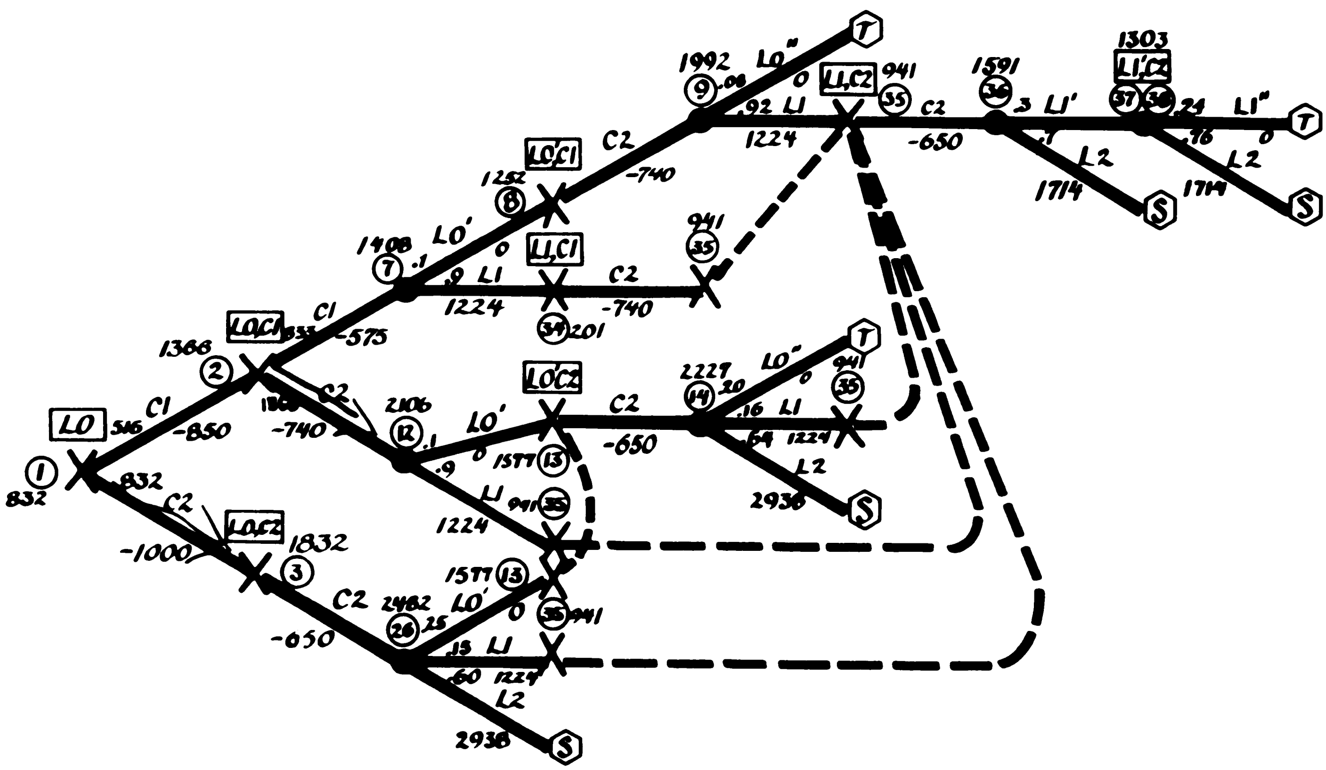

FIG. 3 EXAMPLE OF A DECISION TREE

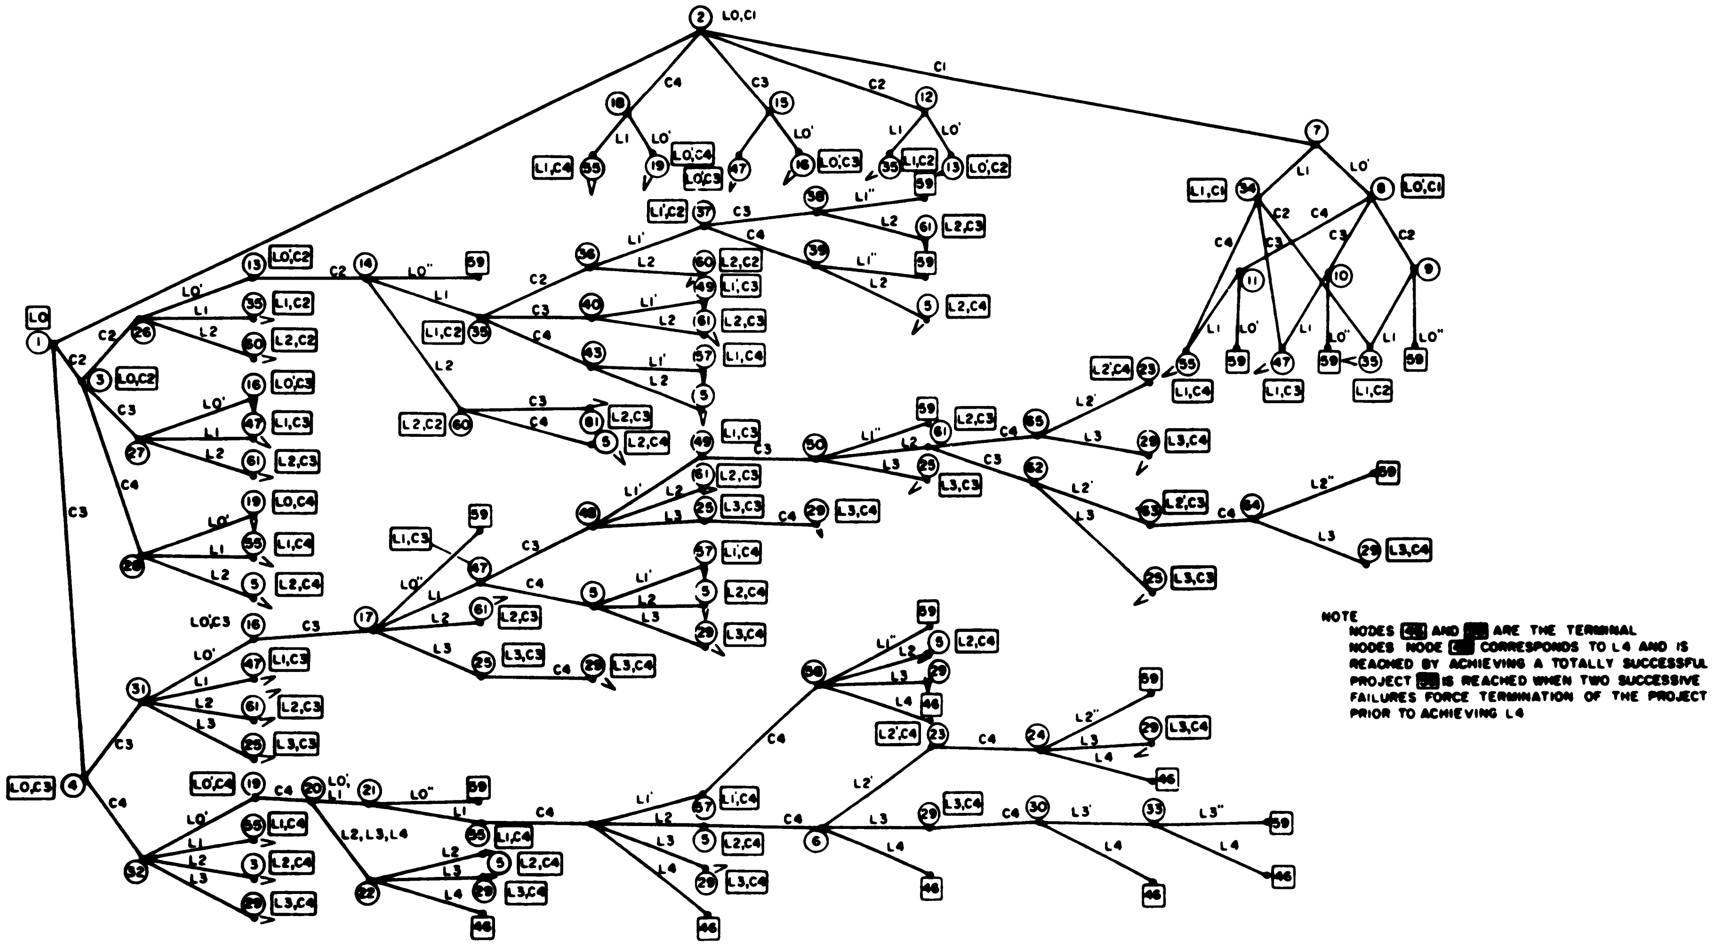

FIG. 4 DECISION TREE FOR PILOT VOYAGER PROJECT

H. Pilot Probability Model

It is, in practice, difficult to estimate the probabilities to be attached to the outcome branches leaving probabilistic nodes. Available data and engineering judgment are more easily applied to the estimation of success probabilities for more elementary operations. A probability model can then combine these inputs into the desired outcome branch probabilities.

For the purpose of the pilot analysis, probability estimates were obtained by applying engineering judgment to the various phases of the mission "top level function flow diagrams" pictured in Figure 5. In

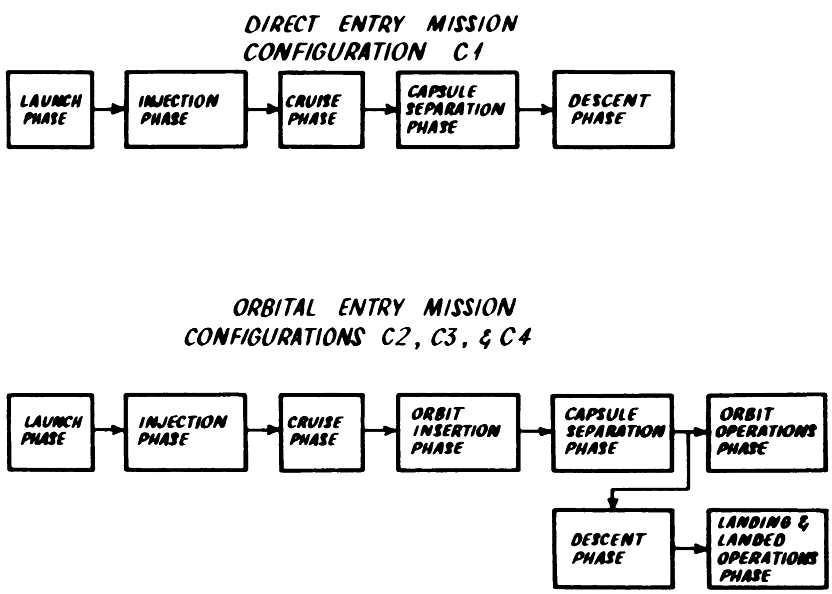

FIG. 5 VOYAGER MISSION FLOW DIAGRAMS

addition to estimating initial probabilities, it is necessary to estimate how subsequent mission probabilities will be affected by results obtained on early flights. The results of this analysis for configurations C1 and C2 are tabulated in Figure 6. Noting that a prime(') attached to a level symbol indicates failure in an attempt to better that level, we see that the probability of failing to reach L1 from L0 with a C1 is 0.10; the probability of failing to reach L1 with a C1 after failing once previously is 0.08, etc. The general philosophy upon which this model is based is that success increases the probability of achieving desirable outcomes, but so does failure (designs are improved due to information gained in the failure, etc.), although to a lesser extent. Thus, the probability of failing with a C2 starting at L0 is 0.25; the probability

