


DECISION ANALYSIS FOR THE PROFESSIONAL

PETER MCNAMEE

JOHN CELONA



FOURTH EDITION

Decision Analysis for the Professional

Peter McNamee
John Celona

FOURTH EDITION

SmartOrg, Inc.

www.smartorg.com

This work is released under the terms of the
Creative Commons Attribution-Noncommercial-No Derivative works 3.0 License.



Decision Analysis for the Professional

Peter McNamee

John Celona

This work in all its entirety is copyrighted and protected by the Creative Commons Attribution-Non-Commercial-No Derivative Works 3.0 license (<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>). You are free to distribute this work **ONLY** under the terms of this license. Distribution in classrooms under the terms of this license is also to be considered as non-commercial usage (no commercial profit from the distribution is permitted). If you wish to distribute this book commercially, or create derivative works, you will need to obtain express written permission from SmartOrg. Please consider buying this book from Amazon or SmartOrg to encourage the release of such works to the public under similar licenses.

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

ISBN 0-9710569-0-0

Copyright © 2001-2008 by SmartOrg, Inc.

Electronic Edition, Oct 2008



Publisher: SmartOrg, Inc.

Editors: Mimi Campbell, Bill Roehl, Mary Story

Typography and Layout: Don Creswell and T/D Associates

Cover Design: Rogodino & Associates

Excel is a registered trademark of Microsoft Corporation. Supertree and Sensitivity are registered trademarks of SmartOrg, Inc.

Dedication

We would like to dedicate this book to a number of very special people—Tereza, Patrick, Christina, and Andrew McNamee, and Karen Schwartz, who endured computer widowhood and orphanhood during the long weeks, weekends, and months while this book was first being written. Without their constant support and indulgence, this project would have never come to fruition.

Preface

In the thirteen years since the first edition of this book, the authors have seen a dramatic evolution in the practice of decision analysis. The number of companies using decision analysis as an approach to problem solving has grown rapidly. Our experience during this period has shown that practical as well as analytical skills are needed for successful implementation of a decision analysis program.

As a problem-solving approach, decision analysis involves far more than the use of decision trees as a calculational tool. It is a process of framing a problem correctly, of dealing effectively with uncertainty, of involving all the relevant people, and, above all, of communicating clearly. Accordingly, in addition to the analytical techniques used in decision analysis, this book presents material that the authors hope will assist the reader in integrating these techniques into a practical and effective problem-solving process.

The dialog decision process (DDP) and the language of decision quality have emerged as a powerful tool in the application of decision analysis in a world of delegated decision making and cross-functional teams. The team process combines with the analytical clarity of decision analysis to produce decisions which can be accepted and implemented by the organization.

This edition splits the material into four major sections. The first section addresses the tools of decision making and decision analysis. The second section then shows how these tools can be applied in the complex corporate environment. The third section is new and presents the process and language that has been developed for dealing with teams and delegated decision making. The fourth section deals with more advanced topics which are of interest to the more advanced practitioner.

The book has been rewritten so that it is independent of software. In several examples and problems, we use hand calculations to teach readers what the computer programs do. In principle, this text could be read (and many of the problems done) without using computer software.

However, as everyone knows, it is impossible to do much in decision analysis without the aid of supporting software. Descriptions of how to use some software packages such as Supertree are available from the authors.

We hope this book will lead the student to develop an appreciation of the power, practicality, and satisfying completeness of decision analysis. More and more, decision analysis and the dialog decision process are becoming

PREFACE

accepted as the best way to address decision problems. Being a decision facilitator is an exciting and satisfying occupation. This text is designed to emphasize this. Furthermore, since texts tend to remain on students' shelves, we hope this book will be of assistance long after the course is done.

The text is intended for a short course in decision analysis in business schools. It could also be part of an analytical methods course. The general philosophy of the book, however, is more consonant with extending the course by having the student apply decision analysis to more complex cases, perhaps based on real data or problems supplied by local businesses.

Decision analysis is both young enough that its founders are alive and active in the field and old enough that the literature on the field has grown large. We have chosen not to write a book bristling with footnotes. Rather, we have chosen to list in the bibliography several books in different areas for readers interested in those topics. We ask our colleagues not to take offense if their names or works are not explicitly referred to in this book. We gratefully acknowledge their contribution of accumulated wisdom and knowledge, which has made decision analysis a useful and powerful management tool.

We especially thank all the people who contributed their useful comments and constructive suggestions, including Charles Bonini, Max Henrion, and Myron Tribus. Dr. Bonini kindly allowed us to use the IJK Products and Hony Pharmaceuticals problems that appear in Chapters 4 and 7, respectively. We are also indebted to Dr. Udi Meirav for assistance in the discussion of Options and Real Options. We also thank Yong Tao, who assisted in constructing problems for the book. We are particularly indebted to Ronald A. Howard and James E. Matheson for contributions and insights, which appear throughout this book.

Peter McNamee
pmcnamee@smartorg.com
Menlo Park, California

John Celona
jcelona@sbcglobal.net
San Carlos, California

Contents

Chapter 1 – Introduction 1

Origins of Decision Analysis 1

Decision Making 2

 A Philosophy 2

 A Decision Framework 2

 A Decision-Making Process 3

 A Methodology 3

Dealing with Complex Problems 3

 Discovering the Real Problem 4

 Keeping the Analysis Manageable 4

 Finding the Insights 5

Dealing with Complex Organizations 5

 The Team Approach 5

 Structured Dialog 6

 Decision Quality 6

Advanced Topics 6

 Dealing with Uncertainty 7

 Dealing with Complex Informational Relationships 7

 Obtaining Reliable Information 7

Focus of This Book 7

Problems and Discussion Topics 8

Part I: Decision Making

Chapter 2 – Uncertainty and Probability 13

Probability: A Language of Uncertainty 13

Why Bother with Probabilities? 17

What Are Probabilities? 18

Using Intuition Effectively 19

CONTENTS

The “Divide and Conquer” Approach	19
Passing the Clairvoyance Test	19
Assigning the Numbers	20
Using Trees	20
The Data in the Nodes	21
Drawing the Probability Tree	23
A Value Function	24
Analyzing the Tree	26
How Much Detail Is Enough?	29
Certain Equivalent and Expected Value	30
Encoding Probabilities	32
Obtaining the Data	33
Discretizing the Data	34
Using Range Data	35
Summary	36
Problems and Discussion Topics	37

Chapter 3 – Decisions Under Uncertainty 41

What Is a Good Decision?	41
Recasting the Problem As a Decision Problem	43
Decision and Value Nodes	43
Rules for Constructing Influence Diagrams	45
Constructing the Decision Tree	45
Decision or Uncertainty?	48
Building the Tree	49
Decision Criterion	50
The Value of Nonmonetary, Intangible Goods	52
The Value of Future Money	53
The Trade-off Between Certainty and Uncertainty	54
Analyzing the Tree	55
The Value of Perfect Information	58
The Value of Perfect Control	62
Summary	62
Problems and Discussion Topics	63

Chapter 4 – Probabilistic Dependence 75

Dependence and Independence	75
Dependent Probabilities	76
Obtaining the Data	76
Using the Data	76
Dependent Outcomes	79

Obtaining the Data	79
Using the Data	80
Nature's Tree	80
Indicators and States of Nature	81
An Example from Medicine	82
Bayes' Rule	85
A Prototype As an Indicator	86
The Value of Imperfect Information	87
The Value of Imperfect Control	92
Common Interpretive Problems with Reordering the Tree	92
Decision-Dependent Probabilities	92
Decision-Dependent Outcomes	93
Summary	94
Problems and Discussion Topics	94

Chapter 5 – Attitudes Toward Risk Taking 101

The Inadequacy of Expected Values	101
Toward a Consistent Risk Attitude	102
Order Rule	103
Equivalence Rule	104
Substitution Rule	104
Choice Rule	104
Probability Rule	105
What Is a Utility Function?	106
Using a Utility Function	108
Value of Information with an Encoded Utility Function	109
An Exponential Utility Function	109
Value of Corporate Risk Tolerance?	116
An Approximation to the Certain Equivalent	117
Encoding a Utility Function	119
Deriving the Existence of a Utility Function	120
Risk-Free Discount Rates	124
Summary	127
Problems and Discussion Topics	127

Part II: Dealing with Complex Problems

Chapter 6 – The Complexity of Real-World Problems 141

A Cyclical Approach	142
Basis Development	

CONTENTS

Starting the Process	144
Using Decision Hierarchy	
Using Strategy Tables	145
Using Influence Diagrams	147
Deterministic Structuring: Modeling the Problem	150
Modeling Profit	151
Sunk Costs	151
Shutting Down the Business	152
Inflation	153
Terminal Value	154
Deterministic Structuring: Sensitivity Analysis	154
Probabilistic Evaluation: Building and Pruning the Tree	161
Basis Appraisal: Obtaining Information from the Tree	166
Conditional Distributions	166
Value of Perfect Information	170
Value of Perfect Control	171
Sensitivity Analysis to Probabilities, Risk Attitude, and Value Trade-offs	172
Policy Matrix	174
Time to Prepare and Present	179
Summary	179
Problems and Discussion Topics	180

Chapter 7 – Typical Corporate Applications of Decision Analysis 189

New Product Introduction	189
Litigation Decision Analysis	195
Bidding Strategies	198
Investment and Investment Rollover Decisions	200
Options	
Real Options	
R&D Decisions	
R&D Portfolio	
Corporate Strategies and Business Portfolios	202
Summary	205
Problems and Discussion Topics	205

Part III: Corporate Decision Making

Chapter 8 – A Decision Making Process 223

Decision Projects	224
Should There be a Project?	224

- Choosing the Decision Process 225
 - Structured Decision Process 226
- Dialogue Decision Process 227
 - Framing Dialog 229
 - Alternatives Dialog 234
 - Analysis Dialog 236
 - Decision Dialog 241
- DDP and the Decision Analysis Cycle 243
- Project Staffing and Timing 243
- Presenting Decision Analysis Results 244
- Decision Process Capability Building 245
- Summary 246
- Problems and Discussion Topics 247

Chapter 9 – Decision Quality

- Quality in Decision Making 253
- People Quality and Content Quality 254
- Elements Used in Measuring Decision Quality 255
 - Appropriate Frame 256
 - Creative, Doable Alternatives 257
 - Meaningful, Reliable Information 258
 - Clear Values and Tradeoffs 259
 - Logically Correct Reasoning 260
 - Commitment to Action 261
- Decision Quality and the Smart Organization 261
- Summary 263
- Problems and Discussion Topics 264

Part IV: Advanced Topics

Chapter 10 – Probability Theory 267

- Theory Overview 267
- Definition of Events 267
- Distinctions 268
- Algebra of Events 271
- Mutually Exclusive and Collectively Exhaustive Events 271
 - Mutually Exclusive 272
 - Collectively Exhaustive 272
- Joint Events 273
- Tree Representation of Events 274

CONTENTS

Probability and States of Information	275
Probability Theory	276
Joint, Marginal, and Conditional Probabilities	277
Bayes' Rule	280
Probabilistic Independence	281
Multiply or Add Probabilities?	281
Events, Variables, and Values	283
Representations of Probabilities for Discrete Values	284
Tabular Form	284
Tree Form	284
Mass Density Plot	284
Cumulative Probability Graph	284
Histogram	286
Mean, Median, Mode, Variance, and Standard Deviation	286
Moments and Cumulants	288
Representations of Probabilities for Continuous Variables	290
Problems and Discussion Topics	292

Chapter 11 – Influence Diagram Theory 303

Theory Overview	303
Elements of Influence Diagrams	304
Uncertainty	304
Decision	305
Influence	306
Determined Uncertainty	310
Value and Determined Value	310
Rules for Constructing Influence Diagrams	311
Procedures for Manipulating Influence Diagrams	312
Turning an Influence Diagram into a Decision Tree	313
Summary	316
Problems and Discussion Topics	316

Chapter 12 – Encoding a Probability Distribution 321

Level of Detail in Encoding Probability Distributions	321
Problems in Encoding Probability Distributions	322
Motivational Biases	322
Cognitive Biases	323

Availability Bias 324
Representativeness Bias 324
Adjustment and Anchoring Bias 325
Implicit Conditioning Bias 325
Probability Encoding Process 326
 Stage 1: Motivating 326
 Stage 2: Structuring 327
 Stage 3: Conditioning 328
 Stage 4: Encoding 329
 Stage 5: Verification 331
Experiences and Insights from Practice 331
Summary 332
Problems and Discussion Topics 332

Bibliography 335

Index 337

Authors 341

1

Introduction

Origins of Decision Analysis

Decision-making is one of the hard things in life. True decision-making occurs *not* when you already know exactly what to do, but when you do not know what to do. When you have to balance conflicting values, sort through complex situations, and deal with real uncertainty, you have reached the point of true decision-making. And to make things more difficult, the most important decisions in corporate or personal life are often those that put you in situations where you least know what to do.

Decision science evolved to cope with this problem of what to do. While its roots go back to the time of Bernoulli in the early 1700s, it remained an almost purely academic subject until recently, apparently because there was no satisfactory way to deal with the complexity of real life. However, just after World War II, the fields of systems analysis and operations research began to develop. With the help of computers, it became possible to analyze problems of great complexity.

Out of these two disciplines grew decision analysis: the application of decision science to real-world problems through the use of systems analysis and operations research. Decision analysis is a normative discipline, which means it describes how people should logically make decisions. Specifically, it corresponds to how (most) people make decisions in simple situations and shows how this behavior should logically be extended to more complex situations.

This book is divided into four parts, as described in the following pages. The first part develops the tools of decision making. The second part describes how these tools can be applied to complex problems. The third part presents a process for using decision analysis in today's corporate setting. Some specialized topics are dealt with in the fourth part.

Decision Making

Imagine a decision-maker struggling with a difficult decision problem. The decision analysis approach provides a normative approach that can support the decision-maker.

Decision analysis functions at four different levels—as a philosophy, as a decision framework, as a decision-making process, and as a decision-making methodology—and each level focuses on different aspects of the problem of making decisions.

Part I of this book lays the foundation for all four levels.

A Philosophy

As a philosophy, decision analysis describes a rational, consistent way to make decisions. As such, it provides decision-makers with two basic, invaluable insights.

The first insight is that uncertainty is a consequence of our incomplete knowledge of the world. In some cases, uncertainty can be partially or completely resolved before decisions are made and resources committed. However, in many important cases, complete information is simply not available or is too expensive (in time, money, or other resources) to obtain.

Although this insight may appear obvious, we are all familiar with instances in the business world and in personal life in which people seem to deny the existence of uncertainty—except perhaps as something to be eliminated before action is taken. For example, decision-makers demand certainty in proposals brought before them. Twenty-year projections are used to justify investments without any consideration of uncertainty. Time and effort are spent to resolve uncertainties irrelevant to the decision at hand. And this list could, of course, be greatly extended.

The second basic insight is that there is a distinction between good decisions and good outcomes. Good outcomes are what we desire, whereas good decisions are what we can do to maximize the likelihood of having good outcomes. Given the unavoidable uncertainty in the world, a good decision must sometimes have a bad outcome. It is no more logical to punish the maker of a good decision for a bad outcome than it is to reward the maker of a bad decision for a good outcome. (Many types of routine decisions have little uncertainty about outcomes; thus, in these cases, it is not unreasonable to associate bad outcomes with bad decisions.)

This insight, too, may seem obvious. Yet how often have we seen corporate “witch hunts” for someone to blame or punish for unfortunate corporate outcomes?

A Decision Framework

As a framework for decision-making, decision analysis provides concepts and language to help the decision-maker. By using decision analysis, the decision-maker is aware of the adequacy or inadequacy of the decision basis: the set

of knowledge (including uncertainty), alternatives, and values brought to the decision. There is also a clear distinction between decision factors (factors completely under the decision-maker's control) and chance factors (uncertain factors completely outside the decision-maker's control). Moreover, the decision-maker is aware of the biases that exist in even the most qualitative treatments of uncertainty. He or she knows these biases exist because people are not well trained in dealing with uncertainty and because they are generally overconfident in describing how well they know things.

A Decision-Making Process

As a decision-making process, decision analysis provides a step-by-step procedure that has proved practical in tackling even the most complex problems in an efficient and orderly way. The decision analysis cycle provides an iterative approach that keeps the focus on the decision and that enables the decision facilitator* to efficiently compare the decision alternatives. Modeling, both deterministic and probabilistic, reduces the problem to manageably sized pieces and allows intuition to function most effectively. Knowledge of the biases in probability estimation enables the decision-maker or facilitator to take corrective action.

A Methodology

As a methodology, decision analysis provides a number of specific tools that are sometimes indispensable in analyzing a decision problem. These tools include procedures for eliciting and constructing influence diagrams, probability trees, and decision trees; procedures for encoding probability functions and utility curves; and a methodology for evaluating these trees and obtaining information useful to further refine the analysis.

It is a common mistake to confuse decision analysis with constructing and manipulating influence diagrams and decision trees. The real contribution and challenge of decision analysis occur at the much more general level of defining the problem and identifying the true decision points and alternatives. Many decision analyses never reach the point of requiring a decision tree. Nonetheless, obtaining a full understanding of the philosophy and framework of decision analysis requires some familiarity with the methodology, including trees.

Dealing with Complex Problems _____

Imagine a decision-maker struggling with a decision problem that involves a complex set of interactions. A decision may affect several products in several different markets. There may be many different alternatives which should be under consideration. Information may be difficult to obtain and may be of uncertain validity.

*The term *decision facilitator* is used throughout this book rather than the more traditional *decision analyst* to emphasize the many roles the individual must play in bringing clarity to a decision.

There is a temptation in problems of this type to go to either of two extremes in using decision analysis. Either the analysis is done at a superficial and often simplistic level, resulting in inadequate insight for the decision-maker and perhaps in incorrect conclusions. Or the analysis attempts to include all possibly relevant detail, resulting in the ultimate abandonment of decision analysis because it is perceived as lengthy and expensive.

Part II of this book presents a method to steer between this version of Scylla and Charybdis. The method helps the facilitator discover the real problem, keep the analysis manageable, and find the insights.

Discovering the Real Problem

Finding the real problem is often the most crucial task facing the decision-maker and the decision facilitator. Problems worth extended analysis often come to the surface because many people see only parts of a problem or opportunity. There is confusion as to how things interact, what the possibilities are, what the threats are, what is important and what is irrelevant to the decision. Information is fragmentary, alternatives have not been thought out, consequences have not been fully identified.

The decision basis provides a structure that cuts through much of the confusion and helps identify the real problem. The decision basis is composed of the answers to three questions: What are the possible alternatives? What information do I have to describe these alternatives? What value (decision criterion) do I want to use to choose between the alternatives? When the decision basis is developed, the underlying problem is usually well identified.

The decision analysis cycle then refines the decision basis through a series of approximations. Start with a simple analysis and use the tools of sensitivity analysis to discover what is important and what is irrelevant. With one or two iterations, the problem is almost always clearly identified.

Keeping the Analysis Manageable

People involved in the decision-making process will usually keep the facilitator from falling into the trap of making the analysis too simple. But what will keep the facilitator from making the analysis too complex? The decision analysis methodology provides guidance.

The decision analysis cycle not only guides the direction in which the analysis grows, but also contains the rules for judging when the analysis should stop and the decision made.

In each iteration of the cycle, various forms of sensitivity analysis determine what information is important and why one alternative is better than another. This guides the next iteration of the analysis, and helps the facilitator avoid the addition of irrelevant detail and complexity.

But when does the process stop? When is the level of detail sufficient? The process should stop when the cost of further refinement (sometimes

money, more often time) is greater than the benefit the refinement would provide to the decision-maker. Value of information and control provide the key concepts used to identify this cost/benefit trade-off.

Finding the Insights

The purpose of the analysis is not to obtain a set of numbers describing decision alternatives. It is to provide the decision-maker the insight needed to choose between alternatives. These insights typically have three elements: What is important to making the decision? Why is it important? How important is it?

The various forms of sensitivity analysis and probabilistic analysis readily identify which factors are important in making a choice and which are not. The decision analysis cycle iteratively focuses the analysis on the important factors and develops an understanding of why these factors matter to the decision and how much they contribute to the difference of value among the alternatives.

Dealing with Complex Organizations _____

Imagine a set of decision-makers trying to identify a set of alternatives, choose between them, and create the conditions required for successful implementation in a multi-organizational environment. Many decisions in the modern corporate world are cross-organizational. They involve decision-makers within the organizations, information shared between organizations, and implementation tasks in each organization.

Part III of this book outlines a decision process that gets the analysis done well, and also manages the organizational dimension in such a way that a decision choice will be implemented effectively. Essential elements of this process are a team approach, structured dialogs, and the concept of decision quality.

The Team Approach

The most effective means of dealing with cross-organizational problems and opportunities appears to be the cross-organizational team. The team normally has someone from each organization to present the information and concerns of that organization. Team members are ordinarily detached from their ordinary duties (either part- or full-time) for a fixed length of time to achieve some well defined goal.

The Dialog Decision Process (DDP) has been developed to combine the decision analysis approach with the team approach. This is important from the decision analysis point of view because the information and alternatives need to come from the organizations, and teams are an effective means of accomplishing this. The team structure is important from the cross-organizational point of view because team interaction and understanding will contribute to the successful implementation of the decision.

The DDP sets up two teams. A decision team is composed of the effective decision-makers. A project team is composed of the people who will supply the information and perform the analysis. A schedule is set up to determine the points of interaction between the teams.

Structured Dialog

The DDP is based on a structured dialog between the decision team and the project team. At several points during the project, the two teams will meet for a specific purpose. At these meetings, the project team members present the results they have developed up to that point and request input from the decision team members.

Several important goals are achieved in the ensuing dialog. First, the developing understanding of the problem can lead to redefinition of the project or redirection of the efforts of the project team. Second, the experience and knowledge of decision team members can contribute to the analysis in a timely fashion. Third, the decision team members will be exposed to the concerns of the other organizations, and this shared understanding will be important later during implementation,

Decision Quality

A single decision-maker can decide when the time has come to stop the analysis and make the decision; decision analysis can provide some guidance, but it is really up to the decision-maker to decide when the decision is “good”—logically consistent with the decision-maker's decision basis (alternatives, information, values).

However, in the cross-organizational, multi-decision-maker environment it is not so easy to determine when a decision is “good.” A more detailed language is needed to facilitate the discussion and indicate when the team is decision-ready.

The language of Decision Quality has been developed to fill this need. It describes both the quality of the analysis and the quality of commitment to action. Decision Quality is measured by a number of quantitative estimates which, although subjective, are less ambiguous than purely verbal descriptions. And Decision Quality can be monitored periodically during the course of a DDP and corrective actions can be taken if required.

Advanced Topics

Clarity of thinking and common sense are the most important skills required of a decision facilitator. As will be seen in the first three parts of this book, many applications of decision analysis do not require complex mathematics or very specialized interviewing skills.

Of course, there are some decision analysis problems that soar into the realms of mathematical complexity and form the substance of Ph. D. theses. And the ability to deal sensitively with people and facilitate group meetings is essential for anyone involved in the decision process.

Part IV of this book is intended to take the facilitator a short way along these paths in three areas: dealing with uncertainty, dealing with complex informational relationships, and obtaining reliable information.

Dealing with Uncertainty

Clarity of discussion through the language of probability is one of the hallmarks of decision analysis. We must confront the reality of uncertainty and be able to describe it, and probability is the natural language to describe uncertainty.

This section develops the concepts and language that facilitate discussion of uncertainty and the linkage between uncertainty and probability. Some of the more important rules for calculating with probabilities are reviewed. The most used representations of probability are defined and motivated. Finally, some hard-to-find results on cumulants are recorded for the expert.

Dealing with Complex Informational Relationships

Influence diagrams are used throughout this book to describe our state of information. Influence diagrams are an intuitively clear way of representing this knowledge, even when states of information are related in a complex fashion.

Influence diagrams are mathematical constructs that obey strict mathematical rules. Definitions and rules that are of practical importance are presented and illustrated through several examples.

Obtaining Reliable Information

One task that faces every decision facilitator is obtaining information about uncertainty. And experience has shown that expressing our state of knowledge about uncertainty is not something that we do well.

The decision facilitator must learn to deal with the problems that occur in eliciting information about uncertainty. The causes of the problems are reviewed, and means of correcting for the problems are discussed. A process for conducting a probability encoding interview is described.

Focus of This Book

Since its birth in the 1960s, decision analysis has developed into several different schools, though differences in schools are mostly differences of emphasis and technique. One school focuses on directly assessing probabilities and the different dimensions of value and spends much effort exploring the trade-offs between the uncertain outcomes. Another school focuses more on the art of bringing an assembled group of people to choose a course of action.

This text concentrates on the approach that grew out of the Department of Engineering-Economic Systems at Stanford University and that was

pioneered as a practical methodology at SRI International. This approach is characterized by models that take the burden of estimating values and outcomes in complex situations off the individual's shoulders. A computer model is constructed to reduce a complex problem into manageable components. An influence diagram or decision tree is used to divide uncertainty into subfactors until the level has been reached at which intuition functions most effectively. This modeling approach is especially appropriate in business decisions, where the expertise of many individuals and groups must be combined in evaluating a decision problem.

This book begins with the archetypal decision problem: a single decision-maker using the knowledge of a number of "experts" to make a business decision based on a single principal value (money). The single decision-maker requirement is then relaxed somewhat to describe decisions in organizations. True multiple decision-maker problems, however, go beyond the scope of decision analysis and bring in elements of game theory.

There are other types of problems to which decision analysis has been applied with some success. These applications tend to fall into three areas—personal, scientific, and societal. All three areas are beyond the scope of this book.

Personal decision-making frequently involves difficult and sometimes complex value considerations, such as the life-death-pain-resources trade-offs found in medical decisions.

Scientific decision-making (e.g., the choice of experiments to be funded) also involves special value considerations, since it involves making a controversial comparison of the worth of different scientific results and of the resources required to obtain these results.

Societal decision-making, finally, provides one of the most frustrating and fascinating applications of decision analysis. Not only is there no single decision-maker (but rather a decision-making process), but there is also no single set of values that characterizes society. Rather, conflicting sets of values characterize groups within society. There are even values attached to the process used to make the decision.

In these problems, as in all decision analysis applications, the analysis aims at providing insight into the problem, at opening channels of communication, at showing where differences in values or information do or do not affect decisions, and at directing future efforts in ways that will most improve decision-making.

Problems and Discussion Topics _____

- 1.1 Describe the difference between good decisions and good outcomes.
- 1.2 Describe your own approach to making important decisions. Do you use a systematic approach in making them? Do you try to make decisions in a consistent manner? Have you been satisfied with the major decisions you've made so far (or just happy or unhappy with the outcomes)?

- 1.3 How did you make your decision on which college to attend? Does hindsight reveal any shortcomings in the decision process?
- 1.4 What concerns would you like a decision-making methodology to address?
- 1.5 Describe a decision you currently face in which uncertainty is an important factor. Will you find out the outcome of the uncertainty before or after you make your decision?
- 1.6 Can an uncertainty be an important factor in a decision when the outcome may never be discovered? Describe why or why not and give an example.

Part I

Decision Making

2

Uncertainty and Probability

Probability: A Language of Uncertainty

Uncertainty is a fact of life in the modern world. Both in business and in personal life, there is an almost universal realization that few things can be counted on as certain, at least in the long run. In the business world, recessions come and go, competition comes up with new and unexpected challenges, consumer preferences change (sometimes seemingly at random), accidents or labor problems unexpectedly interrupt business, lawsuits threaten the existence of the company, and so on. And this is true in personal life as well. How will a new job or personal relationship work out? What are suitable investments against the future? In both areas, the list of uncertainties can be expanded indefinitely.