CONFIGURATION C1

PREVIOUS LEVEL / NEXT LEVEL

	L_0'	L_0''	L_1
L_0	0.1	-	0.9
L_0'	-	0.08	0.92

CONFIGURATION C2

PREVIOUS LEVEL / NEXT LEVEL

	L_0'	L_0''	L_1	L_1'	L_1''	L_2
L_0	0.25	-	0.15	-	-	0.60
L_0'	-	0.20	0.16	-	-	0.64
L_1	-	-	-	0.30	-	0.70
L_1'	-	-	-	-	0.24	0.76

FIG. 6 TRANSITION PROBABILITY INPUTS

of failing with the same configuration starting at L0' is 0.20, with the remaining 0.05 of probability being spread proportionately over the other outcomes.

I. Pilot Cost Model

The costs attached to the alternative branches in the decision tree are constructed by the cost model using representative system cost data, as tabulated in Figure 7. As shown in this figure, we have assumed that second and subsequent copies of hardware systems (bus, landers) can be manufactured at half the initial cost of developing and producing the first system. Also, configuration production experience on a C1, C2, or C3 contributes a 20 percent reduction in the cost of producing a C2, C3, or C4, respectively.

J. Components of Value

The "value" of various accomplishments in the Voyager project can be divided into assigned values and divided values. The assigned values are the values of the Voyager project. The derived values are the contributions of earlier mission accomplishments to the probability of success

PREVIOUS CAPSULE CONSTRUCTED	CAPSULE COST OF C1	C2	C3	C4
NONE	150 M	300 M	600 M	1000 M
C1	75	240	600	1000
C2	–	150	480	1000
C3	–	–	300	800
C4	–	–	–	500

COST OF BUS

FIRST TIME 400 M, SUBSEQUENT TIMES 200 M

OTHER COSTS 300 M

FIG. 7 PILOT COST INPUTS

in later Voyager missions. The derived contributions are incorporated by increases in the probability of success in future missions, hence increasing the probability of ultimately achieving future extrinsic values.

The assigned value itself can be divided into two distinct kinds--direct and indirect. The direct value is the value of the knowledge produced by the outcomes, such as visual records of Mars and characterization of Martian biology. The direct value is achieved independently of the means of gaining the knowledge. For example, we would obtain the direct value even if the knowledge were given to us by our worst enemy. The indirect value is the value of obtaining and possessing the knowledge (rather than the knowledge itself), such as "technological spin-off," national prestige, satisfaction of our "Columbus urge" to explore Mars, and the competitive pleasure of being first in space. Both direct and indirect values provide the total incentive for the Voyager Project; to make project decisions rationally, both values must be included in the evaluation of project outcomes if both affect the level of the decision being considered.

K. The Value Model

To derive a value function, we construct a value tree by considering first the major components of value, both direct and indirect, and then the subcategories of each type identified in more and more detail until no further distinction is necessary. Then each tip of the tree (constructed as above) is subdivided into four categories, each corresponding to the contribution of one of the four levels of achievement within the value subcategory represented by that tip.

The value tree that serves as the value function in the pilot analysis is pictured in Figure 8. The figure 1.0 attached to the node at the extreme left represents the total value of all the objectives of the pilot Voyager Project (thus, the value of achieving L1, L2, L3, and L4). The four branches emanating from this node represent the four major categories of value recognized by the pilot model. The figure 0.62 attached to the upper branch represents the fraction of total value assigned to science. Two branches emanate from the science node, and we see that 60 percent of the science value falls into the category of biological science. The 0.37 attached to the biological science node represents the fraction of total value attached to biological science, and is obtained by taking 60 percent of 0.62 (the fraction of total value

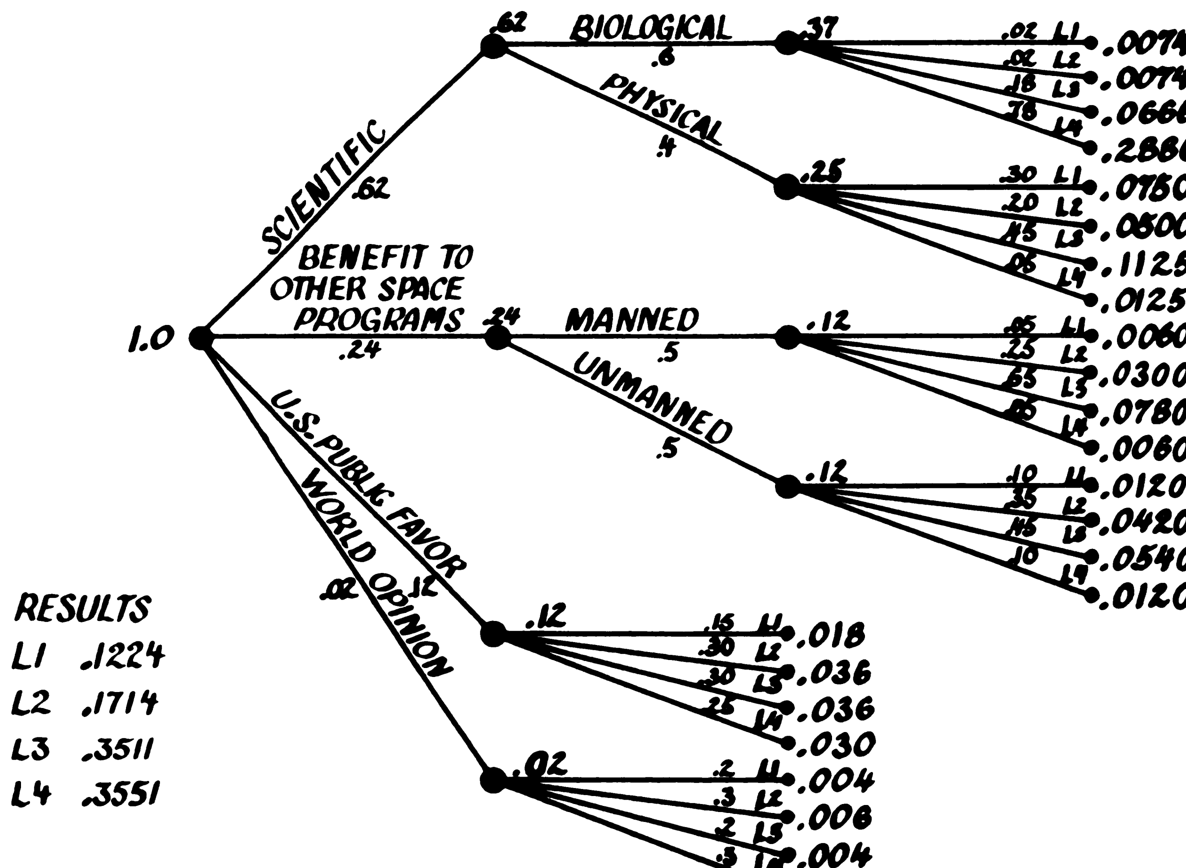

FIG.8 THE VALUE TREE

attached to all science). Finally, the bottom branch following the biological science node indicates that 78 percent of the biological science value is achieved by jumping from L3 to L4.

The final step in value modeling is to obtain the fraction of total value to be attached to achieving each of the four levels. If we add all the contributions to achieving L1 (contributions to world opinion, U.S. public favor, physical science, etc.) the result is the fraction of value that should be attached to achieving L1. The same process is followed for reaching L2 from L1, L3 from L2, and L4 from L3. The results of such a calculation are in the lower-left corner of Figure 8.

L. Rationale for Selecting the Optimal Policies

A "policy" in this context is a setting of each decision node in the decision tree, i.e., it is a complete strategy for conducting all missions in the Voyager Project. The selection of a policy limits the number of paths that might be traversed through the decision tree. However, there are still many alternative paths which may be traversed, which are determined by individual mission successes and failures as the project unfolds. These vagaries have been encoded in the probability model. Thus, for each policy we can derive a probability distribution over all the paths that could possibly be taken. Using the cost and value models, we can also assign a cost and value to each of these paths. Thus, each policy is represented by a lottery on cost and value, as illustrated in Figure 9.

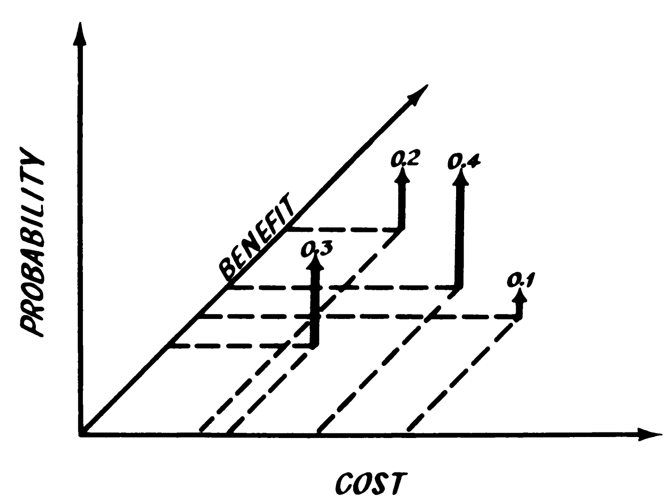

FIG. 9 REPRESENTATION OF A POLICY

To select an optimum policy, we must decide between many different cost-value lotteries. Figure 10 illustrates three different policy lotteries in the cost-value plane. Inspecting this figure, we see that policy 3 tends to have costs similar to policy 2 but lower values than policy 2. Policy 3 also tends to have higher costs than policy 1 but similar values. Thus, policy 3 does not look very desirable. However, policy 2 tends to have both higher values and higher costs than policy 1, and we must therefore determine whether the chances of higher values are worth the chances of higher costs.