Thinking clearly about these uncertainties—whether to plan better, to make better decisions, or to communicate better about plans and decisions—is important. The key to thinking and communicating clearly about uncertainty is the use of probabilities to describe uncertainty. Fortunately, probability language reflects intuitive concepts of uncertainty. A review of the aspects of probability theory that we will use can be found in Chapter 10. Most readers, however, will find that reading Chapter 10 is not a prerequisite to following the development of ideas in this book.

In this chapter, we concentrate on describing and communicating about uncertainty. In Chapter 3, we deal with how to make decisions under uncertainty.

Along with precise language about uncertainty, we will use two equivalent graphic representations that make it easier to express and to communicate about complex uncertainties. The first of these is the influence diagram. Influence diagrams are an efficient, compact, and intuitive way of representing the uncertainties in a problem and the relationships between the probabilities that describe these uncertainties. At the beginning of this chapter, we use influence diagrams to describe uncertainty.

Later in this chapter, we introduce another representation for uncertainty—the probability tree. These two representations serve complementary purposes. Influence diagrams are a natural way to develop and understand the overall picture; probability trees provide a framework that facilitates calculations with probabilities and the development of insight into the solution of a problem.

The example we will use throughout this chapter and the following three chapters is the semifictional one of Positronics, a manufacturer of sophisticated scientific instruments. This example contains elements from many different cases the authors and their colleagues have worked on. It would be rare for one case to exhibit all these elements. All Positronics discussions are set in italics.

Positronics had decided to bid on supplying MegaCorp with 100 instruments. Positronics estimated it could build them for \$4,000 each; the president of Positronics decided to offer the 100 instruments for \$500,000. Positronics' only real competitor was Anode Industries, a company with the same reputation for reliability and quality as Positronics. For this reason, the president of Positronics was sure the order would be given to the lower bidder.

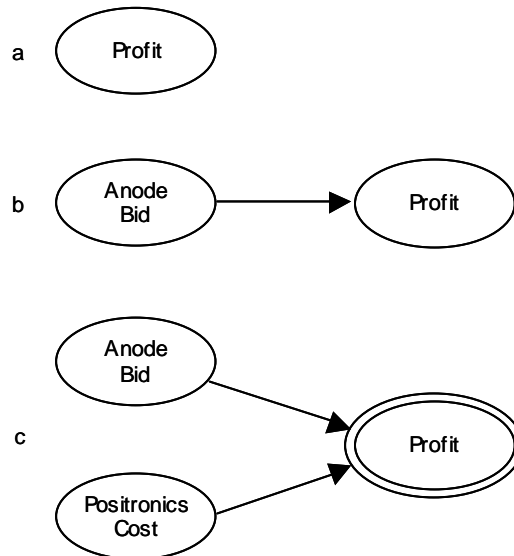
Positronics had formed a team to discuss allocating resources to filling the order should its bid be accepted. At the team's first meeting, it became apparent that everyone was worried about the possibility of losing the bid. The head of marketing put himself on record as saying that it was quite likely that Anode would come in with a bid higher than Positronics'. During the discussion, someone noticed the head of production was looking a little uneasy. Some questioning revealed that he was by no means certain of the costs to produce the instruments and that a winning bid of \$500,000 might well make Positronics lose money.

The team leader decided that there was a deeper problem here than had been expected. To focus the discussion, she elicited from the group an agreement on what they were worried about: everyone was worried about the uncertainty in the profit resulting from the venture. The team leader took a large sheet of paper and wrote the word "Profit" with an oval around it (Figure 2-1a); the oval indicates that profit is an uncertainty. This was the beginning of the "influence diagram" for the problem, a diagram that starts with and focuses on the ultimate, key uncertainty in the problem.

The team proceeded to look for other uncertainties relevant to (influencing) the uncertainty on profit. Most important in most people's minds was the uncertainty in Anode's bid: if Anode were to bid less than \$500,000, Positronics would lose the bid. The team leader added an oval to represent Anode's bid and drew an arrow from the Anode Bid oval to the Profit oval to represent the fact that learning Anode's bid would help answer the question of what profit could turn out to be (Figure 2-1b).

The team then asked if there were other relevant uncertainties. The head of production reminded the team of the uncertainty on production cost. Accordingly, another oval was added to represent the uncertainty on Positronics' cost, and an arrow was drawn from the Cost oval to the Profit oval (Figure 2-1c). At this point, the oval representing Profit was doubled to indicate that, given

Development of the Influence Diagram of the Positronics Bid Venture



Anode's bid and Positronics' cost, there would be no remaining uncertainty—it would be simple to calculate profit.

*The meeting made it clear that there were two key areas of uncertainty to work on: the size of Anode's bid and the amount of Positronics' costs. It was also apparent that if preliminary analysis warranted, the team's task might be extended to finding the optimal level of Positronics' bid. All the other concerns (contingency plans, hiring policy, etc.) would fall into place after these three areas were better understood.**

The influence diagram is a very useful representation of problems involving uncertainty. Not only is it a concise statement of the problem, but it also gives the analyst a valuable tool for finding the structure of problems, for organizing the tasks of an analysis, and for eliciting the data and judgments necessary to analyze an uncertainty. Further, as we will see in Chapter 3, it is straightforward to extend the use of influence diagrams to problems that involve decisions as well as uncertainties.

Three elements of the influence diagram were introduced in the preceding example and are defined below.

- An oval represents an uncertainty. Inside the oval is written a

*The experienced reader might correctly object that things almost never happen this way. Meetings of this sort are often filled with worries and discussions that turn out to be irrelevant to the problem at hand, while the true uncertainties and decisions are often hidden and discovered only by the careful work of someone who listens, observes, and avoids common prejudices and preconceptions.

descriptor (or variable name) to identify the set of events (or the quantity) about which we are uncertain or a question to which we would like an answer. Uncertainties are one of several types of “nodes” that we use in influence diagrams.

- An arrow represents relevance. The arrow between the Anode Bid oval and the Profit oval is read as “Anode Bid is relevant to Profit.” This simply means that if we knew what Anode’s bid turned out to be, it would help us determine what our profit would be. The concept of “relevance” is an important one and will be made more precise in succeeding chapters.
- A double oval represents an uncertainty that ceases to be uncertain once we know how all the uncertainties turn out in the nodes that have arrows pointing to it. This is a “deterministic” node since its value is determined by the nodes that influence it.

The remaining elements concern decisions and will be introduced in the next chapter.

Terminology concerning influence diagrams has shifted during the 1990s. If there is an arrow between two uncertainty nodes, A and B, current best practice is to say “A is relevant to B.” Formerly, common usage was to say “A influences B;” the use of the word “influence” in this context is discouraged because it has the connotation of causality, which is not necessarily present. Influence diagram terminology will be more fully discussed in Chapter 4.

The steps for drawing an influence diagram are summarized below.

1. Determine the one key uncertainty you would like resolved—that is, that you would like an answer to. Write it down and put an oval around it.
2. Ask whether there is another uncertainty that—if you knew how it turned out—would help you resolve the uncertainty you identified in step 1. If there is another uncertainty, write it inside an oval and draw an arrow from this oval to the oval drawn in step 1.
3. Repeat step 2 until all *important* uncertainties influencing the key uncertainty (identified in step 1) are identified. An uncertainty is important in this context only if resolving its uncertainty helps resolve the uncertainty in the key variable. As you repeat step 2, check whether arrows should be added to or from all the uncertainties.
4. Ask whether there are uncertainties that would help resolve the uncertainties identified in step 2. If there are, add them to the diagram. Terminate the process when adding another oval does not help you understand the problem.
5. Check whether any of the uncertainties you have identified are completely resolved (determined) if you have all the information indicated by the arrows. Add another oval around these determined nodes to make a double oval.

Influence diagrams are used throughout this book, and their properties will be introduced as needed. The reader desiring a preview or review of influence diagram theory should consult Chapter 11.

Why Bother with Probabilities?

People often claim that they deal adequately with uncertainty through the ordinary use of language and that the quantitative methods we propose are much too elaborate and unnecessary for any problems except very technical and complicated ones.

However, if we ask people to define just what they mean by such common descriptors of uncertainty as “probably” or “very likely,” we usually find a great deal of ambiguity. The authors and their colleagues have demonstrated this ambiguity in seminars attended by high-ranking executives of many of the country’s large companies. The executives are asked to define a number of terms, such as “likely,” “very likely,” or “probably,” by assigning a range of probability to each term (e.g., “likely” means 60 percent to 80 percent probability). Although the exercise is not rigorously and scientifically conducted, two results occur so consistently that they warrant attention.

1. If we compare the probability ranges of the group for any one term, the ranges vary greatly. If we form some sort of composite group range spanning the individuals’ ranges, there is very little difference between such terms as, for example, “likely” and “very likely.”
2. In any moderate-size group, we find people who assign nonoverlapping probability ranges to the same word. For instance, one person might state that “unlikely” means 20 to 30 percent while another might put it at less than 5 percent. While these people may think they are effectively communicating about the likelihood of an uncertain event, they actually have quite different judgments about the likelihood and are unaware of the differences.

Positronics was just beginning to use decision analysis and had just discovered how differently participants at the meeting were using the same terms. It turned out that when the head of marketing said it was “quite likely” that Anode would bid higher than Positronics, he had meant that there was a 60 to 80 percent chance Anode would bid higher. When the president of Positronics heard this, he was more than a little disturbed, since he had assumed that “quite likely” meant something on the order of 90 to 95 percent. The lower probability argued either for more contingency planning by the production staff or for a lower bid. Of course, a lower bid would increase the likelihood of the newly perceived possibility that high production costs could lead to a loss for Positronics.

Verbal descriptions of uncertainty tend to be ambiguous or ill-defined, but numerical probability statements clearly and unambiguously describe uncertainty. Yet many people are overwhelmed with the idea of using

probability. After all, probability is a rich and complicated area of study. However, conducting a decision analysis requires knowing only a few simple properties. While applying probabilities in decision analysis is at times subtle, common sense and careful reasoning are usually more useful than technical sophistication.

What Are Probabilities? _____

Before we proceed further, let us pose a somewhat philosophical question: what do probabilities represent and where do they come from?

Probabilities represent our state of knowledge. They are a statement of how likely we think an event is to occur.

This simple but profound concept can be illustrated by the example of three people considering an oil-drilling venture in a new area. One person is the president of the company. He has been around for a long time, and his experience tells him that there is a 20 percent chance that there is recoverable oil in the area. The second person is a technical expert who has just finished studying the most recent seismic and geological studies on the area. She assigns a 60 percent probability that there is oil there. The third person is someone out at the drilling site; they have just struck oil, and he would assign a 100 percent probability to finding oil.

Who is right? They *all* are right, assuming that they all are capable of processing the knowledge available to them. (Fortunately, decision-makers in business situations are rarely incompetent or incorrigibly optimistic or pessimistic.) The general acumen of the president, the technical and statistical data of the analyst, and the simple observation of the person at the site—these are different sets of knowledge.

This view of probability as an expression of a state of knowledge has profound consequences. We comment on it in one way or another in every chapter of this book. This perspective on probabilities is called the Bayesian point of view.

Frequently, the nervous decision-maker will ask, “What are the correct probabilities?” Unfortunately, there are no “correct” probabilities. Probabilities represent the decision-maker’s state of knowledge or that of a designated expert source of knowledge. Probabilities thus represent a person’s judgment and experience and are not a property of the event under consideration. We can worry about whether the expert has good information, whether the probabilities represent the information adequately, whether the state of information should be improved, or whether efforts should be made to elicit the probabilities better. These questions examine the quality of the probabilistic information, but they are not attempts to get the “correct” probabilities.

We do not, however, imply that data are irrelevant. One of course wants the best data and statistics available and to consider the advice of the most experienced people available (another form of data). Unfortunately, the hard decision problems usually involve uncertainties for which definitive data are either not available or not completely relevant. And, when the decision needs

to be made, the decision-maker must consider how relevant even the best data are to predicting the future.

Using Intuition Effectively

The “Divide and Conquer” Approach

To help the individual think about uncertainty, decision analysis makes judicious use of the “divide and conquer” approach. To do this, the overall uncertainty is divided into a reasonable set of component uncertainties (and, as we will see in Chapter 3, decisions), which are then treated individually. We have seen an example of this approach in the development of the influence diagram for the Positronics case.

This approach reduces the complexity and scope of the problem to a level at which intuition can function effectively. This is important because, as discussed above, probabilities are statements derived from a person’s state of knowledge. It is very difficult (if not impossible) for a person to give meaningful probabilities on an uncertainty that is complex or that includes factors beyond his or her immediate knowledge and experience.

Dividing the uncertainty is also useful in helping different people in a company communicate about the uncertainty and contribute to the analysis. Since different people bring different expertise and experience to the decision-maker or decision-making process, dividing the uncertainty into small pieces enables each individual to contribute precisely within his or her area of expertise.

In the case of Positronics, the problem already showed a natural division into two subcomponents: the uncertain size of Anode’s competitive bid and the uncertain cost of production if Positronics won the contract. It seemed natural that the uncertainty in the production cost would be best estimated by someone from the manufacturing or production staff, while the uncertainty on the size of the competitive bid would best be estimated by someone from marketing or upper management. Assignments for further study were made accordingly.

Passing the Clairvoyance Test

One of the most common barriers to the use of intuition and to effective communication is lack of clarity in defining the event to which we are assigning probabilities.

To test the clarity of definition, we use the clairvoyance test. The clairvoyant is a hypothetical person who can accurately answer any question, even about the future, but who possesses no particular expertise or analytical capability. Thus, for instance, when asked about production costs in the year 2010, the clairvoyant might “look” at a company’s annual report for the year 2010 and report the answer he sees there.

The clairvoyance test is a mental exercise to determine if the clairvoyant can immediately tell us the outcome of a chance node or if he needs to know

other things first. An immediate prediction means that the uncertainty is clearly defined. If the clairvoyant would have to ask some questions first, then we have not clearly laid out exactly what the uncertain quantity is.

For instance, imagine we ask the clairvoyant what production costs will be in the year 2010. The clairvoyant would have to ask whether the costs are in today's dollars or 2010 dollars; whether they include or exclude depreciation, fixed costs, and allocated expenses; whether they are given in terms of an advance in technology or an evolution of the product; etc. Our response, of course, is to define better what we mean by production costs so they more nearly pass the clairvoyance test.

The clairvoyance test is surprisingly difficult to pass. Yet if it is not passed, we will find that uncertainty compounded with an ill-defined quantity yields results of dubious quality at best. The problem becomes even more acute when information obtained from different people is compared. The time spent in taking—and passing—this test is well rewarded by a lack of confusion and greater insight later on.

The Positronics staff agreed on definitions for the uncertainties that passed the clairvoyance test. Cost was a well-defined measure that excluded depreciation, allocated costs, and truly fixed costs. Anode's bid was interpreted strictly in terms of the deliverables called for in the request for bids.


Assigning the Numbers

Using Trees

To put explicit values on probabilities, we use drawings called “distribution trees” to represent the data in the nodes of the influence diagram—the information on how each uncertainty may turn out. It is called a distribution tree because it has a line (branch) for each possible outcome. In statistical terms, this is the distribution of possible outcome. The data in this form are then combined into a “probability tree” that both graphically describes and

Figure 2-2

Skeleton of a Three-Branch Distribution Tree



numerically analyzes the problem. Each uncertainty node we have identified in the influence diagram becomes an uncertainty or chance node of the probability tree.

What does a distribution tree look like? Shown in Figure 2–2 is the skeleton of a three-branch distribution tree. At the branching point on the left is a circle to indicate that this is an uncertainty or chance node. Associated with each branch is an outcome. The set of outcomes describes all the different events that could occur at that node—all the ways the uncertainty could be resolved.

The set of outcomes at a chance node must be mutually exclusive and collectively exhaustive. That sounds more complicated than it really is. Mutually exclusive means that the way an uncertainty turns out must correspond to only one of the outcomes (i.e., there is no overlap in the list); collectively exhaustive means that the way the uncertainty turns out must correspond to one of the outcomes (i.e., all possible outcomes are included in the list).


The Data in the Nodes

*After some work, the Positronics staff came up with the following probability assessments to represent their perception of the uncertain future. To keep things simple, the staff judged it sufficiently accurate to divide both quantities (Anode's bid and Positronics' cost) into three ranges, thus creating three discrete outcomes for each node.**

First, the three possible states representing Anode's bid were chosen. The Anode bid could be less than \$500,000, between \$500,000 and \$700,000, or greater than \$700,000. The marketing personnel used their judgment to assign the following probabilities: a 35 percent chance the Anode bid will be less than

Figure 2–3


The Anode Bid Node



*Chapter 12 gives much more careful procedures for eliciting probabilities. The informal process used here is appropriate only when an uncertainty is unusually well understood.

Figure 2-4

The Positronics Cost Node



\$500,000, a 50 percent chance it will be between \$500,000 and \$700,000, and a 15 percent chance the bid will be greater than \$700,000. These definitions and probabilities can be thought of as being contained "inside" the Anode Bid oval in the influence diagram. Shown in tree form (Figure 2-3), the outcomes are on each branch, the probability of the branch occurring is at the left of each branch, and the circle at the branching point indicates that it is an uncertainty node. To keep notation compact, all values in the figures are shown in units of thousands of dollars—e.g., \$500,000 is shown as 500.

Similarly, Positronics' cost could turn out to be less than \$300,000, between \$300,000 and \$500,000, or greater than \$500,000. There is a 25 percent chance Positronics' cost will be less than \$300,000, a 50 percent chance it will be between \$300,000 and \$500,000, and a 25 percent chance it will be greater than \$500,000 (Figure 2-4).

The range of values assigned to Positronics' cost is very large, with about a factor of two (<300 to >500) uncertainty. It is rare to find so great an uncertainty on costs, except in one-of-a-kind construction or in very large projects.

Figure 2-5

Influence Diagram Prepared for Drawing the Tree


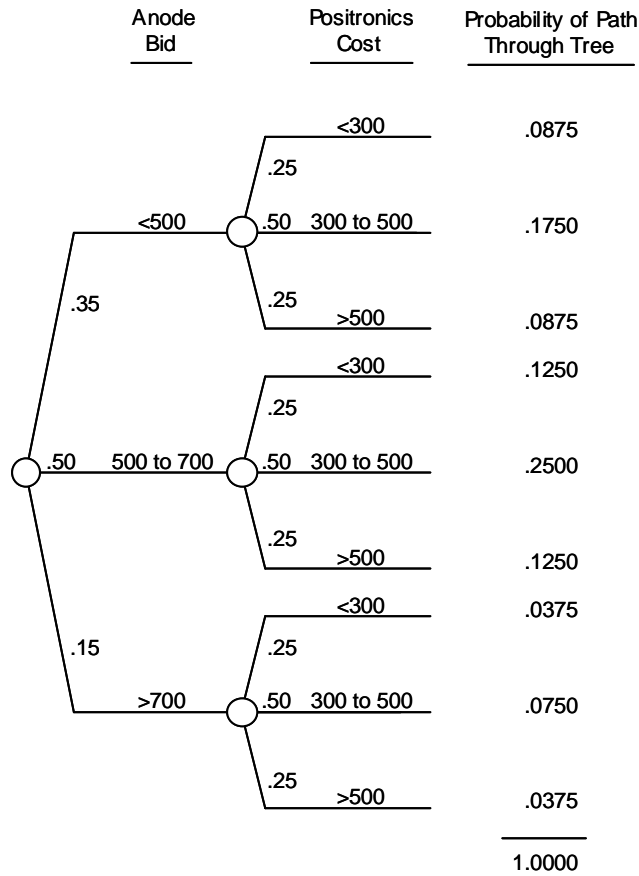


Figure 2-6

Probability Tree Displaying the Probability of Each Path Through the Tree



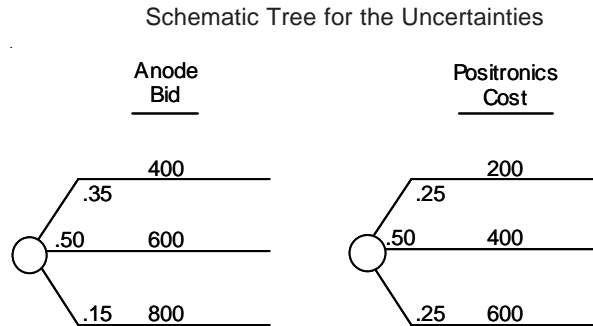
Drawing the Probability Tree

All these data now are put together to create a “probability” tree. This probability tree represents all the different combinations of events that can occur and their probabilities. How are the data joined together to create a probability tree?

Pick an order in which to display the two uncertainty nodes. The order is arbitrary for the present example. The rules applicable in more complicated cases are presented in the next chapter. In Figure 2-5, we choose Anode Bid to be the first node and Positronics Cost to be the second. The deterministic node, Profit, will be at the right-hand side of the diagram.

At the end of each of the branches of the first node (Anode Bid), attach the second node (Positronics Cost), as in Figure 2-6. We will deal with the deterministic node (Profit) shortly.

Figure 2-7



The final column in Figure 2-6 has been added to clarify the meaning of the tree. The probability of each of the nine scenarios (paths through the tree) has been calculated by multiplying the probabilities that occur at each branching. These products are referred to as joint probabilities. Thus, the scenario of Anode bidding less than \$500,000 and Positronics' cost being less than \$300,000 (the topmost path through the tree) is $.35 \times .25 = .0875$. The sum of all the joint probabilities is 1, as it must be.

We can see that however difficult it is to assign the probabilities for each chance node, it would be much more difficult (if not impossible) to directly assign the joint probabilities. The "divide and conquer" technique reduces scope and complexity so that judgment and intuition work effectively.

You can conserve space by drawing the tree in the more compact form shown in Figure 2-7. This form of the tree is often referred to as a "schematic tree." Drawing the tree with nodes following each other like this means that each branch of the left node leads to the next node on the right. Thus, the implication is that each node is duplicated as many times as necessary to produce the full tree structure.

A refinement has been added in Figure 2-7. The outcomes on each branch are now single numbers that represent the ranges of the original outcome. For instance, the Anode bid can be \$400,000, \$600,000, or \$800,000. We have approximated the original ranges by discrete values, resulting in the Anode bid having only three values. (The procedure for doing this is described at the end of this chapter.) Clearly, the more branches we have, the narrower we can make the ranges and the better a single number can represent the range. On the other hand, the more branches, the larger the tree.

A Value Function

We still have not dealt with the deterministic (double oval) node, Profit, which appears in the Positronics influence diagram (Figure 2-5). This is the node for which there is no remaining uncertainty, once we know the outcome of the nodes that influence it. Its value is determined by the outcomes of the influencing nodes. Thus, unlike the other two nodes, the profit node needs

Figure 2-8

Influence Diagram Showing Equation "Inside" the Profit Node





Figure 2-9

Completed Probability Tree



only a value function (a formula) that translates the uncertainties in the chance nodes into the uncertainty on the value of interest to the decision-maker. (We discuss value functions in detail in Chapter 3.)

Positronics had decided that profit was the prime value in this case. The deterministic node for profit can therefore be thought of as having an equation “inside” the node (Figure 2–8).

The completed probability tree is shown in Figure 2–9.

The values in the column on the right are needed to complete the probability tree and allow analysis. Sometimes the value model comes from a subjective assessment where the values for each scenario are assessed directly rather than calculated. Sometimes the value model is a simple relationship (as in the Positronics case) where the value can be calculated by a single equation. However, most often the value model is implemented as a computer program (or spreadsheet) that calculates how revenues and costs develop over time, figures tax effects, calculates a cash flow, and so forth.

Analyzing the Tree

Now that the tree has been completed, we will determine what information can be drawn from it. First, we will examine how the component uncertainties have translated into uncertainty on the value of ultimate interest to the decision-maker.

Because a cumulative probability plot (like the one shown in Figure 2–10) is the most efficient means of presenting information for decision analysis, we use it throughout this book. The plot is worth studying closely, since it answers most of the questions the decision-maker will ask.

The plot gives the probability (the vertical axis) that the venture’s value will turn out to be less than or equal to the value shown on the horizontal axis. Thus, the plot shown in Figure 2–10 indicates that there is:


- No probability of the venture being worth less than $-\$100,000$
- A .1625 probability of its being worth less than or equal to $-\$100,000$ (.1625 that it is equal to $-\$100,000$)
- A .5125 chance of its being worth less than or equal to $\$0$ (.1625 that it is equal to $-\$100,000$ and .35 that it is equal to $\$0$)
- A .8375 probability of its being worth less than or equal to $\$100,000$
- A 1.0 probability of its being worth less than any value above $\$300,000$.

The data for this graph is contained in Figures 2–6 and 2–9. The procedure for constructing the cumulative graph is simple:

1. Create a list of the values in the tree, ordering them from smallest to largest. For Positronics, Profit is the value.
2. Next to each value, put the probability associated with the value. Some values may occur several times in the tree; for these, add the probabilities together.

Figure 2-10

Cumulative Plot of Probability Distribution on Profit in the Positronics Venture



3. In the space labeled cumulative probability, enter the sum of the cumulative probability immediately above and the probability to the left. For the top row in the list, use zero for the cumulative probability immediately above. (The cumulative probability is the sum of all the probabilities on and above each line.)

Value	Probability	Cumulative Probability
-100,000	$.1250 + .0375 = .1625$.1625
0	$.0875 + .1750 + .0875 = .3500$.5125
100,000	$.2500 + .0750 = .3250$.8375
300,000	$.1250 + .0375 = .1625$	1.0000

4. Create a graph with Profit on the horizontal axis and probability on the vertical axis.
5. For each line in the list, mark the point corresponding to that value of Profit on the horizontal axis and the corresponding cumulative probability on the vertical axis.
6. For each line in the list, mark the point corresponding to that value of Profit on the horizontal axis and the cumulative probability for the *previous* value of Profit on the vertical axis. For the first value of Profit in the list, use zero as the cumulative probability for this point.

7. Join the points with horizontal and vertical lines.

The staircase nature of the graph is, in this case, imposed partly by the discrete approximations used to turn an originally continuous range of bids and costs into nine (3×3) different cases. If the ranges had been broken into more segments, the distribution would have been smoother. But there would still have been a step at \$0 corresponding to the probability that the bid was lost.

There is another way to represent the probability distribution for the venture's value: the histogram plot. In a histogram plot, the horizontal (value) axis is divided into a number of equal-sized ranges or bins. The height of the bar drawn for each bin is the probability (vertical axis) that the value falls in the bin.

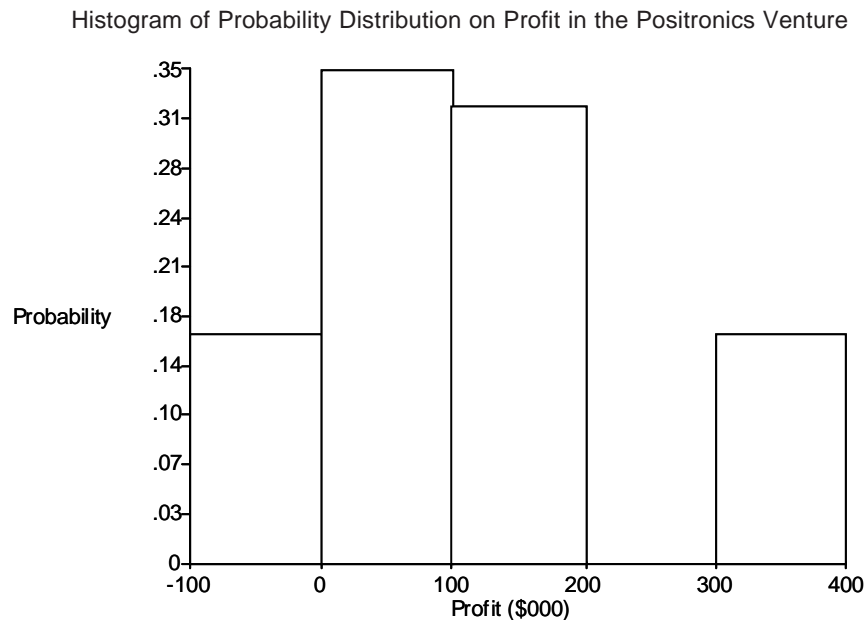
The histogram in Figure 2-11 has five bins ranging from $-\$100,000$ to $\$400,000$. The plot reveals that there is:

- A .1625 probability of the bin ranging from $-\$100,000$ to $\$0$
- A .35 probability of the bin ranging from $\$0$ to $\$100,000$
- A .325 probability of the bin ranging from $\$100,000$ to $\$200,000$
- A .1625 probability of the bin ranging from $\$300,000$ to $\$400,000$.