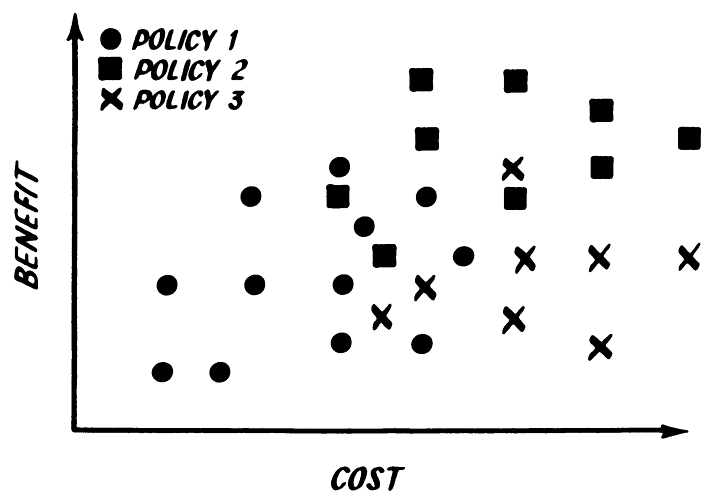

FIG. 10 COMPARISON OF POLICIES

In general, the risk aversion preferences of the decision maker must be encoded in order to make this decision. To gain insight into this example, we have assumed that the decision maker wishes to base his decisions on the expected value and cost of each policy. Thus, we replace each policy lottery by a single point at the expected cost and the expected value of that policy. These points are exhibited for nine hypothetical policies in Figure 11.

The policies in Figure 11 may be separated into three classes: totally dominated policies, marginally dominated policies, and dominant policies. A policy is totally dominated if there is at least one other policy that has both a lower expected cost and a higher expected value. Policies (6), (7), (8), and (9) are totally dominated policies. Totally dominated policies can be dropped from further consideration. This simplifies decision making because the bulk of the possible policies are of this type.

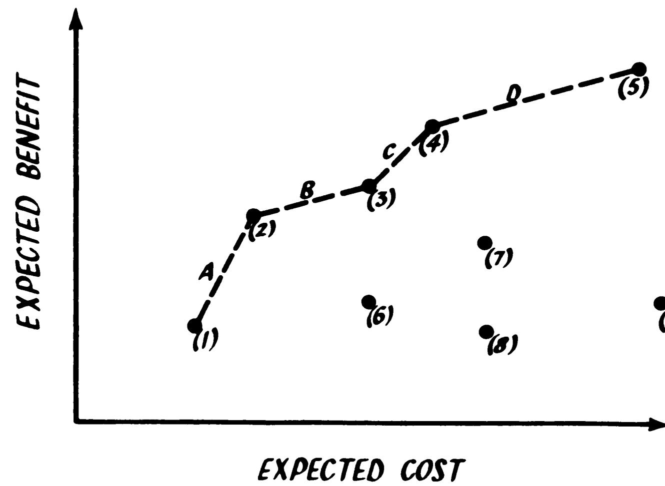

FIG. 11 EXPECTED VALUES OF ALTERNATIVE POLICIES

The remaining policies, connected by a dashed line in Figure 11, are called envelope policies and are of two types: marginally dominated policies and dominant policies. Policy (3) is an example of a marginally dominated policy. The slopes of the lines A, B, and C in Figure 11 show the marginal return from increasing the funding level from (1) to (2), (2) to (3), and from (3) to (4). The marginal return B is less than the marginal return C. This means that increasing the funding from (2) to (3) brings less return per unit cost than the increase in funding from (3) to (4). Since each program is competing for funds with other programs, it would be unusual to find that directing funds from other programs is worthwhile up to point (3), but is not worthwhile beyond (3). For this reason marginally dominated policies such as (3) are eliminated from contention.

The remaining policies, (1), (2), (4), and (5) are called dominant policies. The selection has been reduced to this set, which is shown in Figure 12.

Analysis of the decision tree produces this set of dominant policies and their associated expected costs and benefits. Presumably one policy from this dominant set is best. But which one? In order to make this policy selection, the decision maker must compare these policies with alternate projects competing for the same funds.

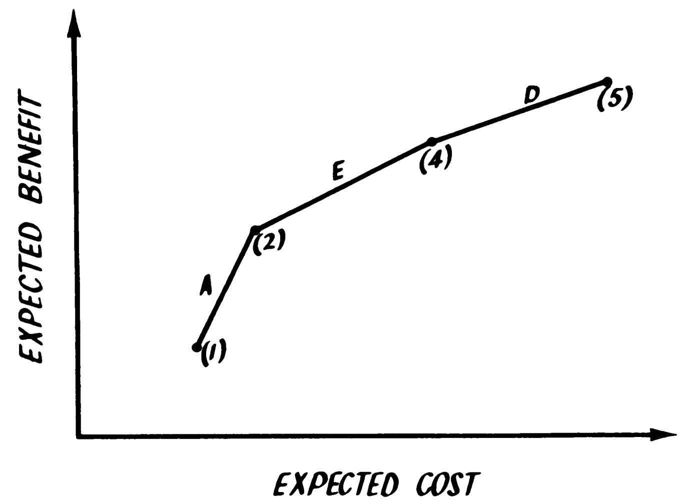

FIG. 12 DOMINANT POLICIES

M. Determination of Dominant Policies and Profit Interpretation

To determine the dominant policies, we must compare a great number of alternative policies. For example, in a 60-node tree there were approximately 2,000 policies. However, the "roll-back" process of the decision tree presented earlier provides a method of finding the dominant policies without evaluating each policy explicitly. Recall that in the decision tree we subtracted expected costs from expected values. To do this, we converted the value points to dollars by multiplying each point by a conversion factor λ. Since the entire value tree has been assigned one point, λ is a dollar value assignment to the Voyager Project.

Figure 13 shows the construction of expected profit from an expected cost-expected value point and an assignment of λ. We can picture this result as being obtained by shining a light beam across the figure from the direction of the upper right hand corner, so that the rays of the beam have slope λ^{-1}. The expected profit for the policy is determined by the shadow of the light beam on the expected cost axis, with increasing profit to the left.

If we shine the light beam on all of the policy points simultaneously, as in Figure 14, the shadows of the policies on the expected cost axis give the expected profit for each policy. The policy with the leftmost shadow is the policy of maximum expected profit, and thus it is the policy that results from evaluation of the decision tree for the given λ. If we shine the light more vertically from the top, lower cost policies will have the leftmost expected profit shadow; as we shine the

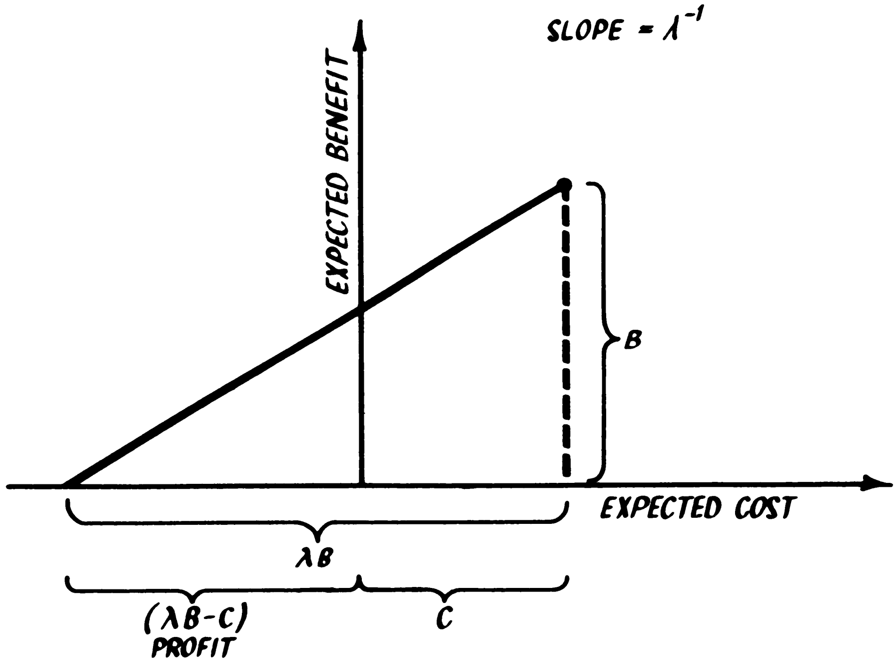

FIG. 13 CONSTRUCTION OF PROFIT

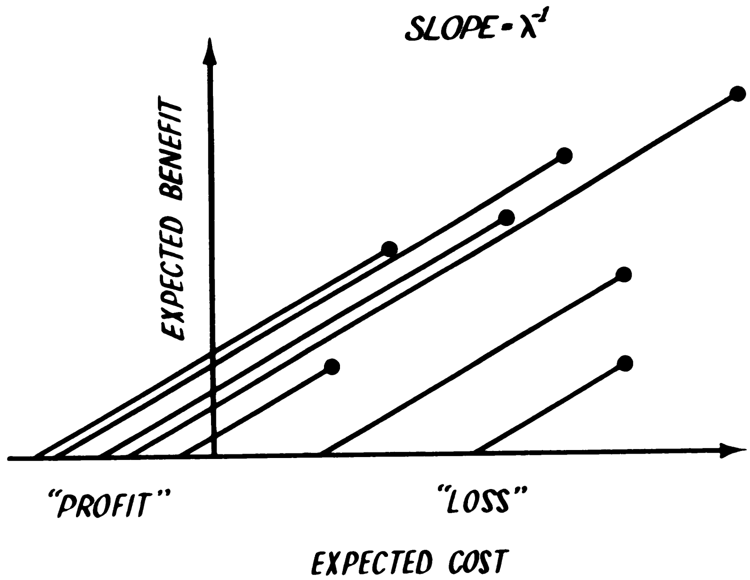

FIG. 14 SHADOWS OF POLICY POINTS

light more horizontally from the right, higher cost policies will have the leftmost expected profit shadow. Consideration of the geometry of this problem shows that sweeping the light from vertical to horizontal, that is, sweeping λ from zero to infinity, will produce exactly the dominant set of policies that were illustrated in Figure 12. Thus, successive evaluations of the decision tree for different λ assignments will sweep out the dominant policies. Algorithms have been developed that sweep out the entire set with high efficiency.

III RESULTS

Figure 15a lists several policies, as well as interpretations of the policies, and shows their expected costs and expected value fractions. Policy 1 (P1) is interpreted as--fly an initial C3, followed by C4's until the program terminates, either by achievement of L4 or by two failures in succession. Policy P2 is an initial C1 followed by C4's, and P3 is two C1's followed by C4's. These are the three optimum policies, and the interpretations show that as higher dollar value is attached to the project the policies become more conservative.

Policy P6, the stepwise policy, is to start with a C1, and advance the configuration sophistication by one step whenever possible. This policy has the lowest expected fraction of total value. Policy P10 is the most complicated policy. It begins with two C1's. If the first C1 achieves L1, then the project continues with C4's. If the first C1 achieves L0' (it fails), then C3's are flown until L2 is attained, and then C4's follow. Interpretation of other policies may be read off Figure 15a.

All possible policies for the pilot Voyager decision tree lie within the dashed lines of Figure 15b. Policies P1, P2, and P3 are the dominant policies. Policy P1 is the minimum expected cost policy, whereas P3 is the maximum expected value fraction policy. Policy P2 is a trade-off policy that has expected values and costs between those of P1 and P2. As one would expect, P2 is optimum for intermediate project dollar value, whereas P3 and P1 are optimum for high and low extremes, respectively.