The data for this graph is contained in Figures 2-6 and 2-9. The procedure for constructing the histogram is simple:

1. Define the equal-size ranges for the value and create a list with as many lines as there are ranges.
2. Determine the range into which the value at the end of a branch falls. By convention, a value falls with a bin if its value is greater

Figure 2-11



than or equal to the lower bound and less than the upper bound.

3. Add the probability for the branch to the line in the list for the range.

Lower Bound	Upper Bound	Probability
-100,000	0	$.1250 + .0375 = .1625$
0	100,000	$.0875 + .1750 + .0875 = .3500$
100,000	200,000	$.2500 + .0750 = .3250$
200,000	300,000	0
300,000	400,000	$.1250 + .0375 = .1625$

4. Create a graph from these ranges and probabilities.

How Much Detail Is Enough? _____

In an actual case, the value model might be more complex, and there might be more nodes in the tree or more branches on the nodes. While such complexity may be appropriate, more detail may not add insight. Hence, the results shown in the cumulative and histogram plots may be close to representing the essentials of the situation. Before succumbing to the temptation to make an analysis more complicated, we should pose the following questions:

- Will a more complicated value model really change the values that much?
- Will more nodes or more branches at a node really change the probability distribution that much?
- Will the added detail really add insight?
- Is the focus shifting from “a model adequate to make a choice” to a “good model of reality”?

Often, the essential insight into a decision problem is conveyed by results almost as simple as those shown here. That a decision has proved much simpler to grapple with than one would have supposed is an insight a decision-maker really appreciates.

With the initial analysis completed, the decision analysts were preparing to present their results to the president of Positronics. They decided to avoid puzzling him with the complications of understanding a cumulative probability graph and instead used the histogram. During the presentation, the president began asking questions like “What’s the chance we will make a substantial profit?” and “What’s the chance we will lose money?” which the analysts realized could be read directly off the cumulative probability graph. Being experienced presenters, they had a copy of the cumulative probability graph ready in the set of backup materials they had prepared for the presentation.

To everyone’s astonishment, the president was genuinely surprised by the results. After some discussion, everyone agreed that he was surprised because the explicit judgments inherent in the probabilities had opened up new

channels of clear communication. For one thing, the president had not appreciated just how worried production was that the cost of producing the instruments might be much higher than expected. The analysis also enabled the president to see how the uncertainties in cost and in Anode's bid combined into an overall uncertainty. Looking at the cumulative and histogram plots, he now had a good idea of the uncertainty and risk in his current bid strategy. He did, however, ask whether there was a single number that could be used to characterize the worth of the venture when he was comparing it with other ventures or speaking to members of the board (who were probably not interested in details of this rather minor facet of the business).

Certain Equivalent and Expected Value _____

One of the goals of decision analysis is to give a single number that characterizes a probability distribution. This number is called the certain equivalent and is the value that, if offered for certain, would represent to the decision-maker an even exchange for the uncertain venture described by the probability distribution. (This topic is discussed extensively in Chapter 5.) There is, however, a quantity closely related to the certain equivalent but easier to use and understand: the expected value. The uncertain venture is described by a probability distribution—a set of possible outcomes and the probability that each outcome occurs. Figures 2–10 and 2–11 show the probability distribution for profit in graphical form. The expected value of a probability distribution is obtained by multiplying each outcome by its probability of occurring and then by adding the products. The data plotted in figures 2–10 and 2–11 are shown in the trees in figure 2–9.

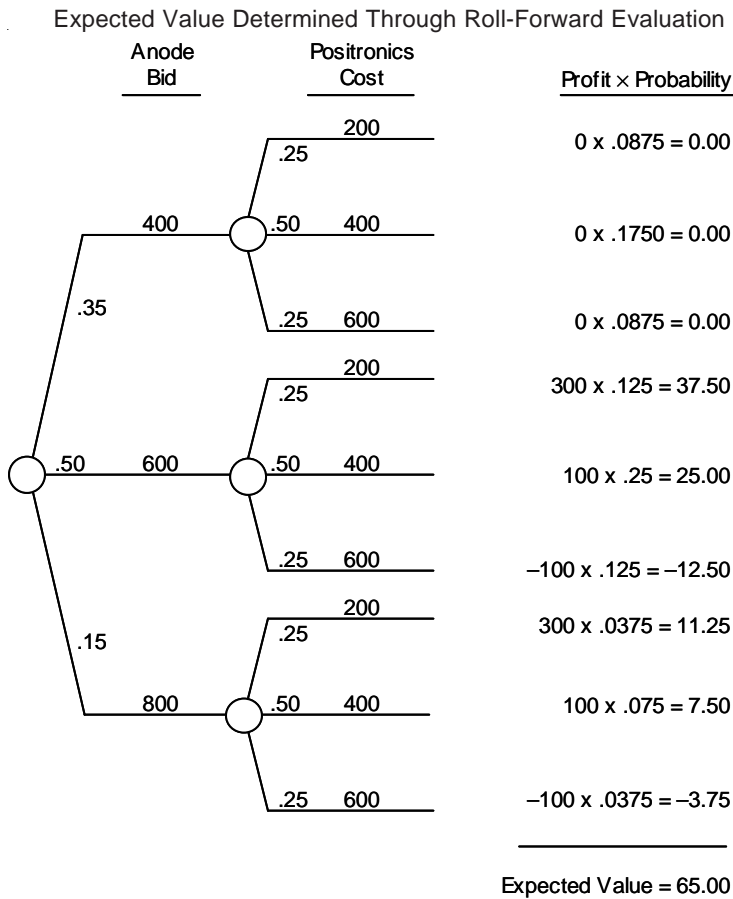
The data in these trees can be combined in what is called the “roll-forward” evaluation (Figure 2–12). The roll-forward evaluation rolls all the information forward—i.e., to the right-hand side of the tree. This means that for each path through the tree, there is a value and a probability (Figure 2–9). Multiply the value by the probability at the end of each path and add the results to obtain the expected value.

The expected value can also be obtained through the “rollback” evaluation, a much more general technique that is easier to use. In the rollback, we start at the right-hand edge of the tree and replace each chance node with its expected value. (Multiply the value on each branch by its probability and sum the products.) We then move to the left and continue the process, except that what was the rightmost column of nodes has now been replaced by expected values. This procedure continues until all the nodes have been replaced and we are left with the expected value of the tree (Figure 2–13)

In Figure 2.13, we see how the Cost node expected values are the branch values for Anode Bid, which lies to the left of Cost. For instance, the branch with Anode Bid equal to 800 replaces the Cost node with its expected value, which is calculated below.

$$(300 \times .25) + (100 \times .50) + (-100 \times .25) = 100 \quad (2-1)$$

Figure 2-12



The expected value of the Anode Bid node is then:


$$(0 \times .35) + (100 \times .50) + (100 \times .15) = 65 \tag{2-2}$$

Given a large number of identical uncertain ventures, the average value would be the expected value. For instance, in a coin toss that yields a prize of \$1, the expected value is:

$$(.5 \times \$1) + (.5 \times \$0) = \$0.50 \tag{2-3}$$

We would expect to be about \$50 richer after flipping the coin 100 times. Of course, the problem is that we are often presented with one-of-a-kind situations. For the coin toss example, if we are offered a single coin flip with a prize of \$1, we will either walk away with \$1 or with nothing. Still, it is not unreasonable to value the opportunity at the \$0.50 expected value for decision-making purposes. For much larger values (e.g., a single coin flip with a prize of \$1,000), the expected value may not adequately give us the value of the uncertain venture. If it does not, we should use the probability distribution. For most problems we will encounter in practice, the certain

Figure 2-13



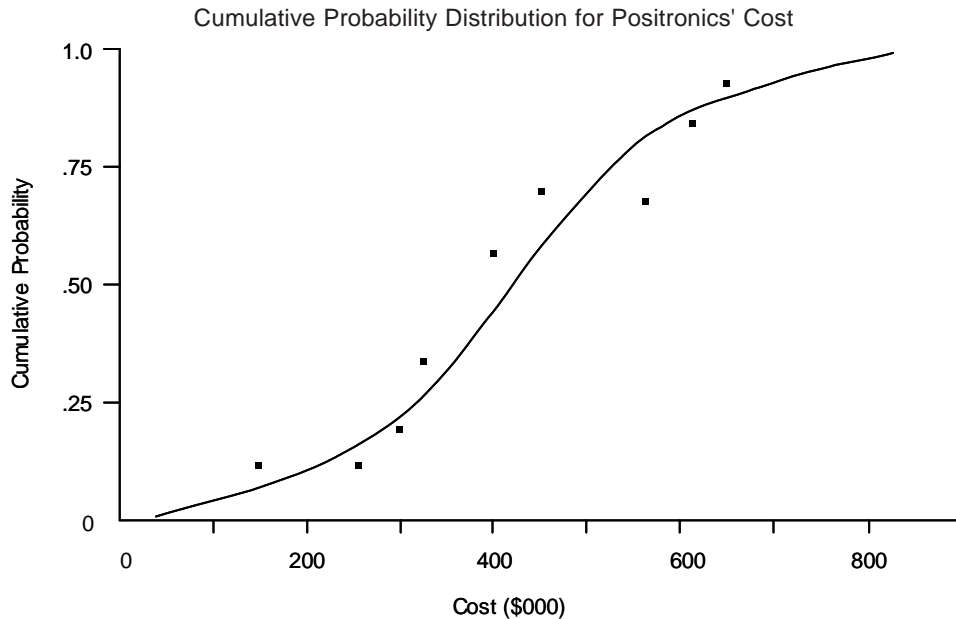
equivalent will not be very different from the expected value. In addition, the methods of tree evaluation are similar for both measures. Thus, it will prove simple to make the transition from one measure to the other.

Encoding Probabilities*

How does one obtain probabilities in practice? Obtaining probability distributions for use in decision analysis is no easy task. People are not trained in thinking about uncertainty, so the exercise of assessing probabilities can be uncomfortable and difficult. In this section, we briefly describe the encoding process and how to discretize the resulting continuous distributions. In Chapter 12, we further discuss the encoding process and give an actual procedure to follow. We also further discuss the commonly encountered biases that must be counteracted (as much as possible) to obtain

* Reading this section may be postponed until after Chapter 5 has been read.

Figure 2-14



accurate information.

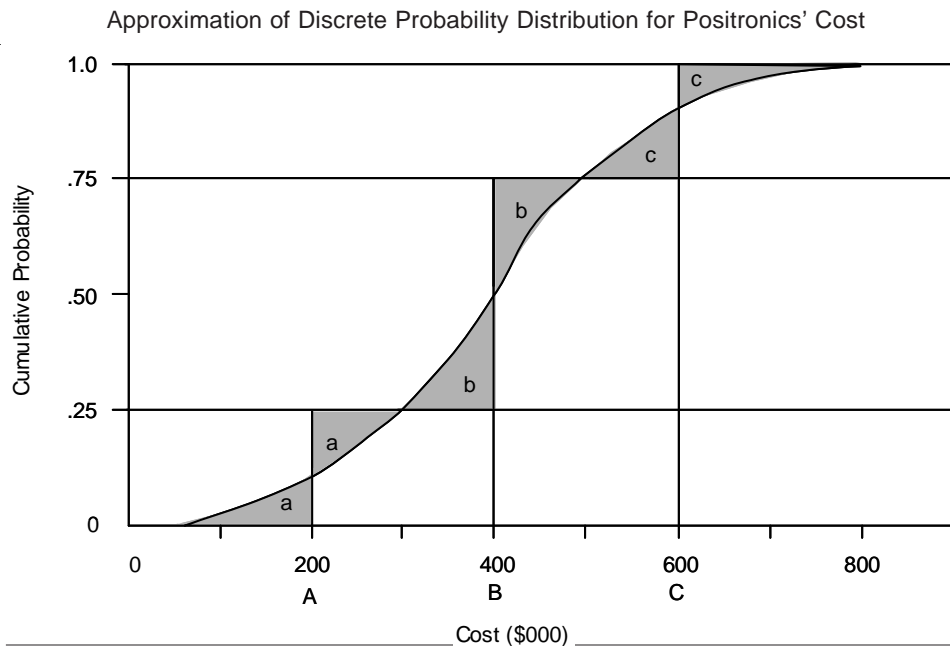
To obtain probabilities, a tree can be set up and people can be directly asked to write probabilities on it. Another quick way to get a probability distribution is from existing “range data.” Unfortunately, while both these procedures are easy and take relatively little time, they can yield “bad” probabilities if the analyst does not take some care to counteract common biases. Earlier, we emphasized that probabilities represent a person’s state of knowledge and that there are no “correct” probabilities. Biased probabilities are those that do not adequately represent what a person really knows; something prevents him or her from using or expressing this knowledge correctly.

The most common types of bias cause people to think they know things better than they do. In other words, their probability distributions are usually much too narrow and understate the uncertainty. The best way to counteract this bias is, if possible, to address extreme outcomes (“How high could the value be? What would have to happen for the value to be even higher?”) before getting base case, best guess estimates.

Another type of bias occurs because most people are unfamiliar with probability theory and make false analogies or draw false conclusions about probability. The best way to counteract this bias is to be clear and explicit in drawing influence diagrams and trees.

Obtaining the Data

Figure 2-15



In a complete encoding of continuous numerical variables, the decision facilitator usually assesses the cumulative probability distribution directly. This process will allow the facilitator to counteract the biases during the assessment. The cumulative probability graph is obtained by plotting responses to such questions as “What is the likelihood that costs will be less than \$500,000?” or “Less than \$650,000?” (Actually, it is best to ask these questions using a reference device such as a probability encoding wheel—see Chapter 12.) The result of such an assessment process will be a set of points and a smoothed curve such as the one in Figure 2-14 for Positronics’ cost.


Discretizing the Data

In tree analysis, we usually make a discrete approximation by setting up nonoverlapping (mutually exclusive) ranges that encompass all possible values (collectively exhaustive), by finding the probabilities that the values fall in these ranges, and by then choosing a value to represent that range. (This is the approximation.) By doing this, we have converted a continuous variable to a discrete variable and a probability density function to a probability mass function.

Given a continuous probability distribution such as the one shown in Figure 2-14, how does one perform this approximation? One widely used technique is to select the number of outcomes and the values of the probabilities you want and then draw a horizontal line at these probabilities. In Figure 2-15, we have chosen the number of outcomes to be three. We have also chosen the probabilities .25 for the lower range (line at .25), .5 for the

Figure 2-16

Discrete Probability Distribution in Tree Form



middle range (line at $.25 + .50 = .75$), and $.25$ for the top range (line at $.25 + .5 + .25 = 1$).

Next, we draw a vertical line at A, choosing point A so that the shaded area to the left of the vertical line is equal to the shaded area to the right. (The eye is surprisingly good at doing this.) These two areas are marked by the letter “a.” Then we pick a point, B, at which to draw a vertical line with the shaded area to the right being equal to the shaded area to the left. Finally, we pick the third point, C, at which to draw the vertical line balancing the two shaded areas.

The procedure sounds much more complicated than it is in practice. The result is that we now have approximated the continuous probability distribution; the discrete probability distribution is shown in tree form in Figure 2-16. The actual values are $A = 200$, $B = 400$, and $C = 600$. These values are used for Positronics’ cost in this chapter. In general, the values for A, B, and C will not come out evenly spaced.


The reason the procedure works is that we divided the continuous probability distribution into ranges with associated probability when we drew the horizontal lines. In Figure 2-17, we see that the first range was from negative infinity to x and had probability $.25$. The second range was from x to y and had a probability of $.5$. The third range was from y to infinity and had a probability of $.25$. (For this example, $x = 300$ and $y = 500$, corresponding to the ranges in Figure 2-4.) Picking point A in such a way that the shaded areas are equal is a visual way of finding the expected value, given that you are in the lowest range. (Proving that the expected value makes the shaded areas equal is a nice exercise in calculus in problem 2.15.) Choosing the expected value to represent the range is a natural approximation and is commonly used. There are, however, other possible choices.

Using Range Data

We see that the final result was three numbers, with probabilities of $.25$, $.50$, and $.25$. There is a shortcut for assessing these values directly. This method is frequently used in the early stages of analyzing a problem when you have obtained ranges within which the values of variables are expected to fall. (See Chapter 6 for a discussion of Sensitivity Analysis.) The ranges are typically defined by best guess (50 percent chance it could be less), low value (10 percent chance it could be less), and high value (90 percent chance it could

Figure 2-17

Ranges Associated with Approximate Discrete Probability Distribution for Positronics' Cost



be less or 10 percent chance it could be greater). As a quick way of setting up a tree, these three values can be used for B, A, and C, respectively. But remember that you want to encode probabilities more carefully later on!

Figure 2-17 shows why this procedure works. There is about a 10 percent chance the value will be less than A and about a 10 percent chance that the value will be greater than C—and this is how we defined the ranges! (This property is almost exact for a Gaussian or normal distribution and is quite good for most of the distributions you find in practice.) Finally, the best guess can be used for B unless you know the distribution is very asymmetric. Thus, in summary, we can take the “10 percent, best guess, 90 percent” range numbers and use them in a tree with .25, .50, and .25 probability, respectively.

Beware, however, of “best guesses” that come from a business plan and are far from the median value. Also, be alert for ranges that are too narrow or are given without sufficient thought and reflection.

Summary

Probabilities are a precise way of expressing a person’s information about an uncertain event. Judicious use of the clairvoyance test and the divide and conquer approach make the use of probabilities an intuitive as well as precise process.

Influence diagrams and probability trees constitute a framework to hold and process probability statements. The influence diagram is a powerful tool

for structuring a problem. The probability tree is the tool used most often for calculating with probabilities.

The primary output of the analysis is the probability distribution on value. This information can be represented graphically by a cumulative probability plot or by a histogram. For the Positronics case, the graphs showed the uncertainty in profit.

A single measure to represent the value of an uncertain venture is the expected value of the probability distribution. For the president of Positronics, this number provided a starting point for putting a value on his uncertain venture.

Finally, care and skill are required in the assignment of probabilities. A simple example showed how information on Positronics' cost can be used for input to the probability tree.

Problems and Discussion Topics

- 2.1 Consider today's weather forecast. Are the chances of rain or sunshine expressed verbally or with probabilities? What is your probability of rain given that weather forecast? If your probability is different from the forecast, is the difference because you and the forecaster have different states of knowledge or because of some other reason?
- 2.2 In the section "What Are Probabilities?" are statements like "Probabilities represent our state of knowledge." Such statements are sometimes misinterpreted to mean that probabilities are arbitrary numbers between 0 and 1. In fact, probability is a well-defined concept with very strong implications. For example, if two people have exactly the same information (knowledge) about an event, they should assign the same probability to this event. Furthermore, if new relevant information becomes available, the prior probability assignment will have to be changed (updated).

What else can you infer from the statement "Probabilities represent our state of knowledge"?
- 2.3 Why is an influence diagram (or similar method) necessary for understanding complex uncertainties? How do the procedure and graphical form of an influence diagram deal with the problem?
- 2.4 What is the relationship between each component of an influence diagram (arrows, ovals, and double ovals) and the components of a probability tree?
- 2.5 Can you draw a probability tree directly without first drawing an influence diagram? When would this be a bad or good idea? Does your answer depend on the level of expertise of the person doing the analysis?
- 2.6 In the section "Using Intuition Effectively", we discussed how to define uncertainty clearly and the role of the clairvoyant in the clairvoyance

test. How can the clairvoyant help you? Can the clairvoyant change your future?

- 2.7 What do you do when an uncertainty fails the clairvoyance test? How might this change your influence diagram? How might this change the structure of your probability tree?
- 2.8 The expected value is sometimes described as the mean value you would expect to achieve if you undertook the same venture many times. Unfortunately, since many decisions are one-of-a-kind decisions, there is no opportunity to repeat them and establish a historical mean. Suppose, though, that a venture had just been resolved (all the uncertain events had happened) and you were now faced with an identical one. Would your prospective expected value for the second venture be the same as it was for the first? Why or why not?
- 2.9 You are going to the movies tonight with a new date. You plan on treating, but your date may want to go Dutch treat (each person pays) or treat you. You figure the three outcomes are equally likely. The cost for the movie is \$5 per person. You plan to at least buy popcorn if your date wants it (with a 4 out of 5 chance that he or she will). However, you have forgotten how much the large popcorn costs. You would give 5 out of 10 that it costs \$2 and split the rest of the probability between \$1.50 and \$2.50.

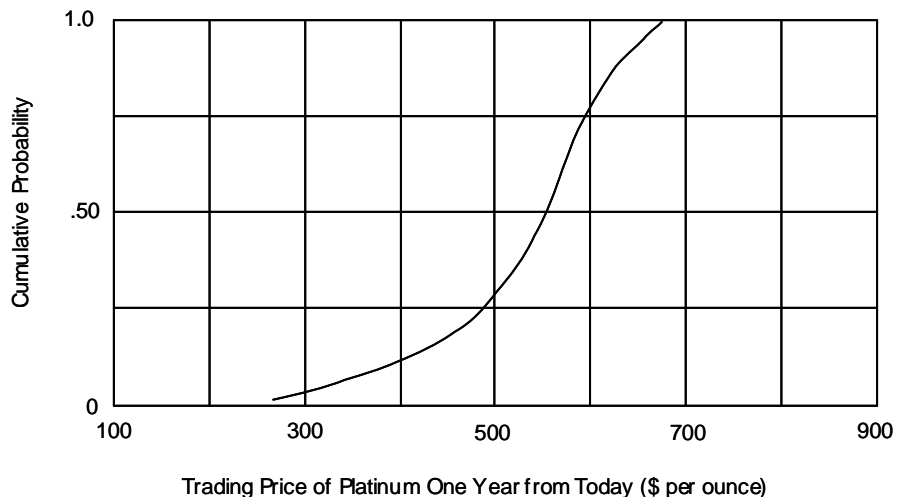
You just discovered that you only have \$10 cash right now. What is the expected cost of going to the movie tonight? What is the probability that it will cost you more than \$10? What is the probability that it will not cost you anything?

- 2.10 Your prize Rapid Ripe tomato plant has flowered and is ready to start producing fruit. If all goes well, your plant will produce tomatoes by the end of the week. It will then produce a new set of flowers and blossoms next week. Unfortunately, your area is subject to blossom wilt, which causes the tomato flowers to fall off. If the blossoms fall off, a new set of blossoms will not emerge until next week, and tomatoes will not be ready until the end of that week.

Luckily, each time blossoms fall off, the plant builds up its resistance; the probability of each succeeding blossom falling off is then only half as much. You estimate that the probability of this first set of blossoms falling off is .40.

- a. Draw the influence diagram for this problem and then draw the probability tree.
- b. What is the probability tomatoes will be ready in the third week?
- c. What is the expected number of weeks you will have tomatoes over the next three weeks?

- d. What is the probability you will lose the blossoms one or more weeks of the next five weeks?
- 2.11 You have discovered the lights on Van Ness Avenue in San Francisco are synchronized, so your chances of getting a green light at the next intersection are higher if you have not been stopped by a red light at the last intersection. You estimate that there is a $4/5$ chance of getting a green light if the previous light was green. Similarly, there's only a $1/4$ chance of getting a green light if the last light was red. Because Van Ness is a major thoroughway, you estimate there's a $2/3$ chance of the first light being green.
 - a. Draw the influence diagram for this problem and then draw the probability tree.
 - b. What is the probability that the second light will be green?
 - c. Out of the first three lights, how many lights can you expect to be green?
 - d. What is the probability of the third light being green when the first one was red? (Hint: draw a tree with three nodes representing the first three lights.)
- 2.12 An excess probability distribution plots the probability that the value is greater than a given number. Plot the excess probability distribution for the initial Positronics tree (Figure 2-9). What is the probability the value is greater than \$130,000?
- 2.13 Plot the probability mass function for the initial Positronics tree (Figure 2-9). What is the difference between the mass function and the histogram? (The probability mass function is defined in Chapter 10.)
- 2.14 Suppose the closing trading price for platinum on the world markets today was \$550 per ounce. (Does this pass the clairvoyant test?)



Toward the end of the day, you put in an order to your broker to

purchase one, two, or three contracts to sell 100 ounces of platinum one year from now, depending on how many contracts were available. Given the low volume in platinum contracts recently and given how late you called in, you figure there is about a .3 chance you got one contract, a .6 chance you got two, and a .1 chance that you got three.

You decide to seek out further information on the future of platinum prices. A very nervous metals broker gave you the following distribution on the closing trading price one year from now.

- a. Draw the influence diagram for this problem. Is there any information not reflected in the influence diagram, and, if so, how does the influence diagram relate to it?
 - b. Discretize the distribution above into a chance node whose branches have probabilities of .25, .50, and .25. What is the expected price of platinum one year from now?
 - c. Put together the complete probability tree describing your profit if you hold on to however many contracts you get until fulfillment. Assume you can buy the platinum you need to fulfill the contract(s) just before you need it. What is your expected profit from holding the contract(s)?
 - d. Discretize the distribution again, but this time into a two-branch node with probabilities of .5 and .5; then incorporate this node into the complete probability tree. Now, perform probability sensitivity analysis on the probabilities for platinum price by systematically changing the pair of probabilities—for instance, (1, 0), (.9, .1), (.8, .2), and so on. At what probability of the high level of platinum price does your expected profit from holding the contracts become negative?
 - e. Is this insight reflected in the influence diagram? Why or why not?
- 2.15 Show that balancing areas in the discretization process is equivalent to choosing the expected value given that you are in the range. The expected value, given you are in the range $a \leq x \leq b$, is:

$$\frac{\int_a^b xf(x)dx}{\int_a^b f(x)dx}$$

where f is the probability density, and the cumulative probability density is

$$P(x)_{\leq} = \int_{-\infty}^x x'f(x')dx'$$

(Hint: you will probably need to do an integration by parts.)

3

Decisions Under Uncertainty

What Is a Good Decision?

One common way people distinguish good decisions from bad ones is by looking at the results of the decision. Most people, however, realize this criterion is not very satisfactory. For one thing, while the results may not be apparent until much later, we would like to immediately characterize the decision as good or bad. In addition, all we see are the results of the chosen alternative—there is usually no way to see the results that would have occurred if a rejected alternative had been chosen and, thus, no way to see whether we have really chosen the better alternative. Most troublesome, however, is evaluating the situation when someone “lucks out” with what seems to us undeservedly good results. Should you call the decision “good” in this case and reward the decision-maker accordingly?


One of the fundamental benefits of decision analysis is that it can distinguish good decisions from bad ones. Furthermore, it provides a criterion for establishing whether a decision is good or bad.

Decision analysis starts by defining exactly what a decision is—the commitment of resources that is revocable only at some cost. If we do not commit resources, it is almost meaningless to say we have made a decision. The alternative “Do nothing” is a commitment of resources in the sense that an opportunity has been rejected.

Decision analysis, then, clearly lays out the four elements of rational decision-making (Figure 3-1). The first element is information, or “What do I know about the world and the business or personal opportunity under consideration?” An important component of this knowledge is an assessment of uncertainty (or “What *don't* I know?”) The second element is alternatives, or “What courses of action are open to me?” The third element is values, or “What do I want?” Finally, there’s logic, or “How do I put knowledge, alternatives, and values together to arrive at a decision?”

Figure 3-1

The Elements of a Good Decision Analysis Process



Given these elements, we can now characterize a good decision as one that is logically consistent with the alternatives, information, and values brought to the decision. In the decision analysis process, we can explicitly record the principal inputs to the decision-making process and establish that a good decision was made *before* the results of the decision are known. This minimizes the problem of “Monday-morning quarterbacking,” which can have an insidious effect on the morale of a company. Furthermore, in decision analysis, you can disagree with the information or values the decision-maker brings to the problem to the extent that you would have chosen another alternative—and still say that the decision-maker made a good decision. In corporate practice, other sets of information (and values, if appropriate) can be substituted in the evaluation for the decision-maker’s. By doing this, we can test not only whether the same alternative would be chosen, but also what the loss in value would be from choosing the other alternative. These results could then be filed as sort of a minority report to the record of the decision.

Why is it important to be able to determine whether a decision is good or bad? For one thing, the decision-maker will sleep better realizing that, given the time and resources available, he or she has done the best job possible. In addition, in corporate decisions (such as those made by electric utility companies), it may later be necessary to justify to stockholders, regulators, or other stakeholders that a good decision was made, even though a bad outcome ensued. Finally, we like to identify people who make good decisions so that they can be rewarded and promoted. Judging by results rather than by good decisions discourages people from taking a course of action with any risk in it—unless the results will not be known until the distant future!

One frequent objection to this definition of a good decision is, “That all sounds fine, but, in reality, it’s results that count, and you can’t tell me that in the real world a decision that leads to bad results is a good decision!” For operational decisions, there is much truth in that objection. However, there is seldom significant uncertainty in operational decisions. Consequently, if we correctly process the information available to us and follow up diligently, our decisions will usually lead to good results. It is in decisions involving significant uncertainty (typically strategic decisions) that the distinction between good decisions and good outcomes becomes important.

A second objection is that uncertainty represents inadequate knowledge and that a decision made with inadequate knowledge cannot be a good one: garbage in, garbage out. While the quality of the available information is certainly an important consideration, waiting for more information is frequently not advisable. Specifically, while the option of waiting and gathering more information to reduce the uncertainty should be included among the alternatives considered, information-gathering requires money (incurring information-gathering costs) and time (risking competitors' actions, opportunity lost, revenue lost). Furthermore, information-gathering will usually be incomplete and resolve only some of the uncertainty. Finally, some uncertainties cannot be resolved, either in principle or in practice, until the actual results occur. Decision analysis addresses these concerns explicitly with the value of information and value of control concepts and calculations, which are described in this chapter.


Recasting the Problem As a Decision Problem _____

Decision and Value Nodes

The president of Positronics was by now becoming uncomfortable with his decision to bid \$500,000. Correction—He had not yet submitted the bid and committed himself, so he had not yet really made the decision! To ease his discomfort, he requested the analysis be redone for different bids. The facilitator suggested rephrasing the problem in terms of a decision tree, which would not only automatically handle different potential bids, but would also allow further analysis to clarify the decision. The facilitator first added the decision to the influence diagram, using a rectangle to denote it. Then she drew an arrow from the decision node to the profit node, since the level of bid affects the profit. Finally, she changed the double oval around the Profit node to a double octagon to denote the fact that profit is the criterion to be used when making the decision.

Figure 3-2

Influence Diagram of the Positronics Bid Decision



In Figure 3–2, we have introduced the remaining elements of influence diagrams. In Chapter 2, we saw the uncertainty node (represented by an oval), the deterministic node (represented by a double oval), and influences (represented by arrows.) Here, we introduce the decision and value nodes and extend the concept of influence.

A *rectangle* represents a decision. Inside the rectangle is a descriptor (or variable name) to identify the set of events (or the quantities) among which the decision-maker is choosing.

An *octagon* represents a value node, inside which is a descriptor of the value measure. The value measure is the quantity that the decision-maker uses in making decisions; as discussed later in this chapter, the decision-maker will choose the alternative that maximizes some function of the value measure. The value node is just a chance node drawn with a special symbol to represent its special role.

A *double octagon* represents a value node that has ceased to be an uncertainty because all the uncertainty has been expressed in the nodes that influence it. It is just a deterministic node with a special role with respect to decisions.

The *arrow* now has added meaning. The basic concept of an arrow is flow of knowledge: the person providing the probability distribution for the node at the head of the arrow has information about what happened at the node at the base of the arrow. In addition, when decision nodes are involved, this information flow implies a chronology. The following are the four ways in which an arrow can be used:

- *Uncertainty Node to Uncertainty Node*—The outcome of the uncertainty at the base of the arrow is provided when probabilities are assigned for the node at the head of the arrow. This means that the probabilities at the node at the head of the arrow are conditional on the resolution of the uncertainty at the base of the arrow. (Conditional probabilities are treated extensively in Chapter 4.) For the special case of a deterministic node at the head of the arrow, the influence means that the resolution of the uncertainty at the base of the arrow provides a value to be used in the deterministic node.
- *Decision Node to Uncertainty Node*—The alternative chosen at the base of the arrow is provided when probabilities are assigned for the node at the head of the arrow. This means that the probabilities at the node at the head of the arrow are conditional on the alternative chosen at the base of the arrow. This means that the decision is made *before* the resolution of the uncertainty.
- *Uncertainty Node to Decision Node*—The outcome of the uncertainty at the base of the arrow is known when the decision is made. This means the uncertainty is resolved (and the results learned) *before* the decision is made.
- *Decision Node to Decision Node*—The alternative chosen at the base of the arrow is known and remembered when the decision

is made at the head of the arrow. This implies that the decision at the base of the arrow is made *before* the decision at the head of the arrow.

“Information flow” means that information is available to the person providing the probabilities or making the decision. Thus, we could have a node “Sales, 2005” influencing “Sales, 2003.” This would not imply backward causality in time! Rather, it would mean that the person providing probabilities for Sales, 2003 would be given information on the Sales, 2005 outcome: “If you knew Sales, 2005 were high, what would be the probability distribution for Sales, 2003?”

Rules for Constructing Influence Diagrams

To construct an influence diagram, follow the procedure outlined in Chapter 2. Two elements must be added to this procedure. First, decision nodes (drawn as rectangles) are to be identified in addition to the uncertainties as we work back from the one key uncertainty. Second, one uncertainty should be enclosed in an octagon to indicate its role as the value measure for decision-making.

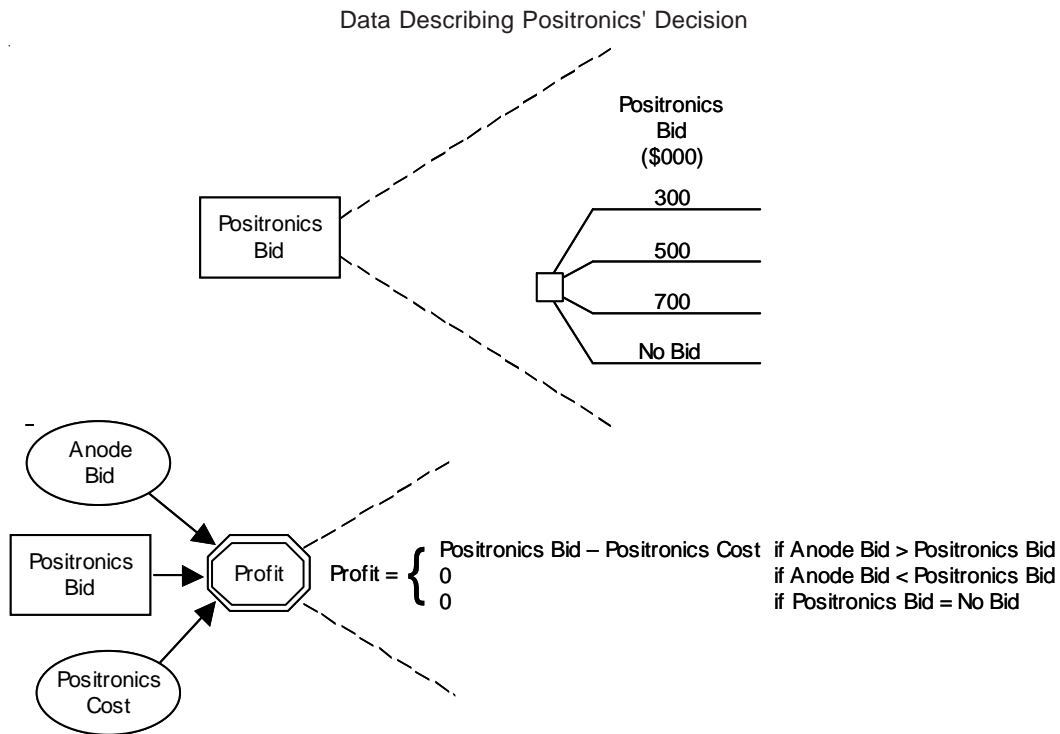
Four rules must be obeyed if you wish to construct meaningful influence diagrams.

1. *No Loops*—If you follow the arrows from node to node, there must be no path that leads you back to where you started. This is most readily understood in terms of the information flow indicated by the arrows.
2. *Single Decision-Maker*—There should be just one value measure (octagon) and all decisions should be made to maximize the same function (expected value or certain equivalent) of this value.
3. *No Forgetting Previous Decisions*—Each decision should be connected to every other decision by an arrow. This will establish the order in which the decisions are made and indicate that the decision-maker remembers all his previous choices.
4. *No Forgetting Previously Known Information*—If there is an arrow from an uncertainty node to a decision node, there should be an arrow from that uncertainty node to all subsequent decision nodes.

Constructing the Decision Tree

After some discussion, the Positronics staff identified four different alternatives for analysis—three levels of bid (\$300,000, \$500,000, and \$700,000) and the alternative not to bid. The square at the branching point in the tree in Figure 3–3 indicates a decision node, and each branch is an alternative. There are no probabilities associated with the branches because the node represents a decision, not an uncertainty. The value node now needs some modification because it is influenced by the decision node.

Figure 3-3



To join the pieces in a tree, we must be a little more careful than we were in the approach described in Chapter 2. The following procedure will transform an influence diagram into a decision tree.

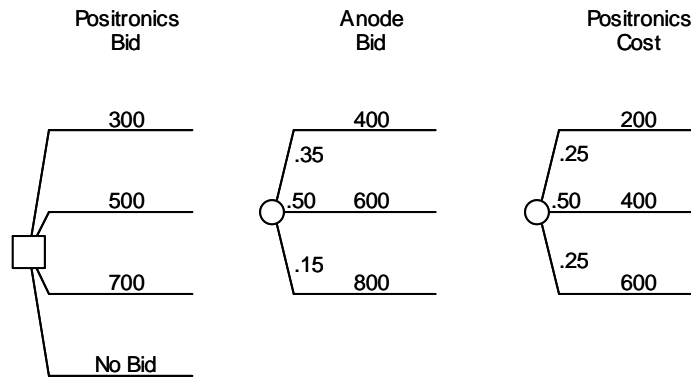
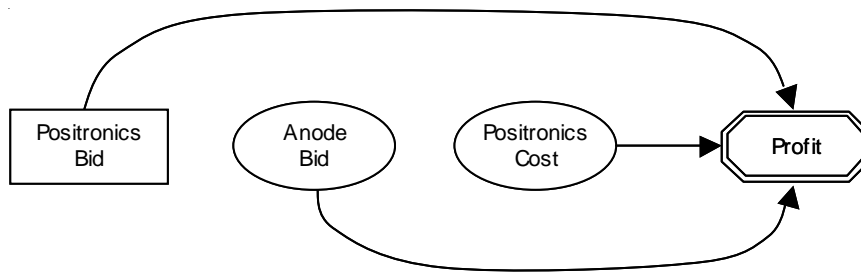
1. Arrange the decision nodes such that all arrows with a decision node at their base or head point to the right-hand side of the page. Arrows pointing to or emanating from decision nodes imply a chronology that must be followed in decision trees. By convention, the chronology of decisions in a decision tree flows from left to right, and therefore, these arrows must point from left to right.
2. Arrange the uncertainty nodes so that no uncertainty node is to the left of a decision node unless there is an arrow from that node to the decision node. In a decision tree, the outcome of a node to the left of a decision node is known to the decision-maker when he or she makes the decision; in an influence diagram, this means there is an arrow from the uncertainty node to the decision node.
3. Arrange, insofar as possible, the uncertainty nodes so that all arrows point to the right. This will cause conditional probabilities

- to be displayed simply on the tree.
4. Make the deterministic value node a tree endpoint node. This will usually involve calculations to find the value associated with each combination of events at the nodes that influence the value node.
 5. Now convert the diagram to a decision tree. Draw the content of the leftmost node in tree form. To obtain the complete tree, draw the content of the next node at the end of each branch of the preceding node, and repeat this process until all the nodes are drawn. To obtain the schematic tree, draw the contents of the node in order, with the understanding that each branch of a node leads to the node to its right.

Figure 3-4 is the rearranged influence diagram for Positronics and the schematic tree. Some common sense modification of this last step is needed in asymmetric situations. For instance, the No Bid branch of the decision need not be joined to the Anode Bid node or the Positronics Cost node, since both these nodes are irrelevant if no bid is made.

Figure 3-4

Translating the Influence Diagram into a Schematic Decision Tree



Decision or Uncertainty?

Graphic conventions are important in forming both influence diagrams and decision trees. As we have noted, a decision node is represented graphically by a square or rectangle, while chance nodes are represented by circles or ovals. This graphic convention helps to avoid confusion about what is a decision variable (completely under the decision-maker's control) and what is an uncertainty (or chance) variable (completely out of the decision-maker's control).

The distinction between decision and chance variables is occasionally not as clear cut as we might think. Sometimes the facilitator can choose which variables are decision variables and which are chance variables. For instance, the price of a product and the volume of sales cannot normally *both* be decision variables. You can set the price and then let the market determine the quantity sold or you can set the volume sold and then let the market determine the effective price realized. If nothing in the structure of the problem dictates which is the decision variable, the facilitator must decide. In making the choice, the facilitator should consider the impact of the choice on modeling the problem and on assessing probabilities. What is the easy or natural way to model? How do people think about the problem? How will it be easiest to assess probabilities?

There are several important differences in handling decision nodes and probability nodes. A decision node does not have any probabilities associated with the branches. Consequently, for the rollback procedure for tree evaluation introduced in Chapter 2, instead of replacing the node with its expected value (as we do with a chance node), we replace a decision node with the value of the branch that has the maximum value to the decision-maker. In other words, at a decision node we choose the alternative (branch) that gives us the most of what we want.

The order of nodes in the probability tree was fairly arbitrary, dictated mainly by the requirements of how people think best about the uncertainties. In a decision tree, however, the ordering of nodes implies a sequence of events. Time proceeds from left to right in the tree. If two decisions occur in a tree, the decision to the right is made after the one to the left; and the decision-maker remembers the choice made in the first decision when making the second decision. In addition, the uncertainty of any chance nodes that occur to the left of a decision node is resolved before the decision is made. That is, the decision-maker knows what actually happened before having to make the decision. However, node ordering is important only relative to the decision nodes. The ordering of the chance nodes relative to one another is still arbitrary. Thus, in Figure 3-4, Positronics makes its bid decision before it knows Anode's bid and before it knows what its own costs will be.

The outcomes or branches of a decision node should be a list of significantly different alternatives. Unlike the outcomes at a chance node, the list of alternatives at a decision node need not be mutually exclusive and

collectively exhaustive (mutually exclusive: you cannot choose more than one branch; collectively exhaustive: every possible alternative is represented by a branch). However, if the alternatives are not mutually exclusive, confusion may arise. It is very important to have a list of significantly different alternatives. (A collectively exhaustive list of alternatives is impossible for all but the simplest problems.)

Finding even a few truly different alternatives requires creative thinking before and during the decision analysis process. All too often, the decision process is hobbled by a lack of genuinely innovative possibilities. In the Positronics bidding case, we are looking only at different bid levels. Other alternatives might involve forming a joint venture to bid with another company more experienced in building this kind of instrument, performing some cost studies before bidding (we consider this later in the chapter), or exploring the possibility that MegaCorp would accept a bid of some fixed markup on cost (with appropriate safeguards against irresponsibility in cost overruns).

Building the Tree

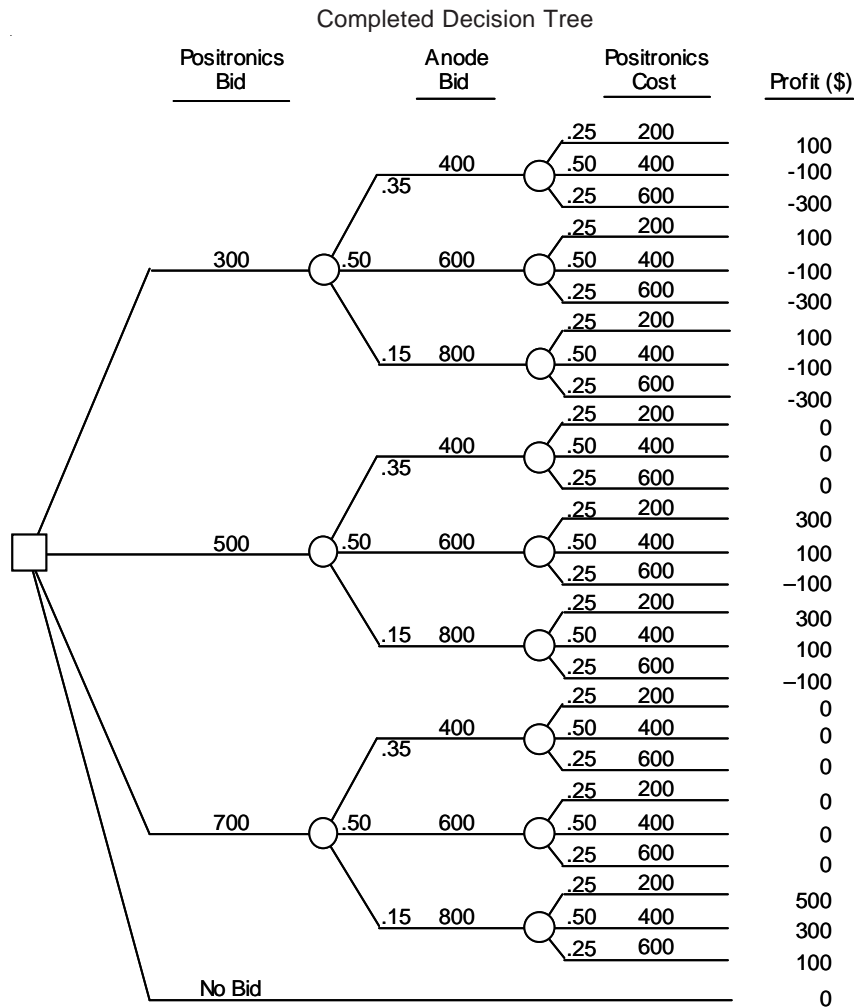
We have already done most of the work required to build the tree. The schematic tree has been created (Figure 3-4) and the value to be placed at the end of each of the branches has been determined. The complete tree is shown in Figure 3-5.

Each bid decision branch has its own probability distribution on Profit. Using the method developed in Chapter 2, we can construct the cumulative probability distributions for our decision tree (Figure 3-6). The staircase appearance is a result of the discrete approximation used in constructing the tree. If more branches had been used, the curves would be smoother. The data in Figure 3-6 show the “risks and rewards” associated with each of the alternatives. The decision-maker now has the data needed to make the decision.

Which alternatives are risky and which are conservative? Often, the high bids (\$700,000 in this case) are considered safer or more conservative because there is no possibility of losing money. However, we could also view the high bid as being the less conservative course of action: we are seeking to make a great deal of money, but since we probably will not make any money at all (because we probably will lose the bid), we are undertaking a great risk.

Notice the ambiguities in the preceding paragraph. “Risk” has so many ill-defined meanings that it is often better to avoid the word entirely. We will, however, use it in several well-defined contexts such as “risk attitude” and “risk penalty,” as discussed in Chapter 5. Another problem word is “conservative,” the meaning of which depends on our viewpoint and on the relevant values. When applied to estimates, conservative means a number the estimator thinks will result in a safe or conservative decision—according to his or her viewpoint. There is no point in giving “conservative” estimates for chance variables; they give an inaccurate estimate of the total uncertainty

Figure 3-5



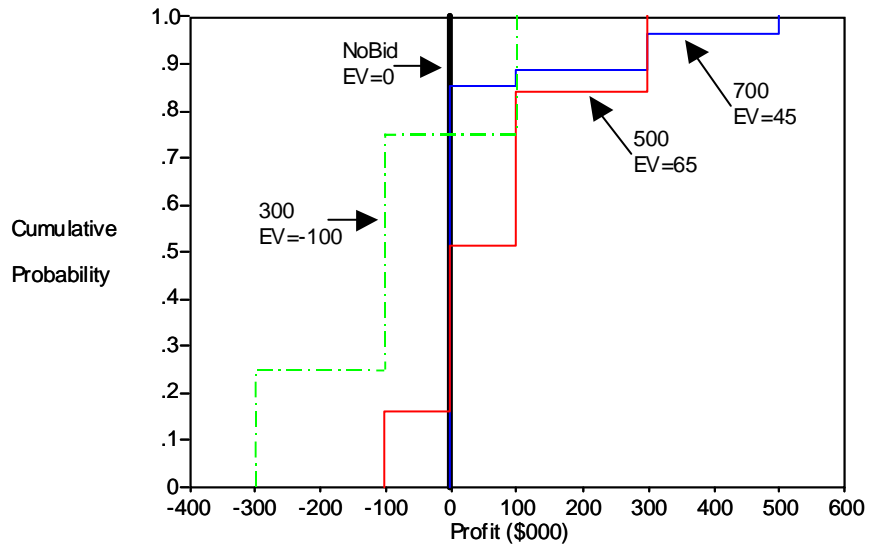
and often distort the results in unanticipated ways. It is much better to put in the best information through the probability distribution and to let the decision-maker make the decision.

Decision Criterion

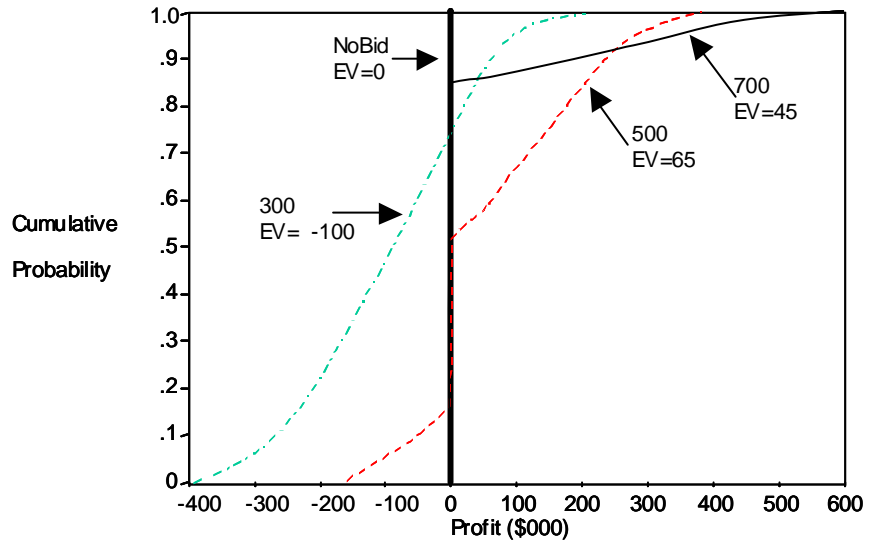
How do we make decisions under uncertainty? For a simple tree with a single decision node (such as Figure 3-5), the decision-maker can stare at the probability distribution on value associated with each of the alternatives and make the decision. Figure 3-6 shows the results of combining all the information logically and consistently. Because only three branches were used for Positronics' Cost, the graph is difficult to read. If the underlying data

Figure 3-6

Probability Distribution on Profits for Alternative Bid Decisions



Probability Distribution on Profits for Alternative Bid Decisions With Many Branches for Positronics Cost Node Discretization



from Figure 2-14 were discretized into very many branches, the second graph in Figure 3-6 would result.

However, for more complicated trees with multiple decision nodes, this procedure is impractical. We need a decision criterion that indicates the best choices and that helps the decision-maker choose consistently. The decision-maker's true decision criterion is embedded in the values he or she brings to the decision. Thus, the decision criterion is highly personal. But there are common characteristics that can help the decision-maker (and the decision facilitator) model these values and obtain a systematic decision criterion.

In most business and personal financial decisions, money is the primary value and maximizing wealth is the principal decision criterion. To express the value function completely in monetary terms, however, we must be able to express nonmonetary items in monetary terms. To do this, we must deal with the following three subjects:

- The value of nonmonetary, intangible goods
- The value of future money
- The trade-off between certainty and uncertainty.

The Value of Nonmonetary, Intangible Goods

When the principal value function of a tree is given in monetary terms, all values to be included in the tree must also be in monetary terms, including such intangible items as goodwill, worker safety, and laying off workers. Usually, obtaining estimates of monetary equivalents for these so-called intangibles is not difficult, though the decision-maker might not want the value needlessly publicized.

Many items that are referred to as intangibles are not really intangible, but rather are difficult to put a value on. For instance, increased goodwill may result in direct economic effects from increased sales or decreased transaction costs. This is not an intangible but rather an effect that is difficult to estimate. After this has been separated out, you can address the intangible part of the value that stockholders and managers attach to goodwill.

Establishing values for intangibles is not an attempt to disguise moral or social irresponsibility. Moral and social responsibility should be built into the initial choice of acceptable alternatives. Nor are intangible value trade-offs an attempt to put a price on intangibles such as happiness. The goal is to appropriately represent the trade-offs people make in practice in order to include them in economic decision-making.

In most business decisions, intangibles with even the most generous trade-offs have little effect on the value of the decision alternatives. In the rare case in which intangibles are important in the problem, the decision-maker must examine his or her values with much more care than outlined here.

In societal decisions, determining the value function is much more difficult. Societal decisions, such as road design or worker safety legislation, typically involve a trade-off between resources measured in monetary terms and in terms of human life and welfare. While there are a number of ways

to deal with these decisions rationally, they are beyond usual business decision analysis practice.