The remainder of the policies shown are all totally dominated policies. Of course, there are many more policies that have not been examined. All of these policies, however, are either totally dominated or marginally dominated. The methods described in this paper allow determination of the dominant policies (P1, P2, P3, in this case) without explicitly considering most of these dominated policies.

Figure 16 illustrates optimal policies for two different value-tree assignments, the balanced value assignment previously described and an assignment of all value to the life experiment (L4). The dominant

POLICY	INTERPRETATION	EXPECTED COST	EXPECTED FRACTION OF TOTAL VALUE BALANCED	LIFE
1	C3,C4's	4124	0.802	0.669
2	C1,C4's	4443	0.815	0.720
3	C1,C1,C4's	5031	0.827	0.730
4	C2,C3,C4;s	4488	0.769	
5	C2,C3's x C4 @ L1	4816	0.759	
6	STEPWISE	5885	0.729	
7	C3's TILL L1,C4's	4555	0.796	
8	C1 THEN AS IN 7	4609	0.805	
9	C3's TILL L2,C4's	4630	0.795	
10	C1,C1,[illegible] TILL L2,C4's	5092	0.825	
11	C1,C1,C3,C4's	5167	0.816	
12	C1,C1,C3 TILL L2,C4's	5720	0.813	
13	C2,C3,C4's	6475	0.735	
14				
15				

FIG. 15a COMPARISON OF SEVERAL POLICIES

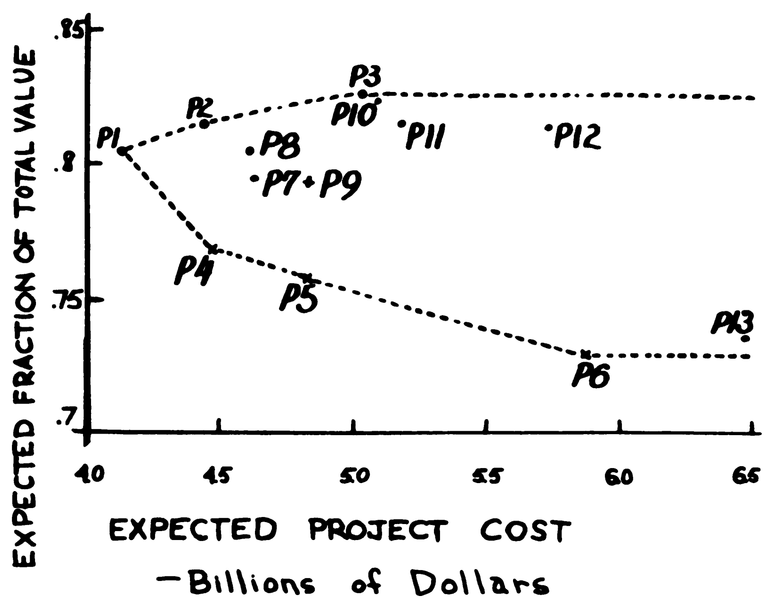

FIG. 15b COMPARISON OF POLICIES FOR BALANCED VALUE ASSIGNMENTS

policies were found by varying the dollar value assignment to the Voyager Project (λ) as described in the previous section. For either value assignment, the dominant set of policies consists of policies denoted P1, P2, and P3.

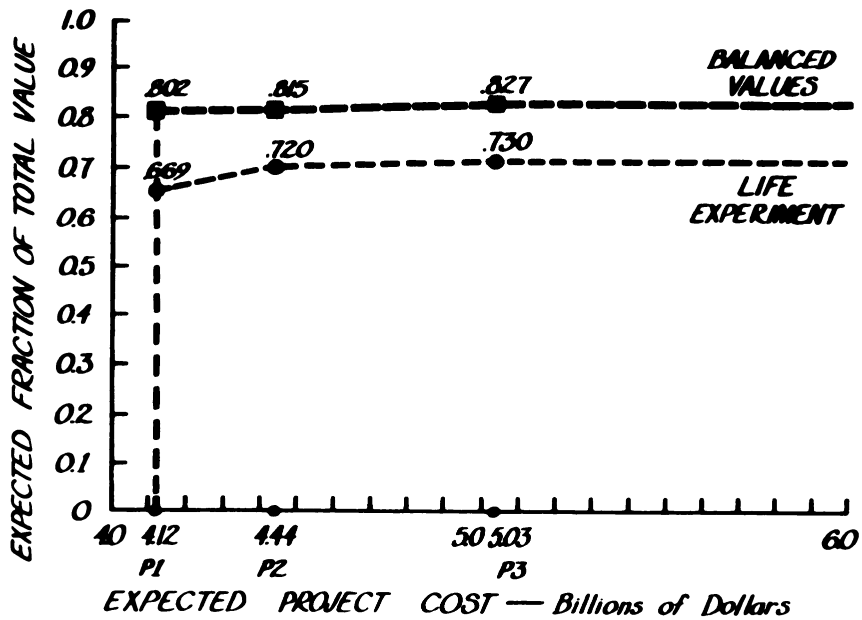

FIG. 16 EXPECTED COSTS AND VALUES FOR DOMINANT POLICIES

IV CONCLUSIONS AND FURTHER WORK

Work with the pilot model has provided encouragement that this approach can be a truly useful one. It has been found that exercising the model with varying inputs provides good insight into the relative importance of the many factors involved and their interrelationships. The people using this model have become sufficiently proficient so as to be able to predict in advance the effect of changing value, cost, or probability inputs with reasonable accuracy. The fact that their ability to do this has improved greatly indicates the improvement in their understanding of the interrelationship of the many factors involved. Providing the decision maker with better understanding of the relative importance of the many parameters influencing his decision will certainly be valuable.