The Value of Future Money

In the previous section, we expressed the value function in monetary terms. This reduces multiple time-varying uncertain results to a single time-varying uncertain result. This result is typically expressed as a stream of money stretching many years into the future. The value of money received in the future is not the same as that of money received today. Some people have immediate, productive uses for their capital while others do not; and if they do not, they can always lend (invest) their money to someone else for a fee (interest). Banks and other financial institutions establish a marketplace between present and future funds. The relationship between present and future funds is expressed in terms of a discount rate or an interest rate. The present value of cash, c_i , received years i from now is:

$$\frac{c_i}{(1+d)^i} \quad (3-1)$$

where d is the discount rate the decision-maker wishes to use when trading off between present and future. (We discuss in Chapter 5 values that might be appropriate for the discount rate.) Using the present value formula, we can reduce the time series of cash extending m years in the future to a net present value (NPV):

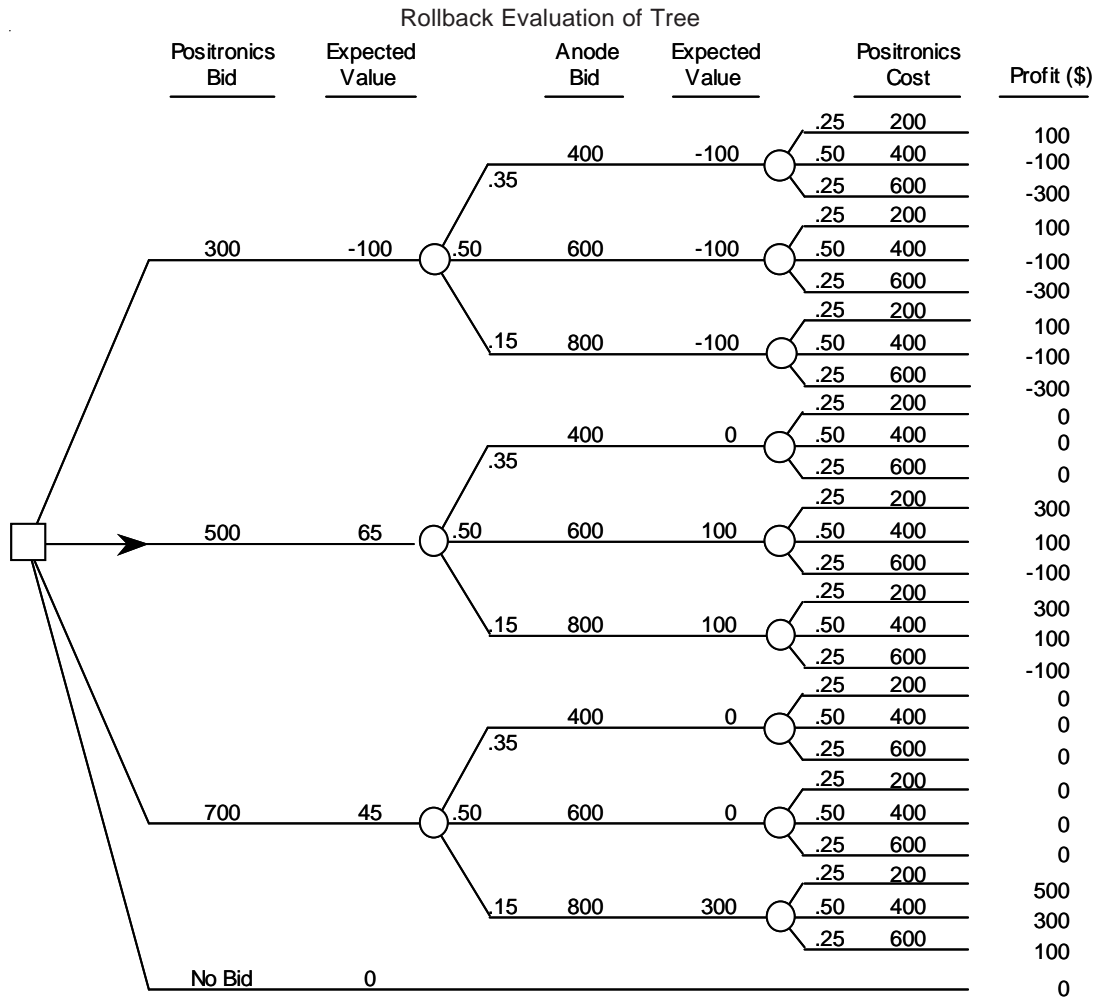
$$NPV = \sum_{i=1}^m \frac{c_i}{(1+d)^i} \quad (3-2)$$

thus eliminating time from the decision criterion.

Our decision criterion is to maximize the NPV of future cash flows. We could also maximize the NPV of other monetary measures, but cash flow is ordinarily a very good measure of the net change in wealth of a company or an individual. Other monetary measures (net income, earnings per share, dividends, payback period, return on investment) tend to focus only on parts of the problem and can lead to bad decisions if used as criteria without sufficient care.

The internal rate of return (IRR) criterion, although similar to the NPV criterion, is difficult to apply meaningfully in an uncertain situation. This criterion can be used to choose the best alternative given certainty (i.e., the same choice as from an NPV), but it does not show how much more valuable this preferred alternative is than the others—and in choices under uncertainty, we *must* be able to balance how much we might win against how much we might lose.

Figure 3-7

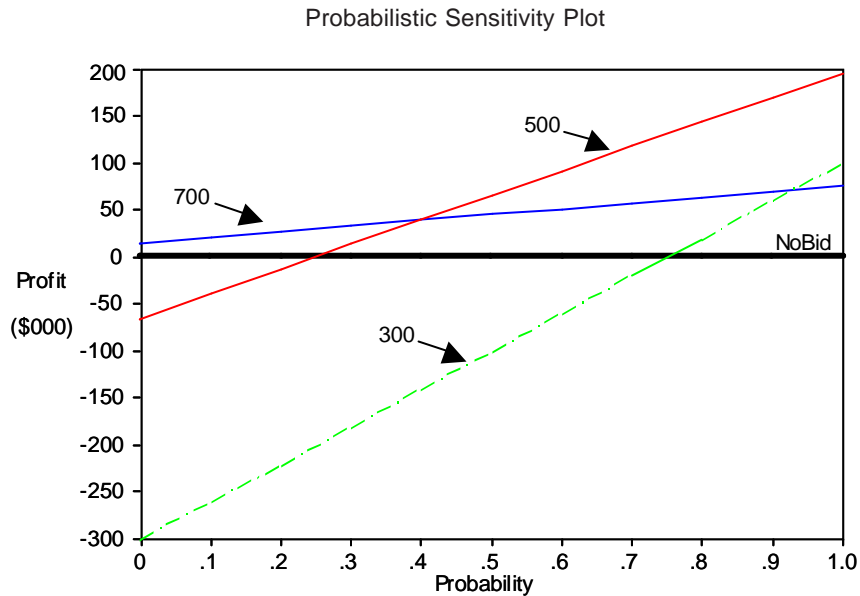


The Trade-off Between Certainty and Uncertainty

In the preceding two steps, we have reduced multiple time-varying uncertain outcomes to an uncertain NPV. As we saw in Figure 3-6, uncertainty means that there are many possible NPVs with associated probabilities. How can we now choose between alternatives with different probability distributions on NPV?

In Chapter 2, we discussed the expected value: the sum of probabilities multiplied by their respective values. The expected value represents the average return we would expect to receive if we were engaging in many identical but uncorrelated ventures. If the stakes are not very high, it makes sense to use the expected value as the value of a probability distribution. As the stakes go up, most people begin to show risk-averse behavior: the value

Figure 3–8



they assign to a probability distribution is less than its expected value. (This behavior is discussed in Chapter 5.) Most business decisions, however, are not really that large relative to the total value of the company, and maximizing the expected value of the NPV of cash flow is a good approximation of the actual decision criterion.


The president of Positronics had become rather befuddled during the discussion of values and trade-offs. It had been a long time since he had been forced to examine just what he was trying to accomplish in running Positronics. However, it turned out that the simple value function used in the tree was substantially correct. Winning the bid would not greatly alter the probability of obtaining future work with MegaCorp, so there was no follow-on value to be added. Losing the bid would not entail any layoffs, because managing the normal labor turnover would be sufficient. The job was routine enough that no intangibles, such as added reputation, were attached to winning the bid. The job would be done within a year, so discounting was not a problem. Finally, at least for the time being, the president was willing to use expected value of profit as a decision criterion. After all, the possible profits or losses from the job were not that important to the company's long-term success.

Analyzing the Tree

The complete tree is shown in Figure 3–7. The rollback procedure has been used to present the expected value for each of the decision alternatives. What should happen at the decision node? If expected value is the chosen decision criterion, then the decision rule is to choose the alternative with

Figure 3-9

Sensitivity of Preferred Decision to Changes in the Probability Assessments for Positronics Cost and Anode Bid



the maximum expected value. For this decision criterion, the best alternative is to bid \$500,000. Note the arrow to indicate the chosen alternative.

We have said that there is no such thing as a “correct” probability. The only worry is whether the probability assessment adequately reflects the decision-maker’s state of knowledge. Nevertheless, there is a natural desire to know how important a probability assessment is in the choice of the alternatives in the decision problem. To show this, you can perform probabilistic sensitivity analysis (also known as stochastic sensitivity analysis). In Figure 3-8, the horizontal axis is the probability assigned to the first (top) branch of node Positronics Cost; the rest of the probability is assigned to the bottom branch, and the middle branch is given zero probability. This two-branch approximation is used for ease of interpretation. There are other approximations which are sometimes appropriate. For instance, the probability of one branch can be varied while keeping the ratio of probabilities of the other branches fixed, as in the probability for Anode Bid in the two-node sensitivity example below.

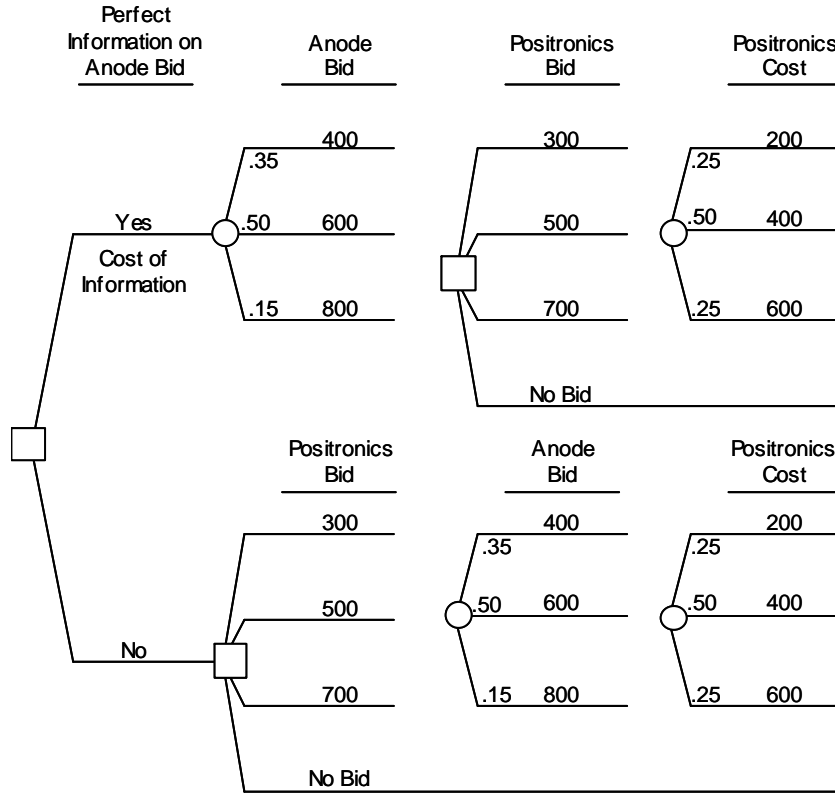
The plot in Figure 3-8 shows that when the probability of low cost (\$200,000) is high, the \$500,000 bid is preferred. Below about 40 percent probability of low cost, the \$700,000 bid (alternative 3) becomes preferred because the losses associated with the \$500,000 bid become much more likely.

Probabilistic sensitivity to two nodes can be evaluated by systematically changing the probabilities at two nodes to determine the probability at which the decision choice switches.

Positronics felt that the relative likelihoods of the \$800,000 and \$600,000 values for the Anode bid were reasonable. The staff was much less certain about the probability of the low bid (\$400,000). Some thought a bid this low was not typical of Anode; others thought Anode had enough experience in this area that it could hold costs very low and hence afford to bid low. To see how changes in

Figure 3-10

Decision Tree to Determine the Value of Perfect Information on Anode's Bid



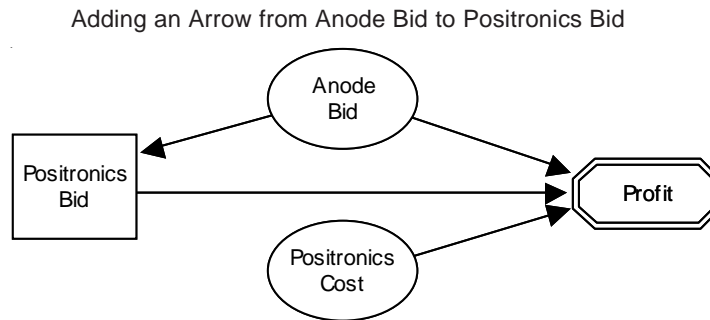
the relative likelihood of the \$800,000/\$600,000 bid level and the \$400,000 bid level affected their decision, they decided to do a probability sensitivity analysis with the constraints shown in equations 3-3 and 3-4.

We use the notation $p(X|S)$ to indicate the probability that X occurs, given the state of knowledge S . However, in performing probabilistic sensitivity analysis, we are continuously varying probabilities—implying that we are similarly varying the underlying state of knowledge S on which a probability depends. Because the state of knowledge is varying, we have replaced the S notation here by S' to indicate that a different state of knowledge is implied by each different probability.

$$\frac{p(\text{AnodeBid} = 600|S')}{p(\text{AnodeBid} = 800|S')} = \frac{.5}{.15} \tag{3-3}$$

$$p(\text{PositronicsCost} = 400|S') = 0 \tag{3-4}$$

Figure 3–11



In the display of sensitivity results (Figure 3–9), the X marks a point corresponding to the current state of information. [Since Positronics Cost has been changed from a three-branch node to a two-branch node, we have set $p(\text{Positronics Cost} = 200 | S) = .5$ to get the same expected value as the original tree.]

The graph shows that when $p(\text{Positronics Cost} = 200 | S)$ is less than .4, bidding \$700,000 is preferred. When $p(\text{Positronics Cost} = 200 | S)$ and $p(\text{Anode Bid} = 400 | S)$ are both high, bidding \$300,000 is preferred.

Probabilistic sensitivities are often useful to settle arguments or to obtain consensus. If the same decision alternative is optimal for everyone's set of probabilities, then there is really no argument about what to do. If someone's probability indicates that he or she would choose a different alternative, then the plots show how much is "lost" by choosing the "wrong" alternative.

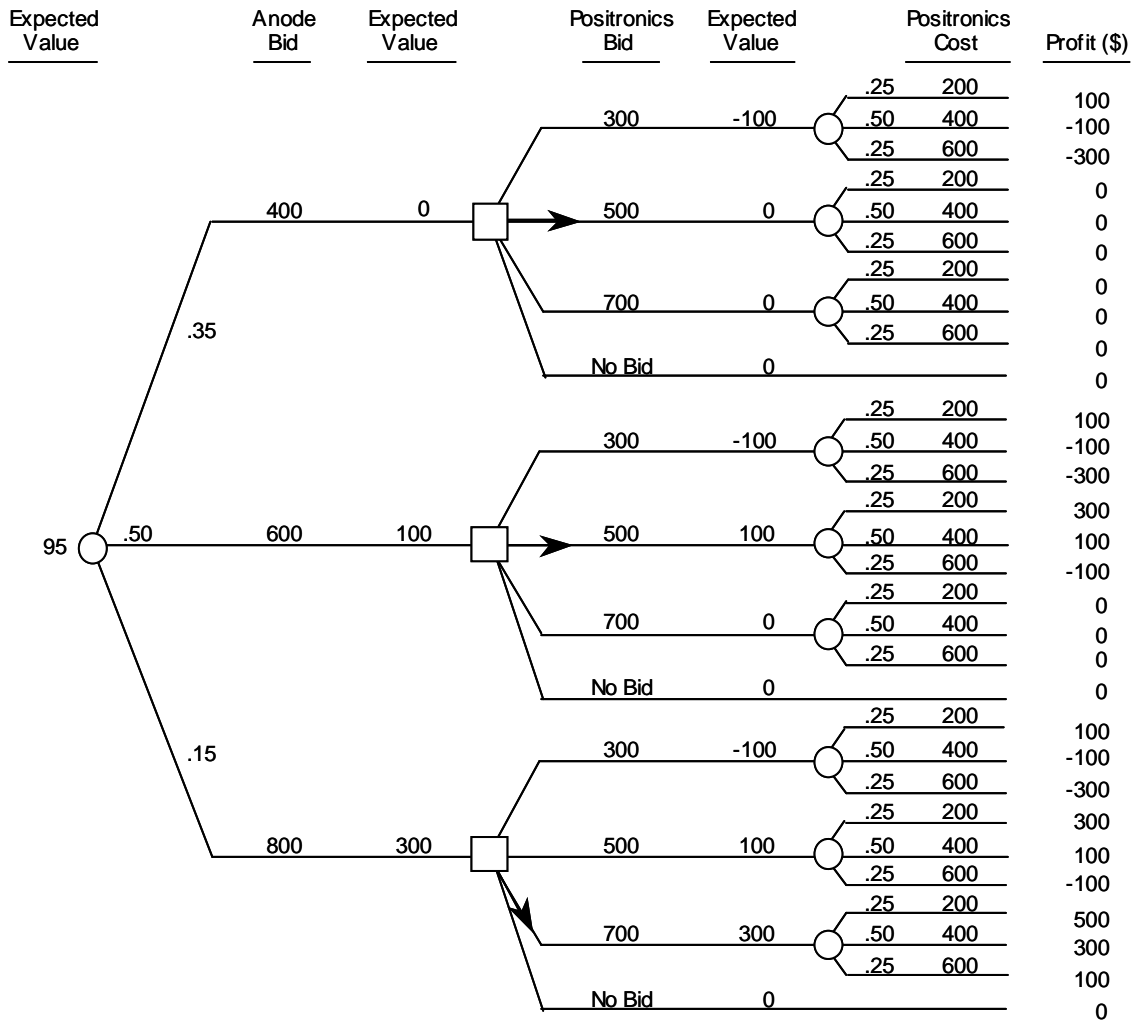
The Value of Perfect Information

The president of Positronics was satisfied by the results he had seen so far. He requested that the facilitators add more branches to the tree so he could zero in on the optimal bid. However, the facilitators suggested that they first examine other insights available from the tree: the value of information and the value of control.

One of the powerful features of decision analysis is the ability to show the value of resolving an uncertainty *before* making the decision. The simplest approach is to calculate the value of perfect information (also called the value of clairvoyance). Imagine that we had some means of obtaining information to completely resolve an uncertainty before we made our decision. (Perhaps a clairvoyant's services were available.) What is the maximum we should pay for this information? This value, which turns out to be easy to calculate, is the upper limit of the amount we should be willing to spend on any information-gathering effort, which generally yields imperfect rather than perfect information. Incidentally, we often find companies spending more on information-gathering than it is worth.

Figure 3-12

Subtree for "Yes" Branch for Perfect Information on Anode Bid

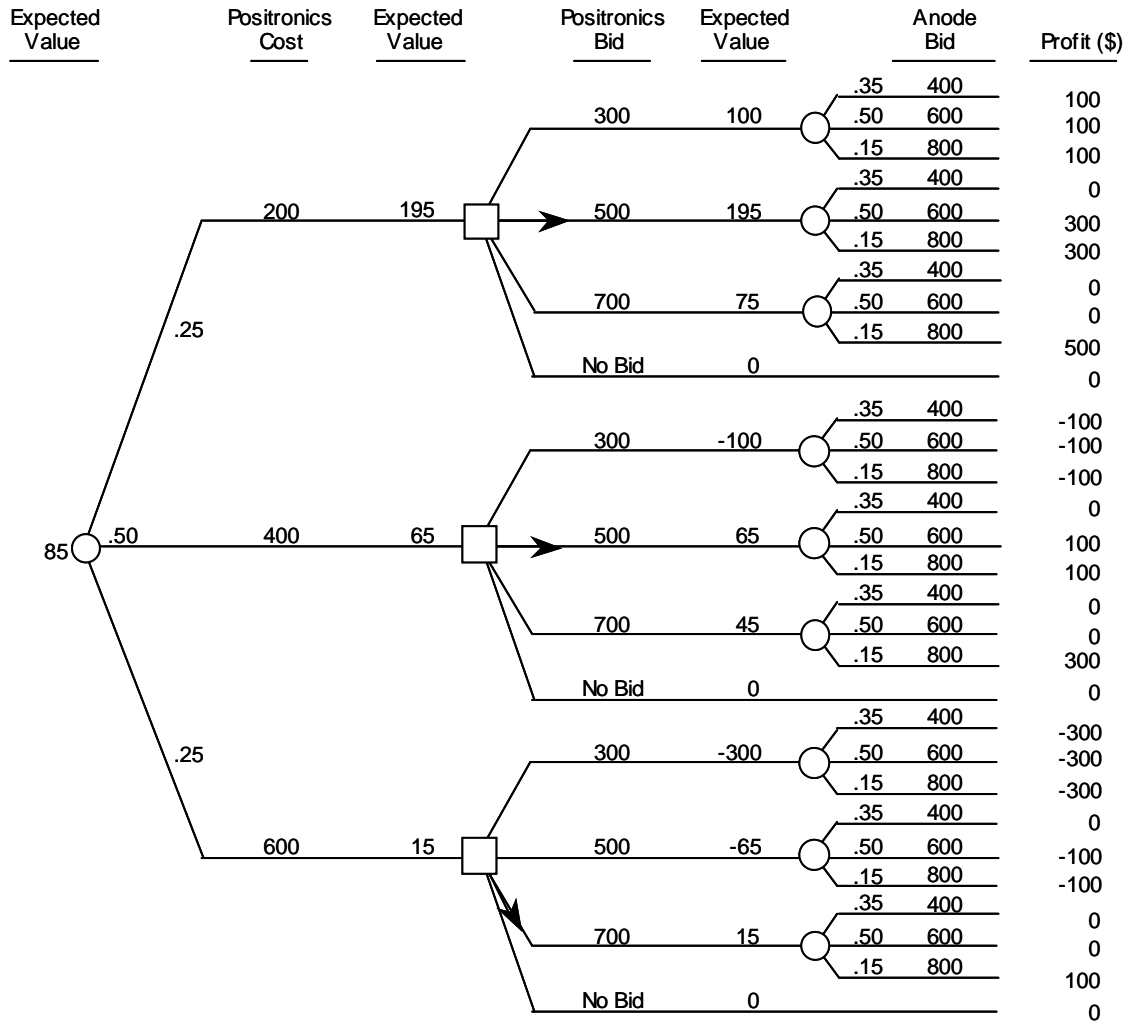


To calculate the value of perfect information on Anode's bid, Positronics introduced a new decision at the front end of the tree (Figure 3-10). The first node in this tree is the decision whether or not to obtain the perfect information. If Positronics does not obtain the information, it faces the tree following the "No" branch, which is identical to the tree evaluated in Figure 3-7. Thus, for the "No" branch, Positronics already knows it should choose to bid \$500,000, resulting in an expected value of \$65,000.

If the company decides to obtain the information (the "Yes" branch), it must pay the cost of the information. It then faces a tree similar to the

Figure 3-13

Subtree for "Yes" Branch for Perfect Information on Positronics Cost



original one, but with the first and second nodes reversed, because Anode's bid is known before the decision is made. A way to represent this subtree at the end of the "Yes" branch is to add an arrow from the Anode Bid node to the Positronics Bid node in the influence diagram (Figure 3-11). This indicates that the uncertainty on Anode's bid is resolved before the Positronics bid decision is made. The corresponding tree is shown in Figure 3-12.

Note that the probabilities for Anode Bid are the same as before—if we thought there is a 35 percent chance of Anode's bid being \$400,000, then

we should think there is a 35 percent chance that the perfect information-gathering activity (or the clairvoyant) will reveal that Anode's bid is \$400,000. The tree in Figure 3-10 could be evaluated for varying costs of information until the expected value of the "Yes" (obtain information) alternative is equal to the expected value of the no-information alternative. This would then reveal the value of obtaining the (perfect) information—the maximum cost Positronics would be willing to incur to obtain it.

However, there is a simpler approach. The expected value of \$95,000 for the tree in Figure 3-12 is the value of the venture with perfect information. We can then calculate the value of the perfect information as follows:

$$\begin{aligned}
 \text{Value of Information} &= \text{Value with Information} \\
 &\quad - \text{Value Without Information} \\
 &= \$95,000 - \$65,000 \\
 &= \$30,000 \qquad \qquad \qquad (3-5)
 \end{aligned}$$

Similar calculations show that value of the venture with perfect information on Positronics' cost is \$85,000 (Figure 3-13). The value of perfect information on Positronics' cost can be calculated to be \$85,000 - \$65,000 = \$20,000.

Note that the value of information on two or more variables is *not* the sum of the values of information of each variable separately—it can be smaller or greater. In this case, the sum of the values of information on both variables separately is \$30,000 + \$20,000 = \$50,000. This is smaller than \$51,300, the value of information on both variables simultaneously.

We can gain further insight into the value of information by examining the portion of the tree as it would appear for the "Yes" branch of Figure 3-10. This subtree is displayed in Figure 3-12. Note that the only change that Positronics would make would be to change its bid from \$500,000 to \$700,000 if it knew that Anode were going to bid \$800,000. Making this change in bid changes the expected value in that case from \$100,000 to \$300,000, for a net improvement of \$200,000. But there is only a .15 chance of finding that Anode will bid \$800,000. Therefore, the expected improvement in value is $.15 \times 200,000 = \$30,000$. This is the value of the information for Anode Bid—the same answer as obtained above.

Thus, information has value if knowing the results could lead us to change our decision. If none of the possible information-gathering results change the decision, there is no value to the information.

The value-of-information tree illustrates a number of powerful concepts. We believe that for many people, a change in perspective must occur before they understand the value of information calculation. ("Oh, is that what you mean?!") Understanding the concepts behind value of information is crucial to understanding and effectively using decision trees.

Positronics' president was impressed with the ideas implied in calculating the value of information. Since new information on Anode's bid was not obtainable by any ethical means, value of information was of only academic interest; most sources of information, such as information on prior bids, publicly

available information on Anode's backlog, etc., had already been used in the probability estimates. Some information on what Positronics' cost might be was obtainable for considerably less than the value of perfect information of \$20,000. This information would be obtained by using some off-the-shelf components to build a prototype instrument. The president encouraged production to go ahead with its prototype, provided the cost of the program was well under \$20,000.

The Value of Perfect Control

Decision trees produce one more piece of information: the value of control (also called the value of wizardry). If a clairvoyant (someone who can foresee the future) is valuable to us, a wizard (someone who can change the future) is even more valuable. The effect of perfect control is that we could pick which one of the future (or otherwise unknown) possibilities (branches) we wanted to happen. Thus, the value with perfect control for Positronics' cost can be obtained by inspecting the tree with the Positronics Cost node moved to the front of the tree (Figure 3-23).

Inspecting the top branch shows that if there were some way to force Positronics' cost to be \$200,000, it would yield an expected value of \$195,000. (A more straightforward but time-consuming way of obtaining this result is to change the probability that Positronics' cost is \$200,000 to 1.0 and then roll back the tree.) The calculation of the value of perfect control is similar to the calculation of the value of information.

$$\begin{aligned}
 \text{Value of Perfect Control} &= \text{Value with Perfect Control} \\
 &\quad - \text{Value Without Perfect Control} \\
 &= \$195,000 - \$65,000 \\
 &= \$130,000
 \end{aligned}
 \tag{3-6}$$

How can we interpret the value of perfect control? The probabilities for cost were obtained under the assumption that Positronics would use its normal cost control procedures. Control might, therefore, correspond to the alternative of instituting some extraordinary cost control procedures to make sure costs stay as low as possible. The value of perfect control gives an upper limit on the amount Positronics should be willing to spend on such cost control procedures.

Summary

Given our understanding of uncertainties and their different possible outcomes, it is possible to define a good decision as one that is logically consistent with the information, alternatives, and values brought to the decision.

We introduced a decision on the bid level into the Positronics example and showed the probability distributions on profits for each bid level. It is clear that whatever bid level is chosen, there is some possibility of a bad outcome; and yet it would seem wrong to characterize a bid of \$500,000 or \$700,000 as

a bad decision. The problem is choosing the better of the two levels.

We then discussed criteria for making a decision, including the value of nonmonetary, intangible goods; the value of future money; and the trade-off between certainty and uncertainty. (The trade-off between certainty and uncertainty is discussed at greater length in Chapter 5.)

By defining these valuations and trade-offs, we were able to calculate the preferred decision for the Positronics decision tree and to examine the sensitivity of the value of the preferred decision to changes in the probability assessment.

The analysis was expanded by the introduction of the value of perfect information and of perfect control. As illustrated by the modified trees used to calculate these values, information or control only has a value when it changes the preferred decision from what it otherwise would have been.

For Positronics, a small information gathering exercise on cost appeared justified. Much more valuable would be a cost control program.