The model can be exercised in many ways to provide this better understanding. Most importantly, sensitivity studies can be carried out to determine how the optimum decision policy is affected with changes in the input parameters. There will always be uncertainty in many of the parameters assigned. If the optimum decision policy is the same over the range of uncertainty, the decision maker need not be too concerned with determining the exact numerical values. On the other hand, if the optimum decision changes over the range of uncertainty, the project manager would like to determine how much it is worth to reduce the uncertainty. For example, if the values of two projects that are optimum, depending upon the exact value of a given parameter are nearly equal, the decision maker can select the more conservative approach with little sacrifice. Alternatively, it could be that a much more valuable project could be carried out if he were sure a given probability number was nearer to 0.9 than to 0.7. In this case, he is concerned with determining the cost of a development program that is required to reduce his uncertainty, and is assessing whether this cost is warranted in terms of the potential increase in the project value.

An important additional benefit of this analysis is that it provides a language for communicating the structure of the space project and the data factors relevant to the project decisions. It provides a valuable mechanism for discourse and interchange of information, as well as a means of delegating the responsibility of determining these factors.

Based on the promising results of working with the pilot model, a more complete model has been developed which encompasses nearly all of the factors involved in selecting the actual Voyager mission configuration. It includes nearly all of the realistic configurations available within the context of using a Saturn V launch vehicle. It provides a more precise structure for the assignment of initial values, probabilities and costs, and for updating probabilities and costs based on results achieved. Figure 17 shows a summary comparison of the complexity of the pilot model with the more complete model. In addition to the

PILOT	PARAMETER	FULL SCALE
4	MISSION CONFIGURATIONS	14
13	MISSION OUTCOMES	56
5	PROJECT OUTCOMES	56
5	CAPSULE OUTCOMES	14
NONE	ORBITER OUTCOMES	4
OPEN	LAST POSSIBLE FLIGHT	1981
60	DECISION TREE NODES	≈ 3000
≈ 1000	DECISION POLICIES	$\rightarrow \infty$

FIG. 17 DECISION TREE COMPARISON

features shown in Figure 17, some of the ground rules postulated for the pilot model that affect the available choices at each decision node were found to be undesirable and these have been modified in the more complete model.

This advanced model has been programmed and is being exercised currently. Too few runs have been made at this time to allow presentation of meaningful results.

If the model is as useful as expected, it can be a valuable tool throughout the life of Voyager Project. As the project progresses, the knowledge of costs, probabilities and values will improve as a result of development programs and flights. The effect of this improved knowledge can be factored into the decision process each time a configuration must be selected for the next opportunity.

ABOUT SDG

Strategic Decisions Group (SDG) specializes in helping capital-intensive, risk-intensive, and research-intensive companies analyze their most critical decisions, develop strategies, and create business opportunities. Formed in 1981 and located in Menlo Park, SDG comprises a staff of consultants who have been pioneers in strategy development and decision analysis for more than fifteen years. The SDG combination of broad business experience, specific industry knowledge, and technical expertise in decision analysis has proven outstandingly effective to clients in developing strategies, creating business innovations, allocating resources, managing risk, and selecting business portfolios. SDG consultants work in joint teams with clients to transfer knowledge and to build the strong client commitment necessary for successful implementation. By blending this participative consulting style with an ongoing professional development program, SDG enriches its clients' own corporate capabilities. *Readings on the Principles and Applications of Decision Analysis* is a product of SDG's continued dedication to advancing the principles of good decision making.

ABOUT THE EDITORS

Dr. Ronald A. Howard, Professor of Engineering-Economic Systems in the School of Engineering of Stanford University since 1965, directs teaching and research in the Decision Analysis Program, and is the Director of the Decision and Ethics Center. Dr. Howard defined the profession of decision analysis in 1964 and has since supervised several doctoral theses in decision analysis every year. His experience includes dozens of decision analysis projects in virtually all fields of application, from investment planning to research strategy, and from hurricane seeding to nuclear waste isolation. He has been a consultant to several companies and is presently a founder and director of Strategic Decisions Group. He has published three books, written dozens of technical papers, and provided editorial service to seven technical journals. He has lectured on decision analysis at universities in several foreign countries, including the Soviet Union and the People's Republic of China. His current research interests are life-and-death decision making and the creation of a coercion-free society.

Dr. James E. Matheson, a founder and director of Strategic Decisions Group and a leading figure in developing professional decision analysis, has supervised hundreds of decision analysis applications in such areas as corporate strategy, capital investment, research and development, environmental safety, contract bidding, space exploration, and public investment. He created the SRI International Decision Analysis Group and directed it for fourteen years. He is responsible for innovations in methodology that have made decision analysis a powerful tool throughout many fields and industries. In 1967, Dr. Matheson was appointed to the consulting faculty of Stanford University where he is currently Consulting Professor of Engineering-Economic Systems. He continues to help businesses develop their own decision analysis capabilities through SDG's extensive executive and professional development program. His current research interests are managing risk in business portfolios and integrating research and development decisions into business strategy.

STRATEGIC DECISIONS GROUP
3000 Sand Hill Road • Menlo Park, California 94025