Problems and Discussion Topics _____

- 3.1 It is very common for people to judge the quality of a decision by looking at its outcomes. Is this an unbiased point of view? How can you best deal with uncertainty to make a decision?
- 3.2 A good decision is defined as one that is logically consistent with the information, alternatives, and values brought to the decision. Give an example from your own experience (or perhaps from history) of a good decision/bad outcome and bad decision/good outcome. How could the bad decision have been improved? Remember that a good decision only has to be consistent with the information available at the time of the decision, though value of information should reveal the importance of the things you do not know.

Possible examples from history are: good decision/bad outcome—sailing on the Titanic (Who would have thought that the safest ship afloat would sink on its maiden voyage?); bad decision/good outcome—assigning Michelangelo to paint the Sistine Chapel (Who would have thought that a sculptor could paint on such a vast scale?).
- 3.3 Describe the information, alternatives, values, and logic you used in deciding where to eat dinner the last time you went out. Afterwards, were you satisfied with the decision? Was the outcome good or bad?
- 3.4 How is getting the expected value different for decision trees than for probability trees?
- 3.5 Net present value (NPV) and expected value are abstract concepts in that people usually will not get the NPV of a cash flow (unless they are buying or selling an annuity) or the expected value of a lottery. NPV is usually understood as the result of a trade-off between future value and present value; expected value is regarded as a way to deal with or

evaluate an uncertainty that is unresolved.

Write a short definition of each term. How are these concepts implemented in decision tree calculations?

- 3.6 List two radically different alternatives to getting your homework done (other than doing it yourself). Are these alternatives worth pursuing further? Why or why not? (Relate them to your values, the probability of getting caught, and the consequences of getting caught; to the effect when you take the midterm exam; etc.)
- 3.7 Today, liability for damages is usually settled with a cash sum, whereas in the past, the rule was often blood for blood. (Consider Oedipus, for instance.) One rationale is not that the victim is being bought off (i.e., that money is equivalent to his or her pain), but rather that given that the incident has already occurred, the victim can use the money for some purpose whose value to him or her can help compensate for any losses. In addition, there are sometimes punitive damages to punish the malefactor.
- Do you agree with this concept of trading money for pain and suffering? Do you find it more, less, or just as moral as the earlier method of compensation? Another alternative is no compensation. Suppose you are a victim. Do these two types of compensation mean anything different for you? Why or why not?
- 3.8 Anyone lending money at interest is establishing a marketplace between present and future money. The relationship between present and future is described by the interest rate. The marketplace works because some people have capital that would otherwise be unproductive if they did not lend it to people with a better use for it. List at least three reasons why people (or companies) might have different time values of money (discount rates).
- 3.9 As with discount rates, people can have different attitudes to trading off certainty and uncertainty. List at least three reasons why people might have different attitudes toward risk. For instance, some people have dependents and some don't. Why might a person's risk attitude change over time? Can education play a role in this?
- 3.10 List some of the principal values that must be considered in making decisions in a profit-making, publicly held hospital. Suggest the structure of a value model that establishes trade-offs between the values (at a deterministic level). State the limitations of the model. Are there ethical questions that cannot be simply resolved?
- 3.11 Explain how (if at all) the purpose of an influence diagram is different from that of a decision tree. Would it ever make sense to draw a tree and *then* the influence diagram? If so, when?

- 3.12 Can an equivalent decision tree *always* be drawn from an influence diagram and vice versa, or are there structural and/or informational differences that would make it impossible to do so without changing the problem or adding information? If there are differences, illustrate them with examples.
- 3.13 There are two main reasons for doing probabilistic sensitivity analysis in a decision-making process. First, probability assignments could change because of new information or could be different if there is more than one decision-maker. In each case, we need to know how the decision will change given changes in probabilities (information).
Second, probability sensitivity can distinguish the major uncertainties that are most influential to the decision from those that are less influential. This may increase understanding of the decision, provide directions for further information-gathering activity, etc.
Interpret Figure 3–8, the probabilistic sensitivity plot. If there are crossover points between alternatives, how do you interpret the corresponding crossover probabilities?
- 3.14 Draw an influence diagram that corresponds to the tree in Figure 3–10. (Hint: Introduce a deterministic node called Information Learned About Anode Bid.)
- 3.15 Suppose you are going skiing this weekend for the first time. However, you are worried about the possibility of breaking your leg during your first time out. Your alternative is to go to the beach. After thinking about the possibilities, you have decided that you value a weekend of skiing with no mishaps at \$1,000 and that you value a broken leg at $-\$5,000$. Going to the beach instead is worth \$500 to you. Finally, after talking to other people, you peg the chance of breaking your leg at 1 in 100.
- Draw the influence diagram for this problem.
 - Structure the decision tree for this problem. What is the preferred decision and expected value?
 - Calculate the values of information and control on breaking your leg. To do this, you will have to symmetrize the tree so that the uncertainty on breaking your leg follows both the ski and beach alternatives.
 - What impact might the values of information and control have on your choice of weekend?
 - For what probabilities of breaking your leg do you prefer going to the beach?
- 3.16 You are considering four different restaurants for dinner with a group of friends tonight. However, before deciding whether or not to eat at a restaurant, you'll want to look at the menu and see what they have. At that point, you can either stay and eat or move on to the next

restaurant. You value the cost of gathering everybody up and driving to the next restaurant at \$20.

When you look at a menu, you have a scale in mind for rating it. You'll assign a score of A, B, or C and value the scores for their contributions to the evening's enjoyment as follows.

Score	Value
A	\$100
B	\$50
C	\$0

You judge your four possible choices as equally likely to get any of the three scores.

- Draw the influence diagram and decision tree for this problem.
 - What is the best strategy for picking a restaurant, and what is the overall expected value?
 - Suppose you can buy a restaurant guide for \$20 that will alert you to any C restaurants. Modify your influence diagram and tree to reflect this choice. Should you buy it?
- 3.17 Raquel Ratchet is working on a Volkswagen Beetle to be sold at a collectors' car sale on Saturday. If she finishes it in time, she'll be able to sell it for \$1,000, at a cost of \$100 in parts. (She got the car from a junkyard.) She thinks there's a 60 percent chance of being able to do this.

She also has the option of installing a turbocharger and intercooler in the car, which would quadruple the horsepower and enable her to sell it for \$10,000. This would cost an additional \$1,000 in parts. She thinks there's only a 20 percent chance of being able to finish this amount of work by Saturday. If she misses the collectors' sale on Saturday, Raquel figures she can only sell the car for \$400 (\$1,400 with turbocharger and intercooler).

Suddenly, the Wizard appears in the form of a good salesman and tells her that he can increase the selling price to \$20,000 if she installs the turbocharger and intercooler. What is the maximum Raquel should be willing to pay the Wizard to do this? What is the maximum amount she should be willing to pay him on a contingency basis (if he makes the sale)?

- Draw the influence diagram and tree for this problem.
 - How much should Raquel pay in advance (nonrefundable) if there's only a .8 probability of his making the sale?
- 3.18 Samuel Steelskull is thinking about whether or not to wear a helmet while commuting to work on his bicycle. He figures his chances of dying in an accident during the coming year without the helmet are about 1/

3,000. The odds of dying go down to 1/5,000 with the helmet.

Sam figures it is worth about \$80 to him for his hair not to be messed up when he gets to work during the coming year, and he is pretty much indifferent between wearing and not wearing the helmet.

- a. Draw the influence diagram and tree for this problem.
 - b. What's the implicit value Sam is putting on his life (i.e., a value that might be used in very small probability situations)?
- 3.19 Your company has recently developed Chewsy, a new sugar-free gum that contains fluoride. Not only does it taste good, it's also good for your teeth. You are faced with the decision of whether or not to introduce Chewsy to the market.

The total sales of chewing gum are expected to be about \$200 million over the next 10 years.

Your marketing personnel feel that with their best efforts and with a front-end marketing expenditure of \$4 million, your company could capture from 2 percent to 10 percent of the chewing gum market with Chewsy. They have given you the following probabilities.

Market Share	Probabilities
High (10%)	.30
Medium (6%)	.50
Low (2%)	.20

Your financial advisors point out that the profit margin on Chewsy is quite uncertain because of unusual manufacturing requirements. They say that there is a 40 percent chance that the profit margin will be only 25 percent of sales revenue and a 60 percent chance it will be 50 percent of sales revenue. The manufacturing requirements will be known before the marketing decision needs to be made.

- a. Draw the influence diagram for this problem.
- b. Structure the decision tree for Chewsy. What is the preferred decision and expected value?
- c. Plot the probability distribution for profit for the Chewsy decision. What are the expected values for both alternatives?
- d. What is the value of information on market share? On profit margin? On both?
- e. What is the value of control on market share? On profit margin? On both? Are there any possible ways of achieving further control over either of these uncertainties?
- f. Does the preferred decision vary with the probabilities for market share or margin? What decisions are preferred for what ranges of probability?

- 3.20 The Southern Power Company is planning to submit a major rate increase request to the state Public Utilities Commission. The Commission has assured Southern Power that it would approve a request for a moderate rate increase. With this moderate rate increase, Southern would receive \$40 million in additional revenues during the next few years (relative to no rate increase).

However, the company is also considering a riskier course of action—requesting a high rate increase that would yield \$100 million in additional revenues, if approved. If the high rate increase is not approved, there is still some chance the Commission would grant Southern a low rate increase, which would mean \$30 million in additional revenues. Of course, the possibility exists that the Commission would simply refuse any rate increase whatsoever if Southern asks for the high increase.

The best information within the company indicates a 70 percent probability the Commission would disapprove the high rate increase request. Given that it does so, the chance it would then grant a low rate increase is believed to be 60 percent.

- a. Draw the influence diagram for Southern's problem.
 - b. Draw the decision tree for Southern's decision problem.
 - c. Find the expected value of each alternative.
 - d. Draw the probability distribution on profit for each alternative.
 - e. Calculate the value of perfect information on:
 - Whether or not the Commission approves the high rate increase
 - Whether or not the Commission would grant a low rate increase given that it does not approve the high rate increase
 - Both of the above.
- 3.21 Your company markets an all-purpose household glue called Easystick. Currently, a sister company in another country supplies the product at a guaranteed delivered cost of \$2.00 per unit. You are now thinking about producing Easystick locally rather than continuing to import it. A staff study indicates that with a projected sales volume of 4 million units over the product's life, local production would cost an average of \$1.50 per unit.

However, two things could significantly affect this cost. First, the government in your country is considering imposing a heavy tax on the primary raw material of Easystick. This would increase the average production cost of Easystick to \$2.25 per unit. You think there is a 50/50 chance the government will impose the tax.

The second factor is a newly developed improvement in the production process that uses the expensive raw material more efficiently. This new process would reduce the average production cost of Easystick as shown below.

	Average Cost Per Unit	
	Old Process	New Process
No tax on raw material	\$1.50	\$1.25
Tax on raw material	\$2.25	\$1.75

Unfortunately, local conditions may make it impossible to implement the new process. Your staff estimates a 60 percent chance of being able to use the new process.

Should you continue to import or switch to local production?

- a. Draw the influence diagram for this problem.
 - b. Draw the decision tree for this problem and calculate the expected value for each option. (For the outcome measure, use the total savings in cost relative to importing Easystick.)
 - c. Draw the probability distribution on profit for each option.
 - d. Calculate the values of perfect information and control on:
 - Whether or not the tax will be imposed on the raw material
 - Whether or not the new production process can be used
 - Both of the above.
 - e. Do a probability sensitivity analysis to determine the preferred decision for different ranges of probabilities.
- 3.22 It is 1986 and Shipbuilder, Inc., has decided to take a long, hard look at its telephone needs. Its present system has one major problem: it is very costly to make moves or changes. A task force has identified the most attractive alternative: Fone-Equip can install a system that will enable moves and changes to be made at almost no cost.

The five-year lease on the current system is up for renewal for the period 1987–1991. Another renewal of the lease would be made in 1991 for the period 1992–1996. Fone-Equip’s system is available only for outright purchase. Under any alternative, Shipbuilder will be in the same position in 1996 (a completely new phone system will be needed), so costs beyond 1996 can be neglected.

Shipbuilder is planning a shipbuilding program called Program A, which will entail considerable phone moves and changes in the years 1987–1991. Under the present system, these changes would cost \$2 million per year. However, if Program A does not materialize, there will be very little in the way of phone moves and changes in this period.

Similarly, Program B would entail \$4 million per year in costs for phone moves and changes in the years 1992–1996. However, if Program B

does not materialize, there will be very few moves and changes in this period.

All costs were estimated in 1986 dollars, with the effects of inflation removed. The costs for the present system are \$1.5 million per year lease and \$0.5 million per year recurring costs. Fone-Equip's system has a purchase price of \$14 million (including installation) and \$1.0 million per year recurring cost.

Shipbuilder judged that Program A has a 50 percent chance of occurring. The 1986 telephone decision will be made before Program A's fate is known.

Program B is judged to have a 60 percent chance of occurring if Program A does occur, but only a 30 percent chance of occurring if Program A does not occur. If there is a telephone decision to be made in 1991, it will be made before Program B's fate is known.

There is also a 20 percent uncertainty on what Fone-Equip's system would cost in 1991.

The decision-maker insisted on using a discount rate of 10 percent. At this rate, the net present value in 1986 is:

- \$1 expense in 1987—\$0.91
- \$1 expense in 1992—\$0.56
- \$1 expense per year for 1987–1991—\$3.79
- \$1 expense per year for 1992–1996—\$2.12.

- a. Draw the influence diagram and decision tree for this problem.
 - b. What is Shipbuilder's best strategy? Show the probability distribution on costs for the alternatives.
 - c. What is the value of delaying the 1986 decision until Program A's fate is known?
 - d. What is the value of delaying the 1991 decision until Program B's fate is known?
 - e. Suppose Shipbuilder could obtain perfect information on Program A and Program B before the 1986 decision is made. What would this information be worth?
 - f. Show how the expected values associated with the 1986 decision vary with the probability of Program A occurring.
- 3.23 The pharmaceutical division of Dreamland Products has been the world leader in the area of soporific drugs. Its major product, Dozealot, is approaching the end of its patent life, and already sales have fallen significantly from the peak because of the inroads of new and superior competitive products. However, Dozealot sales are still quite significant and are considered to be of strategic importance for maintaining

the sales of the entire product line of soporific drugs. Therefore, the research and development (R&D) department has defined two alternative approaches to improve the product quality and, thus, future sales prospects.

One approach, which is quite conventional, is simply to reformulate the product to minimize an undesirable side effect that exists in the current galenical form. The manager in charge of galenical development has created a number of new formulations since the original introduction of Dozealot that have the desired characteristic, but even such a simple change in the formulation will require a development expenditure of 500,000 Swiss francs. By estimating the increase in sales of both Dozealot and the rest of the soporific product range, the product manager for soporific drugs has estimated the value of this improvement. Taking into account the production cost of the new formulation and the minor investments required, the improvement in Dreamland's cash flow would be substantial and yield a net present value (NPV) of 2.5 million Swiss francs (not including the cost of development).

The second approach, which is riskier but potentially more rewarding, involves a new controlled-release technology based on differential microencapsulation. This approach, if successful, would not only eliminate the undesired side effects but would also substantially improve the product efficacy. Market forecasts and cash-flow analyses indicate that this product, B, would be four times as profitable to produce and market as the more conventional product A, described above. The drawbacks, however, are that the microencapsulation development project would cost 3 million Swiss francs and still might fail because of an inability to control the differential layering process within tolerances specified by Good Manufacturing Practice. After a recent review of the microencapsulation process development efforts, the R&D director concluded that there is 1 chance in 2 of being technically successful within the deadlines imposed by the patent life of Dozealot.

- a. Draw the influence diagram and decision tree for Dreamland Products considering it could separately pursue the development of A or B or even pursue both to reduce the risks involved in B. In the latter case, it would naturally market B and not A if B were a technical success.
- b. Compute the expected value of each alternative, assuming B has 1 chance in 2 of being successfully developed. Based on the criterion of expected value, what should Dreamland do? Since the probability of successfully developing product B is difficult to determine, Dreamland's managing director would like to know how sensitive the best decision is to this probability assignment.

- c. Compute the expected NPV of developing B alone and of developing A and B simultaneously as a function of p_B , the probability of technical success of B.
- d. Graphically represent the expected value of the three alternatives as a function of p_B and determine for what range of values of p_B each alternative is best.

Resolving the technical uncertainties surrounding the microencapsulation project early could be achieved by immediately conducting a few critical experiments at an additional cost of 1 million Swiss francs.

- e. Compute the expected value of perfect information on whether microencapsulation could be successfully accomplished, assuming an initial probability of success of .5.
 - f. Compute and graphically represent the expected value of the entire project given perfect information about whether or not microencapsulation would be feasible as a function of the initial probability of success p_B . Graphically show the expected value of perfect information and determine over what range of initial probability of success p_B this value exceeds 1 million Swiss francs.
- 3.24 The internal rate of return (IRR) of a venture is the value that, if substituted for the discount rate, makes the net present value (NPV) zero. There are several difficulties in using IRR as the sole criterion in choosing between alternatives.

Consider the following two ventures:

- A: Invest I in year 0, receive positive cash flow, C , in all years from year 1 to infinity
 - B: Invest I in year 0, receive positive cash flow, K , in year 1, nothing in succeeding years.
- a. Solve for $c = C/I$ and $k = K/I$ in terms of $n = NPV/I$ and d , the discount rate. Use the sum

$$\sum_{i=1}^{\infty} \frac{1}{(1+d)^i} = \frac{1}{d}$$

- b. Assume that both ventures have the same investment and the same NPV. Plot the IRR value against n for the two ventures. (Take the discount rate d to be .1.)
- c. For equal n (equal value to the decision-maker), how do the IRR values of long-term and short-term investments (A and B) differ? Is IRR an adequate decision criterion? What assumptions do you have to make concerning use of funds after year 1 for investment B?

- 3.25 A venture with an NPV of zero is acceptable to the decision-maker. Why is this true? What problems in interpretation does this cause for the decision-maker unfamiliar with the concept? What steps would alleviate these problems?
- 3.26 If financing is available at an interest rate equal to the decision-maker's discount rate, a venture with a positive NPV can be transformed into a venture with an infinite IRR.
- How can this be done?
 - Does a similar situation exist for leasing alternatives?
 - For existing businesses, one alternative usually involves no major new investments. What is the IRR for this strategy? Compare the type of results you might expect for NPV and for IRR for "Invest" and "No New Investment" strategies. In what situations might you prefer NPV or IRR as a decision criterion?

4

Probabilistic Dependence

Dependence and Independence

The probabilities in the preceding chapters were simpler than those one finds in most real decision problems. There was no probabilistic dependence; it made no difference to the person providing the probabilities at one node if he or she learned what happened at another node. Real decision problems usually start with a set of variables for which the probabilities are very interdependent. One of the subtler skills in influence diagram and decision tree construction is finding a set of variables for the nodes that have independent probabilities and that are natural and meaningful. But it is not often that this can be done for all the variables in the problem.

Usually the facilitator faces a situation described by variables that are far from independent. How are uncertainties related? How will resolving uncertainty on one variable affect the uncertainty in another area? Questions like these frequently lie just below the surface of many deliberations, but they can be extremely difficult to enunciate. One of the advantages of influence diagrams and decision trees is that they provide a language that makes dialog in this area possible and relatively easy.

This chapter deals with three related topics. First, the Positronics case is used to demonstrate why and how dependent probabilities are used in real decision problems. Second, dependent probabilities are used to combine new information with prior information. The process for doing this is called Bayes' Rule. Third, Bayes' Rule and decision trees are used to show how to calculate the value of imperfect information.

Dependent Probabilities

Obtaining the Data—

Achieving clarity when dealing with dependent probabilities is often both difficult and important, particularly when information from different people needs to be combined. This is illustrated in the Positronics example.

After a presentation on the MegaCorp bid, the Positronics production manager came to the decision facilitator. After seeing the presentation and learning more about how the probabilities were used, he had reconsidered the probabilities that he had provided for Positronics' costs. If he knew what Anode's bid was, he would assign different probabilities to Positronics' costs. He had great respect for the production staff at Anode and had heard it had much more experience than Positronics in making instruments of the type being bid on. Thus, Anode should have higher quality information and much less uncertainty about costs than Positronics. Positronics should have the same costs as Anode, except that Anode's experience would give it a small cost advantage. The facilitator tentatively drew an arrow from the Anode Bid node to the Positronics Cost node to show this influence (Figure 4–1a). Using this diagram, the production manager would provide the probabilities for the distribution tree and the marketing department would give the probabilities for Anode Bid.

After some thought, however, the facilitator realized they had it backwards. It was the probabilities for Anode's bid that should be reassessed. The original assessment on Anode's bid was done without any real input on Anode's cost and would be a meaningless starting point for an assessment of anyone's cost. However, the production manager had added two crucial pieces of knowledge: Anode will know its costs accurately when it makes its bid, and Anode and Positronics both have approximately the same costs because they use the same process. Thus, given a particular level of cost, the real uncertainty is how much over cost Anode would bid. The people who would best know how much Anode wanted the business and how much of a margin it would figure in were in marketing. Accordingly, the arrow was reversed, and the probabilities in Figure 4–1b for Anode's bid were obtained from the marketing department.

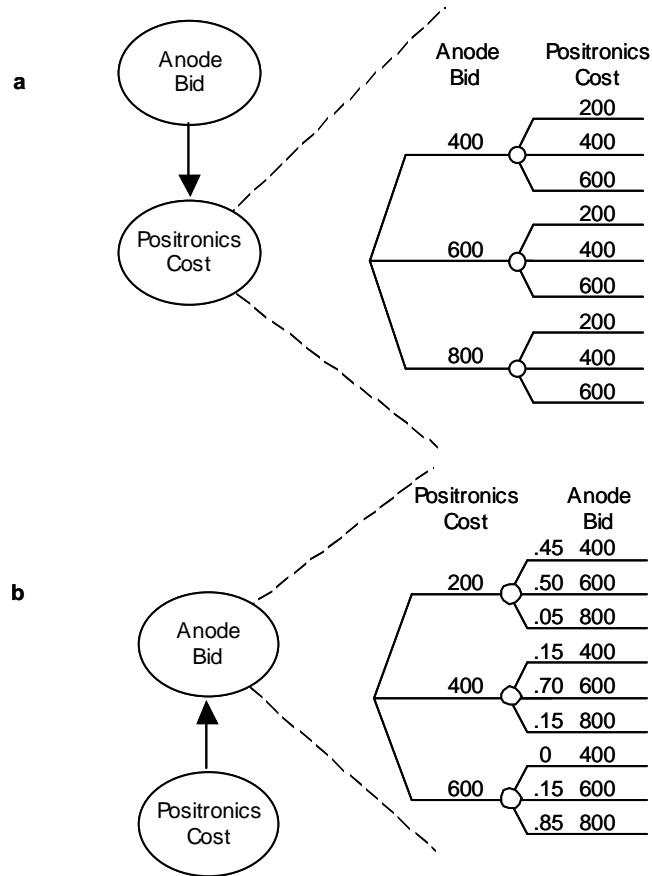
A discussion such as this would be very difficult without the language of influence diagrams and probability trees. Most people untrained in probability theory or decision analysis find it difficult to communicate information concerning uncertainty; interdependence among uncertain factors is almost impossible to talk about. The graphical structures of influence diagrams and trees give a visual, simple, and systematic way of discussing structure and data in this area.

Using the Data

When we have chosen the order of nodes and assessed a set of dependent probabilities, we will have a set of probabilities that depends on the branch or outcome of another node. The second node is “relevant” to the first node.

Figure 4-1

Structuring the Problem and Identifying Data Requirements

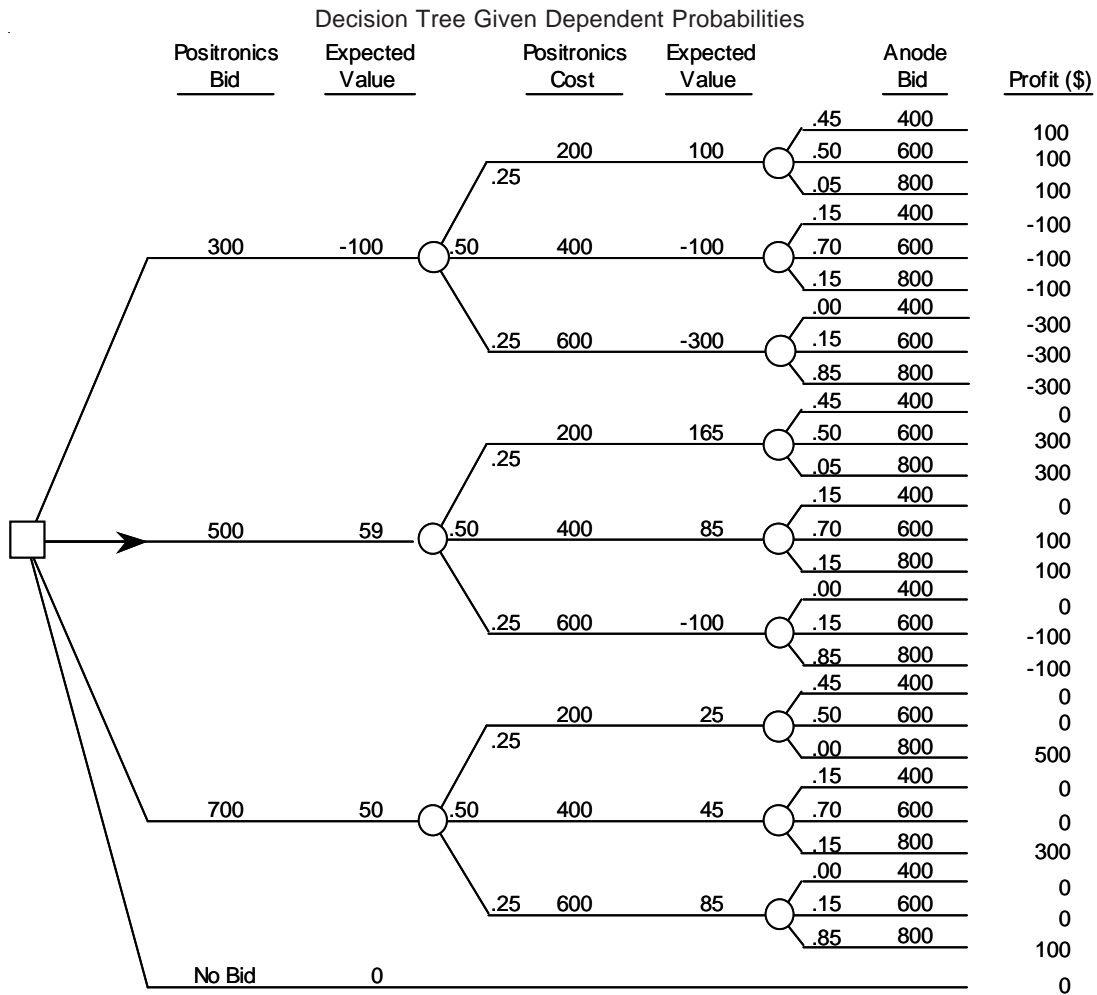


The second node is called the “conditioning” (or, in older terminology, the “influencing”) node and is the one at the base of the arrow in the influence diagram.

It is easy to input these dependent probabilities in a tree. Figure 4-2 shows the tree in completely rolled back form, with expected value used as the decision criterion. Because of the order in which the probabilities were assessed (Figure 4-1), the Positronics Cost node is to the left of the Anode Bid node. Later in this chapter, we will see how to reverse the order of these nodes.


The cumulative probability distributions for this tree are shown in Figure 4-3. Comparing this with the previous cumulative plot (Figure 3-6), we see that the risk has been intensified: there is a larger probability of losing money if Positronics wins the bid. This makes sense, because the dependent probabilities tell us that if Anode bids high, Positronics wins the bid but its costs are likely to be high.

Figure 4-2



This effect would have been missed if the probabilistic dependency were missed. A common pitfall for the novice decision facilitator is missing the interdependence or relevance between uncertainties. These interdependencies can often dramatically increase or decrease the uncertainty in the value measure. The use of influence diagrams helps avoid this problem; arrows indicate probabilistic dependence or relevance, and most people deal naturally and easily with these arrows.

Figure 4-3



Dependent Outcomes

Obtaining the Data

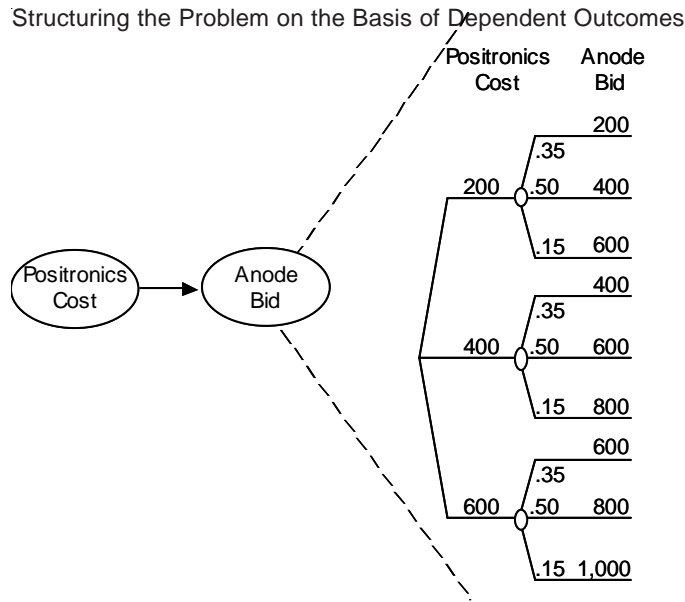
The next morning, the marketing people returned to the decision facilitator and said they were really uncomfortable with the probabilities they had supplied the previous day (the dependent bid probabilities shown in Figure 4-1). After thinking about them, they wanted to be able to systematically change the outcomes rather than the probabilities. They presented the decision facilitator with the distribution tree shown in Figure 4-4.

Dependent outcomes are another way of representing dependencies among uncertainties. In this case, it is not the probability of an outcome that changes, but the outcome itself.

Having dependent outcomes is another way of representing probabilistic dependence for continuous variables. In our discrete approximation, we chose ranges and then assigned probabilities to those ranges. Probabilistic dependence means that the probabilities assigned to these ranges change depending on which branch of the conditioning node we are on. Alternatively, we can readjust the ranges (as we move from branch to branch of the conditioning node) to keep the probabilities the same.

In our experience, using dependent outcomes is frequently a sign of systematic considerations in the thought process of the person supplying the data. To promote clarity of thought and communication, we can use the

Figure 4-4



deterministic value model (rather than probabilistic dependence) to capture this systematic consideration in a deterministic way. For instance, dependent outcomes might result from a thought process like “Shift the value up by X in that case and then put a 10 percent uncertainty in either direction.” Remembering our divide and conquer technique, we can improve the accuracy of the assessment by separating these two effects and assessing them individually.

Using the Data

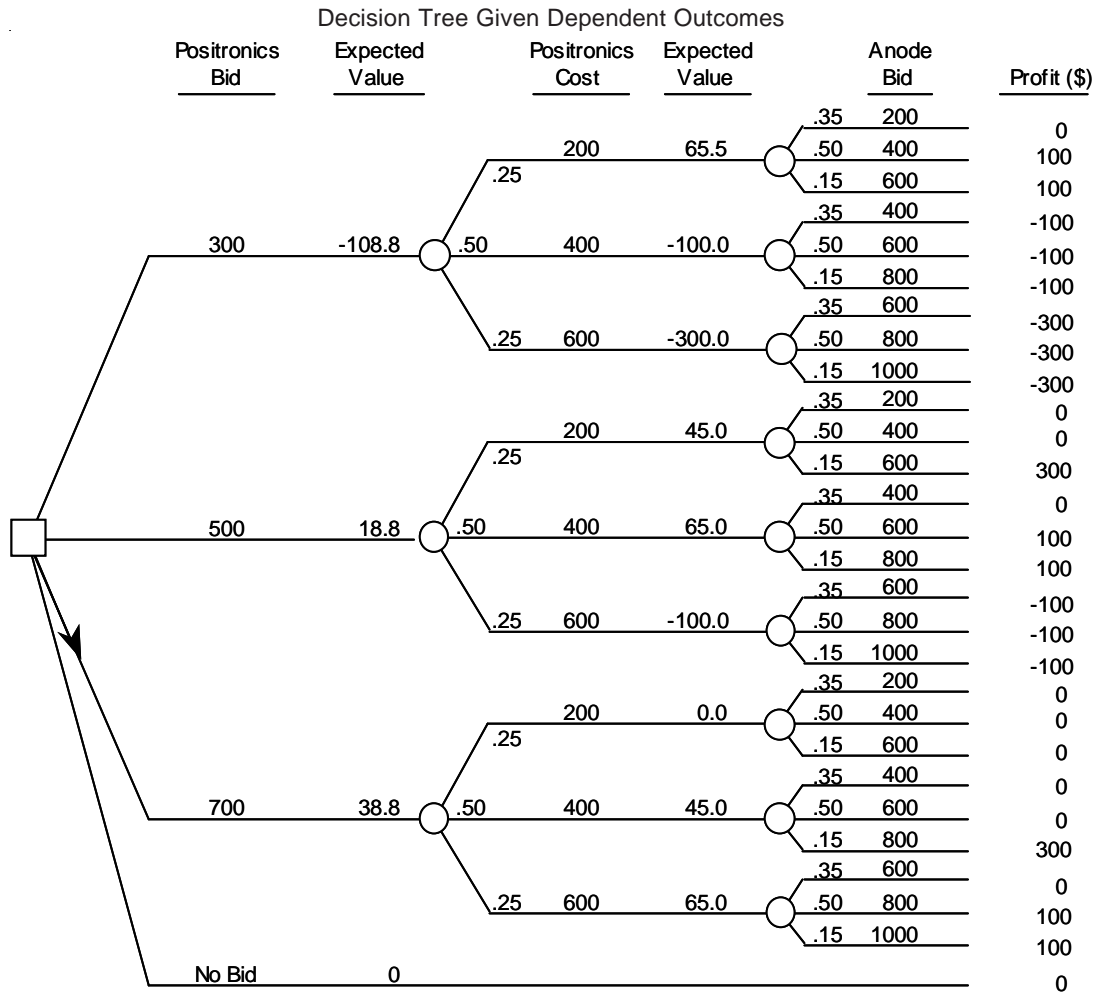
Entering dependent outcomes in the tree is not difficult. Figure 4-5 show the tree rolled back for the expected value criterion. The cumulative probability distributions associated with this tree are shown in Figure 4-6.

We see that the risk has intensified (as indicated by the higher probabilities of low and high outcomes), and the even larger probabilities of losing money lead to an expected value of only \$38,800. The preferred alternative has shifted from a bid of \$500,000 (Figure 3-7 and 4-2) to one of \$700,000 (Figure 4-5). The \$500,000 option has an expected value that has dropped to \$18,800.

Nature’s Tree

Dependent probabilities are important in dealing with a problem we address informally in our daily lives: How do we incorporate new information into our current information base?

Figure 4-5




Indicators and States of Nature

In the section "Dependent Probabilities", we stressed the importance of drawing the relevance arrow correctly when encoding dependent probabilities. There is one type of situation in particular where correct ordering is imperative and where people frequently make mistakes, both in daily life and in decision analysis: when people describe the relationship between an imperfect indicator and the state of nature (e.g., the relationship between medical symptoms and the actual state of health). To establish this relationship correctly, we follow a two-step process. First, we track the indicator to see how accurate its predictions are. Second, these results are incorporated into a tree, which is used to predict the state of nature, given an indicator result.

Figure 4-6

Probability Distribution on Profits Given Dependent Outcomes



The following are examples of state of nature/indicator pairs that occur in business situations.

State of Nature

Market size
 Number of future competitors
 Amount of recoverable oil
 Economic climate
 Production costs
 Drug safety and efficacy
 Software stability

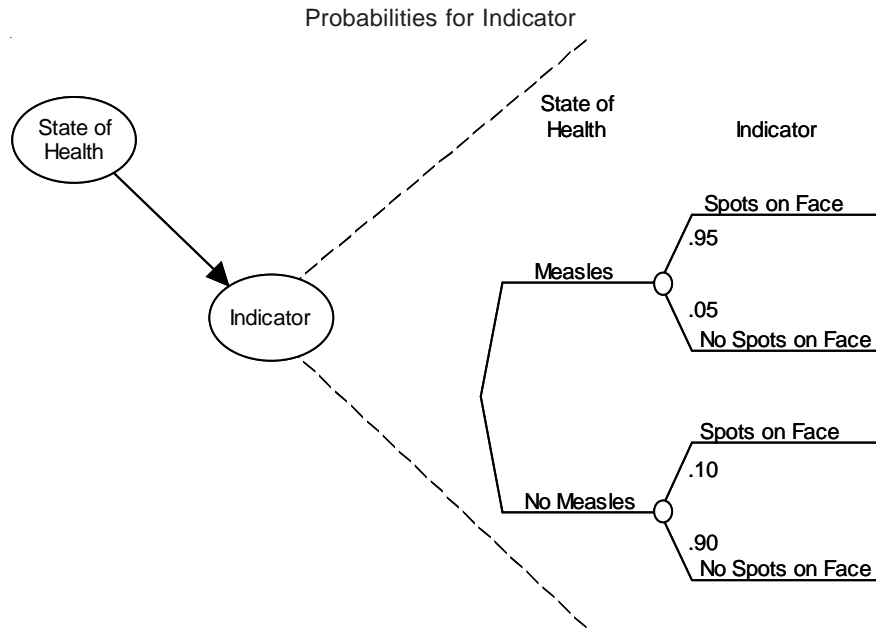
Indicator

Market survey
 Articles in business journals
 Seismic study
 Leading indicators
 Prototype costs
 Clinical trials
 Beta tests

An Example from Medicine

Let us use a simple medical example. Through either experience or a controlled test, we track the number of people with and without measles who have spots on their face. Using this and other data, we can find the probability that someone who has spots on his or her face has measles. Unfortunately, people make mistakes all too often when they try to go through this process. The problem is that they do not correctly combine their prior information with the indicator data. For instance, assume that most people with measles have spots on their face and that few people without measles have spots on their

Figure 4-7



face. If you know (prior or general information) that there is a measles epidemic, spots on a person’s face would lead you to put a high probability on measles. If, on the other hand, there has not been a case of measles in the city for months, spots on a person’s face would lead you to put a much lower probability on measles.

How do we correctly mix our prior or general information with the data (indicator results and track record of the indicator)? The rule for performing this process correctly is to first construct Nature’s tree. In the influence diagram corresponding to Nature’s tree, the state-of-nature node is relevant to the indicator result. The data then supply the probabilities for the indicator results, given the state of nature. We have to supply our subjective (prior, general information) probabilities for the state-of-nature node.

Suppose a doctor has tracked the correlation between the state of health and the indicator for her patients. Suppose she has come up with the Nature’s tree in Figure 4-7. (The probabilities in this distribution tree are for illustration only.)

Now let us assume there has been a measles epidemic, and the doctor judges that the next patient who walks into her office has a 20 percent chance of having the measles (Figure 4-8).

In the probability tree (Figure 4-9), the column of joint probabilities (that is, the probability that both events will occur) has been calculated by multiplying the probabilities at each branching for a path through the tree. For instance, the topmost path (Measles, Spots on Face) has a probability of .19 ($20 \times .95$).

Figure 4-8

Probabilities for State of Nature

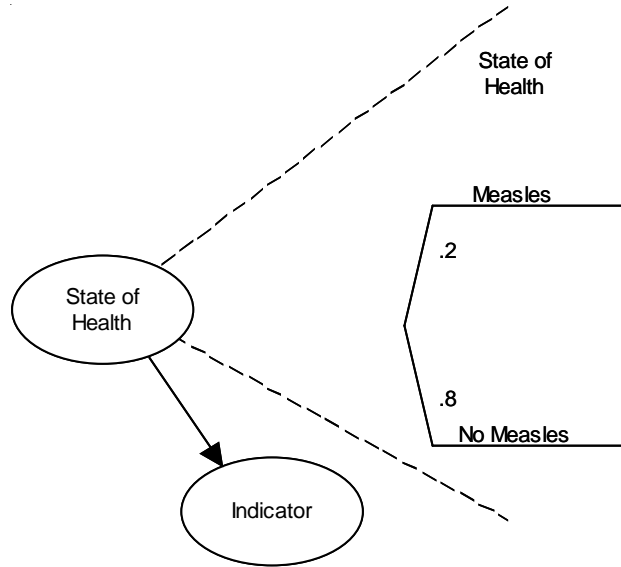


Figure 4-9

Probability Tree for Medical Example

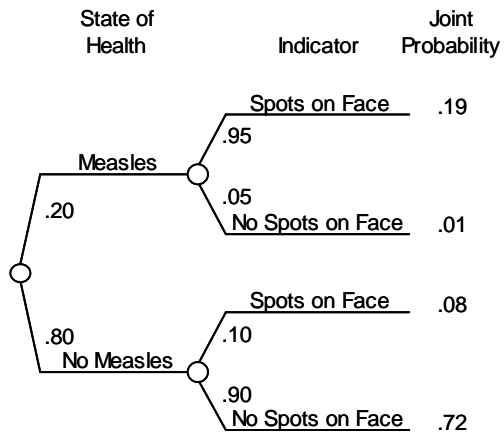
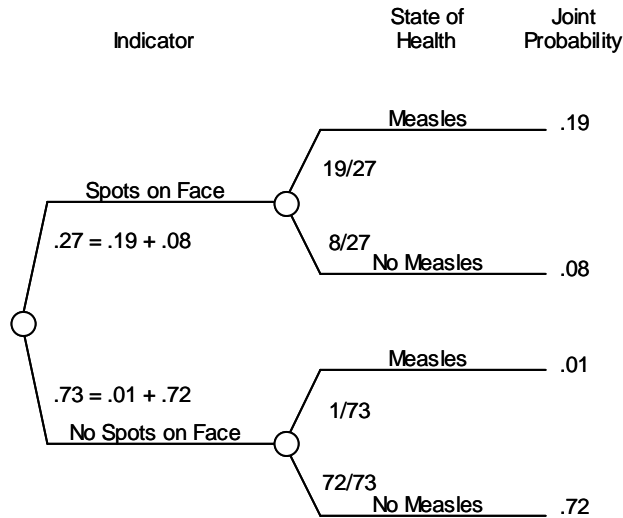


Figure 4-10

Applying Bayes' Rule to Reverse the Tree Shown in Figure 4-9



Bayes' Rule

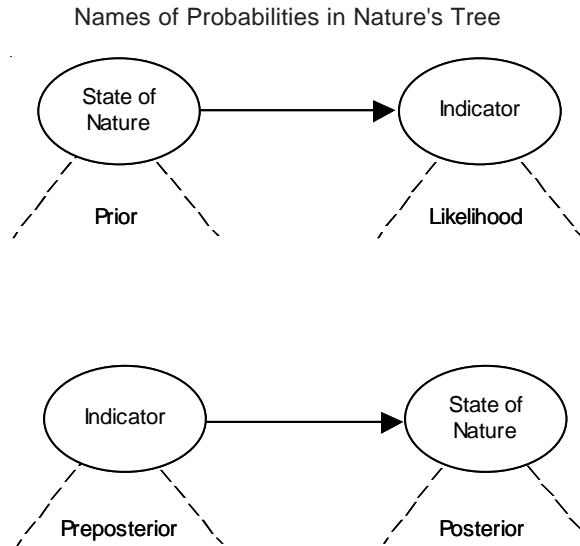
To find the probabilities in the order needed for most decision-making, we use Bayes' Rule to reverse the order of nodes, putting the indicator node first and the true state of nature second. Thus, we can read off the probability of the true state of nature given an indicator result.

To apply Bayes' Rule to trees, we first calculate the joint probabilities from Nature's tree. Second, we draw the tree with the nodes in reversed order and transcribe these probabilities to the ends of the appropriate paths in the reversed tree. Then we can find the probabilities for the node on the left by adding the probabilities of all paths passing through that branch. The probabilities for the branches of the node on the right are found by dividing the joint probability at the end of the branch by the probability for the branch of the node on the left. The process is fairly simple but tedious.

For the measles example, we can use this process to reverse Nature's tree (Figure 4-10). Thus, we see that if a patient with spots on his or her face walks into the doctor's office, there is a $19/27 = .70$ chance the patient has measles. There is a 27 percent chance that the next patient who walks in will have spots on his or her face.

The probabilities in Nature's tree have traditional names (Figure 4-11). In Nature's tree, the probabilities of the state of nature are called the prior probabilities; the probabilities of the indicator (the second node) are called the likelihood function. The rule for reversing the order of nodes in the tree is called Bayes' Rule. Reversing the order of nodes in the tree is indicated in the influence diagram by reversing the arrow between the nodes. The probabilities of the second node of the reversed tree (the state of nature) are

Figure 4-11



called the posterior probabilities; the probabilities of the first node in the reversed tree are called, believe it or not, the preposterior probabilities.

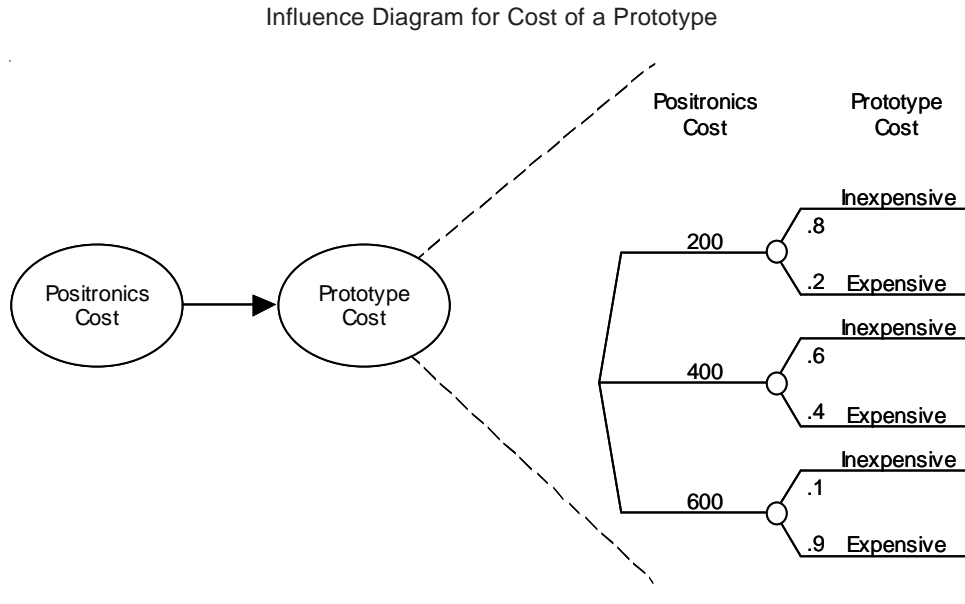
Why is Bayes' Rule so important? It gives the method for correctly adding data and new information to an existing subjective state of information (the prior). This information can then be used to construct the posterior distribution, which represents an updated state of information. People commonly make mistakes in using new information. Although it appears rather formal, Bayes' Rule is the best method for avoiding mistakes in judgments that often have great importance.

In summary, when dealing with indicators, first draw Nature's tree: the tree with the state of nature first and the indicator second. Then, enter the data into this tree. Only in this way will we be able to correctly separate our expectations about what the future will bring from how good the indicator is. The process of reversing the tree is important when we wish to find out how knowledge of an indicator result changes our expectations of what the future will bring. This procedure will become clearer in the next section, "The Value of Imperfect Information."

A Prototype As an Indicator

The production engineer suggested that Positronics could resolve some of the uncertainty on production costs by building a prototype instrument and carefully monitoring the cost of the steps that went into building it. The cost of the prototype would be an indicator for the state of nature, the actual costs. In Nature's tree for this procedure, the first node is the actual costs and their probabilities, as used earlier. The next node is the prototype cost, which was felt to be best described as inexpensive or expensive. (The production engineer gave

Figure 4-12



a technical specification for these rather vague terms.) Given the difference between prototypes and real production, the engineer estimated the probabilities for prototype cost, as shown in Figure 4-12.

We emphasize that the probabilities must be assessed in this order, although there is a real and constant temptation to assess them in the reverse order. The reason for using this order is that it clearly separates expectations about what the future will bring (the probabilities on the first node) from the reliability of the test (the probabilities of the second node). Assessing probabilities with the nodes in the opposite order completely mixes the two sets of ideas.

The Value of Imperfect Information

One of the most important uses of Bayes' Rule and Nature's tree is in obtaining the value of imperfect information. If we obtain imperfect information, we know the outcome of the imperfect indicator before we make a decision—and before we learn the true state of nature. This order requires reversing Nature's tree.

The president of Positronics was well aware that perfect information is seldom available. Most often, the choice is whether or not to undertake a study to reduce the uncertainty. Figures were now available (Figure 4-12) that would permit more quantitative analysis of the value of building the prototype instrument to reduce the uncertainty on cost (see Chapter 3, "The Value of Perfect Control").

Figure 4-13

Influence Diagram Depicting the Positronics Prototype Build Decision

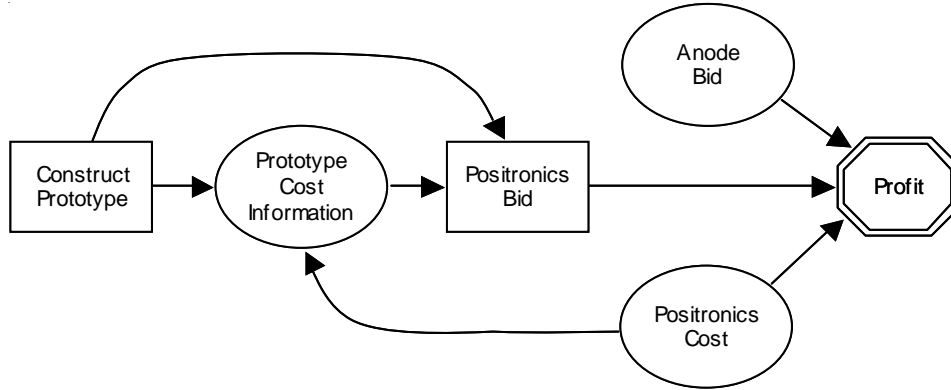


Figure 4-14

Distribution Tree of Prototype Cost Information Node

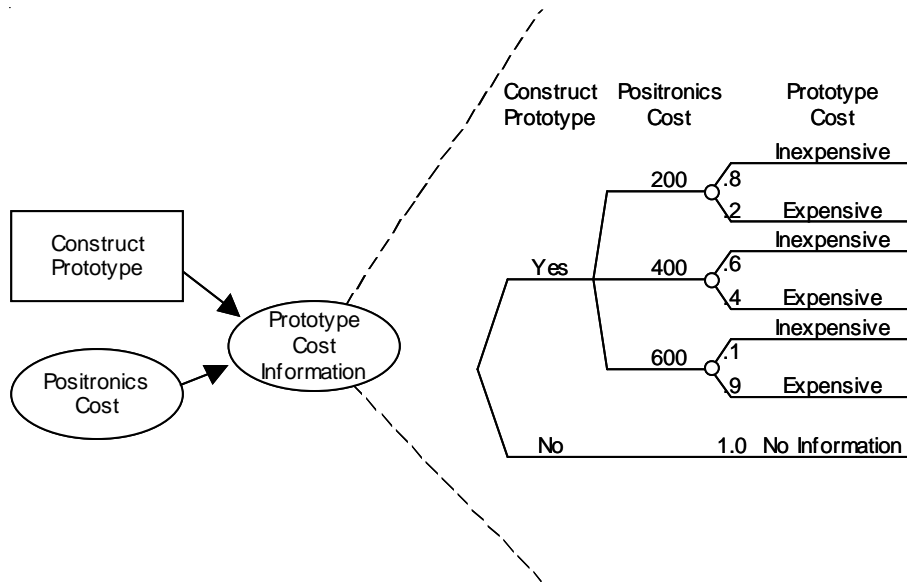


Figure 4-15

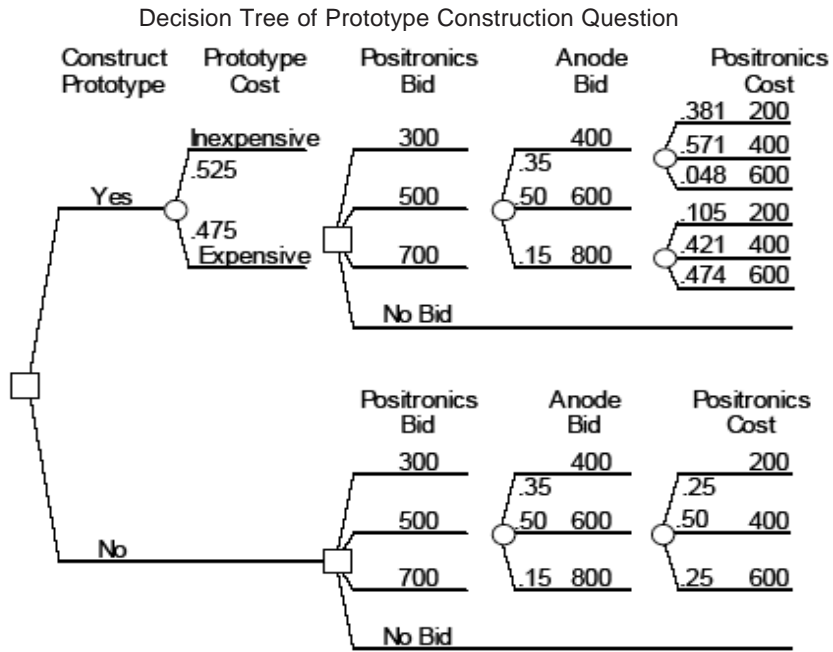
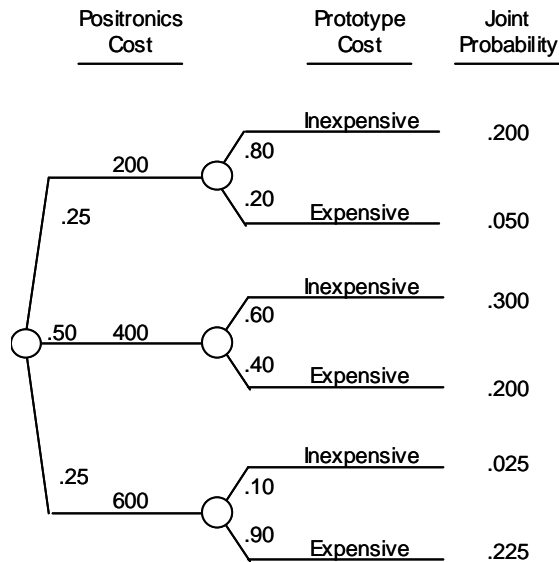


Figure 4-16

Probability Tree for Prototype Example



One way to represent this extended decision problem is shown in Figure 4-13.

The Construct Prototype node contains the alternatives “Yes” and “No.” The distribution tree contained in the Prototype Cost Information node is shown in Figure 4-14.

The decision tree (Figure 4-15) is constructed according to the rules given in the section “Recasting the Problem as a Decision Problem” in Chapter 3. Note the difficulty in representing the probabilities in this tree: For the “Yes” branch of the tree, the upper node for Positronics Cost corresponds to the “Inexpensive” branch of Prototype Cost, and the lower node to the “Expensive” branch.

Where did the probabilities for Positronics Cost in Figure 4-15 come from? Just as with the medical example, we have to use Bayes’ Rule to reverse the order of the nodes. Figure 4-16 shows Nature’s tree, and Figure 4-17 shows the reversed tree. The probabilities in Figure 4-17 are in the form needed for the tree in Figure 4-15.

The easiest way to find the value of imperfect information is by just evaluating the subtree following the “Yes” branch in Figure 4-15. We can do this by adding the node on Prototype Cost onto the front of the existing tree and by then comparing its value with that of the original tree (the subtree at the bottom of Figure 4-15). This tree is shown in Figure 4-18.

The expected value with the imperfect information is \$73,000. We follow the same calculation method we used for the value of perfect information.

Figure 4-17

Applying Bayes’ Rule to Reverse the Tree Shown in Figure 4-16

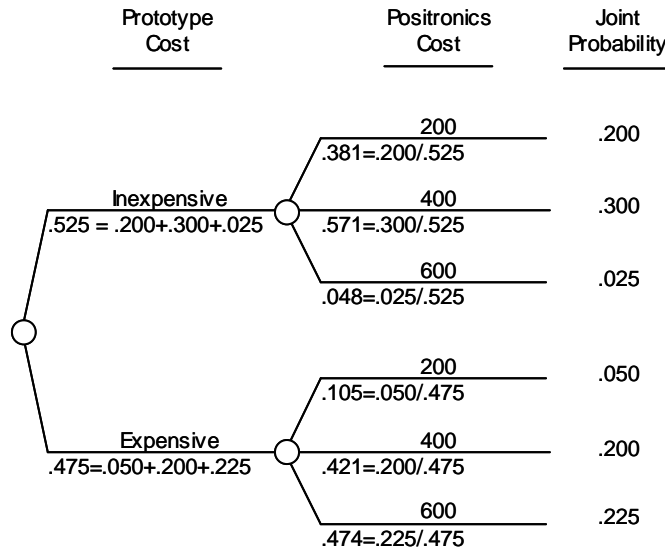
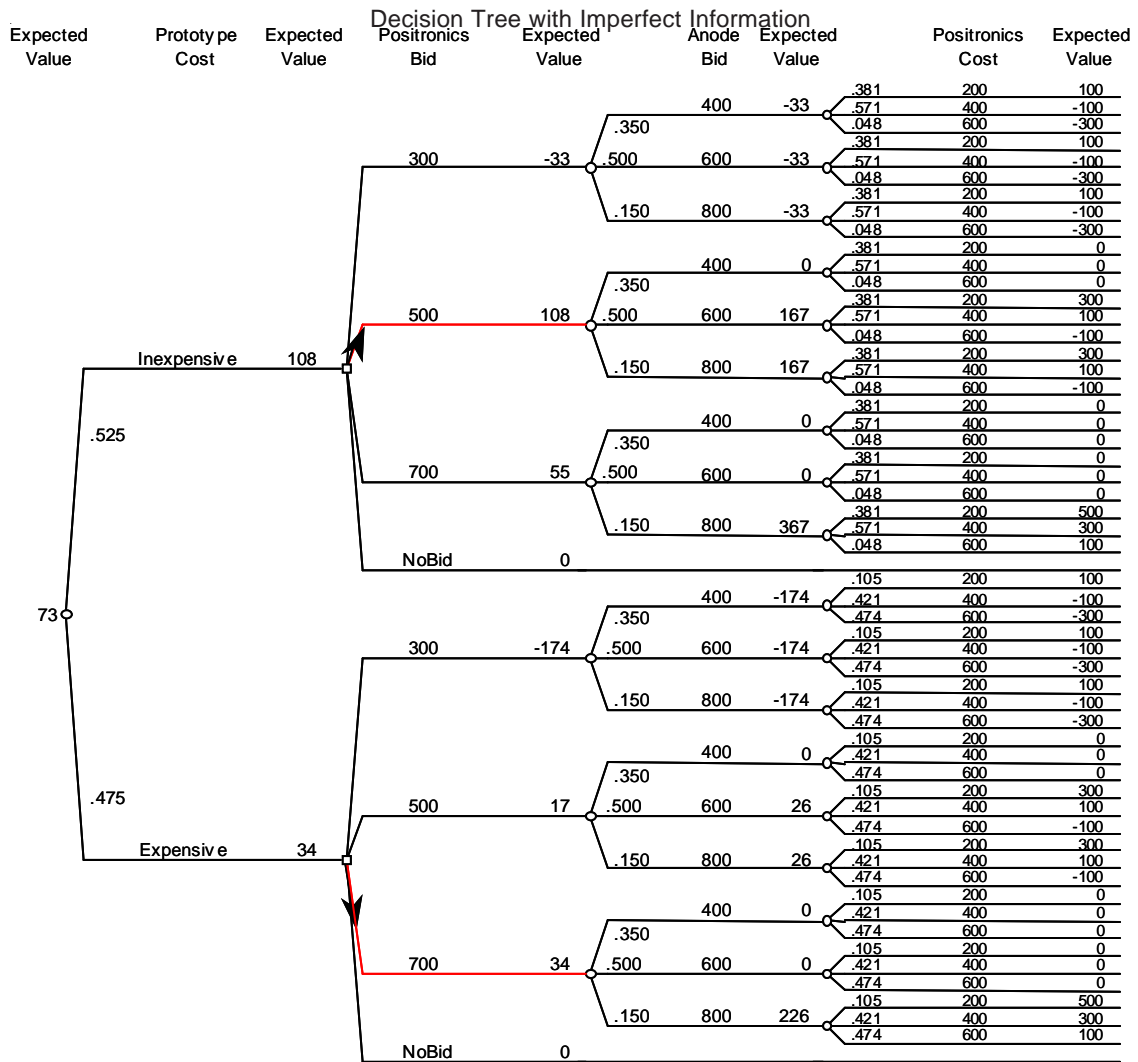


Figure 4-18




$$\begin{aligned}
 \text{Value of Imperfect Information} &= \text{Value with Imperfect Information} \\
 &\quad - \text{Value Without} \\
 &= \$73,000 - \$65,000 \\
 &= \$8,000 \tag{4-1}
 \end{aligned}$$

Thus, the prototype program should not be undertaken if it costs more than \$8,000.

As with perfect information, there is an alternative method for calculating this value. From Figure 4-18, you can see that about half the time (probability .475) the prototype information will lead you to switch your bid (hence, the information has value). If the prototype is expensive, we switch

Figure 4-19

Influence Diagram Showing Probabilities Dependent on a Decision Node



the bid from \$500,000 to \$700,000. This improves the value by \$33,900 – \$17,100 = \$16,800. The expected improvement is then \$16,800 × .475 = \$8,000.

The Value of Imperfect Control

The value of imperfect control is much simpler to calculate. Assume that a feasible cost control program is put into effect. Obtain the probabilities for Positronics' costs given this cost control program. Presumably, the probability of the high-cost branch is smaller than before. Find the expected value of the tree using these new probabilities. This is the value with control for this (imperfect) cost control program. As before, the value of control is the improvement in value.

$$\text{Value of Imperfect Control} = \text{Value with Imperfect Control} - \text{Value Without Imperfect Control} \quad (4-2)$$

Common Interpretive Problems with Reordering the Tree

Care must be taken in reordering decision trees. The decision facilitator is normally concerned with decision-dependent probabilities or decision-dependent outcomes. The information in this section will help you in these cases.

Decision-Dependent Probabilities

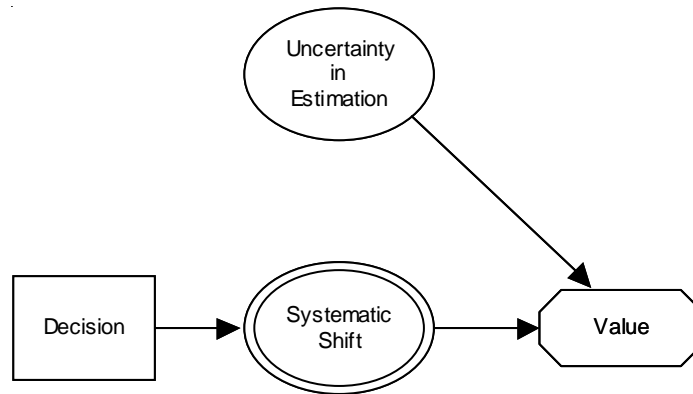
Probabilities can depend on a decision node. For instance, suppose we introduced a decision on whether or not to institute stringent cost control measures. The probabilities for Positronics' costs would then depend on which branch of the cost control decision node we were on (Figure 4-19).

Suppose we now decide to move the cost node in front of the cost control decision node in the decision tree. This would be equivalent to drawing an arrow from the Positronics Cost node to the Cost Control node. Decision analysis theory would object that this was not a valid order for the tree because of decision-dependent probabilities.

There are three ways to understand why the reordering is not allowed. The first way is operational. To reverse the order of nodes with dependent probabilities, Supertree uses Bayes' Rule. However, a decision node has no probabilities to insert into Bayes' Rule!

Figure 4–20

Influence Diagram Illustrating Handling of Uncertainty in Estimation



The second, more fundamental way of understanding the problem is to realize that this would create a loop in the influence diagram. Since arrows represent a flow of information, a loop makes no sense. This conclusion is reinforced by the chronology implied by arrows into and out of decision nodes. (See Chapter 11.)

The third way of understanding the problem is the second way reexpressed in decision tree language. If the Positronics Cost chance node occurs before the Cost Control decision node, it means we know what our costs are going to be before we make the decision. However, knowing that costs were low, for instance, would argue that we are probably going to make the decision to implement the stringent cost controls, which will make costs low. This gets us into a logical loop. No matter how we put it, decision trees will not allow us to interchange the order of a decision node and a chance node whose probabilities depend on the decision node.

Putting all influence diagrams in the Howard Canonical Form (see Problem 11.10) would avoid all arrows leading from decisions to uncertainties. In this case, the oval for Positronics Cost in Figure 4–19 would be replaced by several ovals, Positronics Cost under Cost Control alternative X, one oval for each alternative. In this formulation, the state of knowledge is clearly separated from the choice of alternative, and the problem of decision-dependent probabilities does not occur.

Decision-Dependent Outcomes

For chance nodes whose outcomes depend on a decision node, we would expect the same problem as with decision-dependent probabilities. After all, dependent probabilities and dependent outcomes are both ways of characterizing dependent probability distributions. However, there is an interpretation of dependent outcomes that allows us to interchange the nodes.

The interpretation is that we learn whether we are on the high branch, the middle branch, or the low branch (assuming there are three branches) before we make the decision. We do not know the value of the outcome, only whether we are on the top branch, the second branch, or the third branch. This interpretation is consistent with the observation that dependent outcomes often represent a deterministic shift combined with uncertainty in estimation. The information gained in interchanging the nodes is on the uncertainty in estimation. This interpretation is illustrated in the influence diagram in Figure 4–20.

Summary

A dependency can be expressed with dependent outcomes or with dependent probabilities. These forms are equivalent, and the selection of form usually depends on how the source of information thinks about the uncertainty. This was illustrated in a probabilistic dependence that increased the riskiness of Positronics' bid decision.

We also discussed how correctly handling dependencies is a common problem that arises when dealing with tests or information gathering (Nature's tree). We then applied this to the calculation of the value of the imperfect information that could be obtained if Positronics were to build a prototype.

Finally, we discussed some interpretive difficulties that commonly arise when reordering trees with decision-dependent probabilities or decision-dependent outcomes.

Problems and Discussion Topics

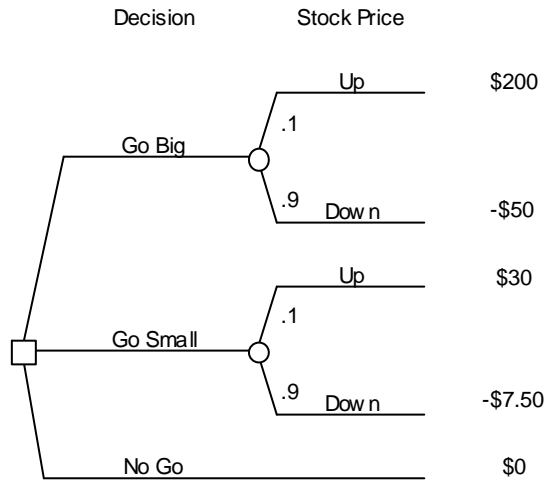
- 4.1 What are joint probabilities? How are these different from conditional probabilities?
- 4.2 Suppose your interview with a decision-maker has revealed a dependency between two uncertainties. How do you determine which uncertainty depends on the other for assessment purposes? What does this imply for the order of assessment?
- 4.3 "Flipping the tree" refers to reordering the nodes in a tree or part of a tree (reversing them if there are only two nodes). Since the flipped tree represents the same state of knowledge (uncertainties) as the original tree, the principle of flipping a tree is that the probability of any event derived from the original tree should be the same as the probability of the same event represented by the flipped tree.

Consider the medical example on page 80. Suppose a medical expert assigns a probability that 50 percent of the total population currently has measles, and that given a person has measles there is a chance of 85 out of 100 that the person has spots on his or her face. If a person

does not have measles, there is a .95 chance that the person has no spots on his or her face.

What is the conditional probability that someone who has spots on his or her face has measles? What is the conditional probability that someone who does not have spots on his or her face has measles?

- 4.4 We already know that changing the order of two chance nodes (called flipping the tree) does not change the knowledge represented by that tree. What happens if we move a chance node from the right of a decision node to the left of it? And vice versa?
- 4.5 The radiator in your car tends to overheat, but you have not fixed it because it is still winter and cold outside. The radiator overheats only 5 percent of the days it is used in cool weather. However, it overheats 70 percent of the time in warm weather. The weather report has just predicted a 1 out of 5 chance of warm weather today.
- Draw the influence diagram for these relationships.
 - What is the chance your radiator will overheat today?
- 4.6 After having had pizza delivered at 11 p.m. several times a week for a number of years, you decide that there is a 70 percent chance that a pizza with a visible amount of cheese has a visible amount of pepperoni. You also figure that the probability that a randomly selected pizza will have visible amounts of cheese and pepperoni is .40.
- Draw the influence diagram for these relationships.
 - What is the probability that a randomly selected pizza has a visible amount of cheese?
- 4.7 Using first dependent probabilities and then dependent outcomes, write down your probabilities on the temperature outside given that it is 9 a.m. or 9 p.m. Assume a .5 probability that the observation is made at either time and make your chance node on temperature have three branches.
- Compare the expected temperature from each method (dependent probabilities or outcomes). How different are they? Which method enabled you to give the better estimate and why? What does this tell you about the underlying process affecting the temperature? What does it tell you about your thought process?
- 4.8 An expected-value decision-maker faces the following short-term investment in a given stock:



- Draw the influence diagram. Do you have enough information to do so?
 - Calculate the expected value of this decision.
 - What is the maximum amount the decision-maker should pay for perfect information on the stock price?
 - Suppose there is a test that can predict the stock price with an accuracy of $.9$. Draw the influence diagram for this. What is the maximum amount the decision-maker should pay for this test?
 - Suppose the test says the stock will rise and this information is given to the decision-maker for free. Now what is the value of perfect information on the stock price?
 - Suppose a wizard comes along. He can make any possible outcome happen the decision-maker desires. Draw the influence diagram for the value with the wizard. What is the maximum amount the wizard should be paid?
- 4.9 Ursa Major Movies (UMM) has been trying a blind test on all its movies before releasing them. The test labels a movie as a “Hit” or a “Dud.” To make the test blind, UMM released all movies regardless of the test result. The test result and the actual history of the movies are shown in the following table.

	Test Result	
	“Hit”	“Dud”
Broke box office records	5	1
Run of the mill	13	7
Disaster	8	9

- a. Draw the influence diagram for the test and movie results. Is it in the order of Nature's tree? Why or why not?
 - b. What is the probability that a film chosen at random out of the studio's past movies was a disaster? Run of the mill? Broke records?
 - c. What is the probability that a disaster had previously been labeled a "Hit"? What is the probability that a box office record breaker had been labeled a "Dud"? Is the test better at detecting good or bad movies?
 - d. The producer thinks that a new movie is really quite good (a 5 in 10 chance of being a box office hit, a 3 in 10 chance of being run of the mill). After learning that the test came up "Dud," how should the producer revise her probabilities?
 - e. The president thinks that this new movie is like all the others, meaning that the historical frequencies above apply. What should he think after learning that the test result for this movie was "Dud"?
 - f. Assume a record breaker gives the company a net profit of \$20 million; a run of the mill, \$2 million; and a disaster, a net loss of \$2 million. What value would the producer place on the new film before learning the test results? After learning the test results? How about the president?
- 4.10 C. Thompson, the credit manager of IJK Industrial Products, considered extending a line of credit to Lastco Construction Company. Lastco was a new company and was definitely considered a credit risk. Drawing on his experience, Thompson said, "There is about a 30 percent chance Lastco will fail within the year, which means a severe credit loss. And the way these construction companies operate, I would say there is another 25 percent chance Lastco will run into serious financial trouble." After being further questioned about other possibilities, Thompson said, "If they don't run into financial problems, there still is less than a 50/50 chance of Lastco becoming a regular customer. I would say the odds are about 5 to 4 that Lastco will end up being a sporadic customer." Thompson also made the following predictions:
- If Lastco failed completely, it would average purchases of \$1,500 before failing but leave an average unpaid balance of \$800, which would be totally lost.
 - If Lastco had severe financial troubles, it would lose its credit but only after purchases of \$2,000, including an unpaid balance of \$1,000, of which \$500 would ultimately be collected.
 - As a sporadic customer, Lastco would average purchases of only \$500 (with no credit losses). However, as a good

customer, it would average purchases of approximately \$6,000.

IJK was concerned about granting credit to Lastco. On the one hand, if it did not extend credit to a potential customer, business was lost. On the other hand, there was a substantial risk of nonpayment (as described above), and since IJK made an average profit (price minus variable cost) of only 20 percent of sales, this exacerbated the problem. In addition, there were collection costs of \$100 per customer for those that failed or were in financial trouble.

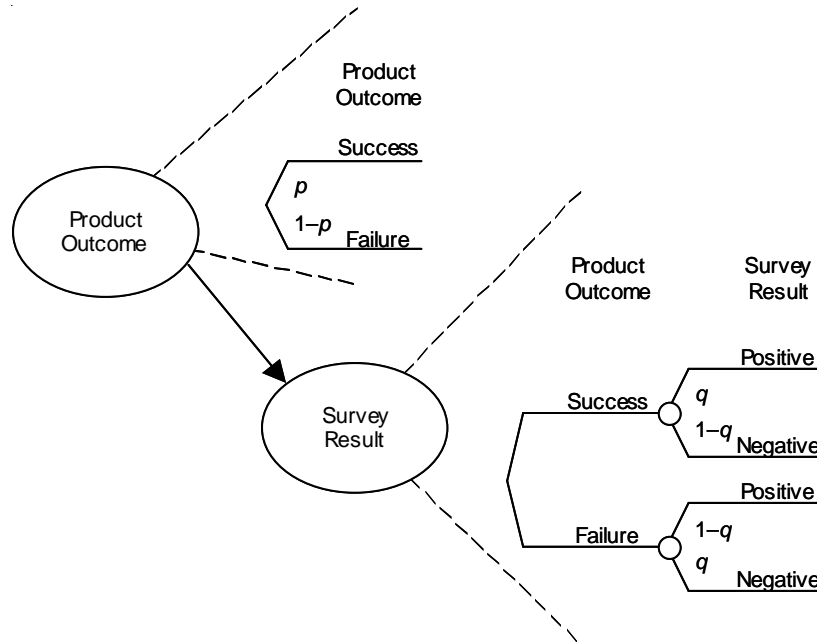
- a. Draw the influence diagram for this case.
- b. Construct the decision tree for this case.
- c. Should IJK grant credit to Lastco?
- d. Suppose a credit rating company could somehow provide perfect information on a potential customer for \$200. Should IJK buy it?
- e. Suppose the fee of the credit rating company was only \$50, but the company could provide only “good opinions” (not perfect information) about potential customers. Suppose also that Thompson has some experience with credit rating companies, which he says applies to the Lastco decision. His rating experience is summarized below as credit ratings by customer classification (percent of total.)

Rating	Failed	Financial Troubles	Sporadic Customer	Good Customer
Good	0	10	40	40
Medium	40	50	50	50
Poor	60	40	10	10
	100	100	100	100

Note: The table should be interpreted as follows: For example, in similar situations of companies that failed, none had been rated good, 40 percent had been rated medium, and 60 percent had been rated bad.

Would it be worthwhile to use the credit rating company? Illustrate your answer with a revised influence diagram.

- 4.11 Most market surveys give imperfect information. The example below shows a symmetric situation—a fraction, q , of product successes had positive survey results, and a fraction, q , of failures had negative survey results. The prior probability of a product success is given as p .



- Flip the tree and calculate the posterior probabilities for Product Outcome and the preposterior probabilities for Survey Result in terms of p and q .
- Plot the probability of a product success given a positive survey result against the value of p . Do this for $q = .5, .6, .7, .8, .9$, and 1.0 . Why are values of q less than $.5$ and greater than 1 not needed?
- How useful is the survey result for values of p near $0, .5$, and 1.0 ? How does this depend on the value of q ? Explain qualitatively what this means in terms of uncertainty and certainty and the accuracy of surveys.

5

Attitudes Toward Risk Taking

The Inadequacy of Expected Values

As noted previously, the expected value is not always an adequate decision criterion, especially for decisions that involve uncertainty and that are truly important to people or companies. Usually (but not always) people shy away from risk and uncertainty.


For example, in seminars to upper-level executives from many of the country's largest companies, participants were asked to bid for the rights to toss a coin where the prize was \$200. Because we are dealing with a coin toss, the expected value is, of course, \$100. It is surprising how few people bid anywhere close to \$100. Most bid in the \$40 to \$60 range—and these were not the people who objected in principle to gambling.

The exercise shows how easy (but often wrong) it is to think of expected value as an adequate decision criterion. When there is money on the table, we find we are anything but expected-value decision-makers. The seminar also revealed that few people behave consistently when making decisions involving risk. If they carried the attitude used in the coin toss over into other decisions, they would be very conservative investors indeed!

Most people and companies exhibit an aversion to risk taking. While they are willing to play the averages for small stakes, they are willing to give up a part of the expected value to avoid the uncertainty and risk for larger stakes. Thus, this attitude should be incorporated in the decision criterion. The effects of risk attitude can be quantified in terms of the certain equivalent, which, as we recall, is the certain amount the decision-maker would accept in exchange for the uncertain venture. This certain amount of cash is thus the venture's minimum selling price.

One way to see the effect of a risk-averse attitude is to look at the risk penalty (or risk premium)—the difference between the expected value and certain equivalent. As the amount at risk increases (e.g., as the prize in the

Figure 5-1



coin toss increases), the risk penalty increases and the certain equivalent becomes a smaller fraction of the expected value (Figure 5-1).

Furthermore, the decision criterion should have the certain equivalent decreasing the most (relative to the expected value) when the uncertainty is greatest. For instance, experience shows that if we hold the prize constant, but vary the probability (p) of winning the prize, we will see the greatest difference between the expected value and the certain equivalent in the region where the uncertainty is greatest (p around .5). (See Figure 5-2.)

Toward a Consistent Risk Attitude


The decision-maker could, in principle, stare at the probability distribution on value for each alternative of each venture, arrive at a set of certain equivalents, and then make the decision. Indeed, this is what we all do informally when faced with a decision under uncertainty. However, this approach has three problems. It is virtually impossible to be consistent from decision to decision, it is difficult to delegate effectively, and it is time-consuming and exhausting when many decisions must be made.

For these reasons, the decision criterion should enable the decision-maker to make explicit this most personal of values. The criterion should be applicable to all decisions, simple to apply in calculations, simple to communicate, and based on solid grounds—not just a rule of thumb of limited applicability.

Fortunately, we can take a series of certain equivalents the decision-maker has given for different ventures and produce a criterion for making other decisions consistently. The argument goes as follows:

- If we choose to behave according to a few reasonable behavioral rules, there is a utility function that describes our attitude toward risk taking.
- This utility function can be used in a simple way to obtain our certain equivalent for any probability distribution.
- The utility function is obtained by asking what our certain equivalent is for a few simple uncertain ventures.
- Over a reasonable range of outcomes, most people's and companies' utility functions for money can be fit by a function with one parameter, the risk tolerance. This makes it simple to speak of and compare attitudes toward risk.

Figure 5-2



The behavioral rules are shown below. We give a mathematical formulation for each rule and then explain how each relates to the way we actually behave—or wish to behave. Note that in the context below $>$ means preferred to and an arrow in the tree drawing indicates this preference.

Order Rule

If $A > B$ and $B > C$, then $A > C$.


The order rule states that if we prefer A to B and B to C, then we prefer A to C. For instance, if we prefer a Mercedes to a Ford and a Ford to a Volkswagen, then we prefer a Mercedes to a Volkswagen. If we were to violate

this rule, an unscrupulous car dealer could turn us into a “money pump”: have us trade in our Volkswagen and buy a Ford; trade in our Ford and buy a Mercedes; trade in our Mercedes and buy (at a higher price) a Volkswagen. In the end, we have returned to our original state, but are poorer. Objections to this rule usually err in focusing on only one attribute of the outcomes.

Equivalence Rule

If $A > B > C$, then there is some value, p , for which the decision-maker is indifferent between the alternatives in the tree shown in Figure 5–3. (The double arrow means indifference between the two alternatives.)

Figure 5–3



The equivalence rule says there is some value for the probability, p , at which we are indifferent between choosing a Ford and choosing the alternative where we might win a Mercedes or we might win a Volkswagen.

Substitution Rule

Given this indifference, B can be substituted for the uncertain venture in the top branch of the tree shown in Figure 5–3 without changing any preferences.

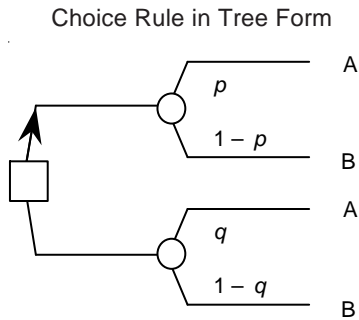
The substitution rule is really here for mathematical purposes. In behavioral terms it means, “Do you really mean it?” Are you willing to use the value p obtained in the preference statement of the equivalence rule as a probability in calculations? You must be able to substitute a certain venture for an uncertain venture with the same certain equivalent without changing any preferences.

Choice Rule

If $A > B$, then the choice shown in Figure 5–4 is true only if p is greater than q .

The choice rule says that we prefer the venture with the greater probability of winning the Mercedes.

Figure 5-4



Probability Rule

The decision-maker is indifferent between the two alternatives shown in Figure 5-5. (Note that the probabilities of outcomes A and B are the same for both alternatives.)


The probability rule is fundamental in that it says that we wish to act rationally and consistently, at least to the level of being willing to use probabilities in our decision-making.

This rule is sometimes called the “no fun in gambling” rule. It states that multiple ventures can be replaced by equivalent single ventures because the decision-maker finds no intrinsic value in the ventures themselves. Much of the “fun in gambling” occurs because multiple ventures spread the excitement out over time. If we can put this entertainment value of gambling into the value function, then even the gambler may be able to satisfy this rule.

We can prove that if we subscribe to these rules, a utility function exists that can be used to find our certain equivalent for any uncertain venture. We will postpone the proof for the existence of the utility function until the end of this chapter.

Figure 5-5

Probability Rule in Tree Form



For convenience, we will restrict the values to a single continuous variable (money), although the utility function can treat discrete variables and multiple attributes. As we discussed in Chapter 3, this restriction will cause no problem for most business decisions.

What Is a Utility Function?

A utility function is a means of describing how much a particular outcome is worth to you. For instance, winning \$1 million probably means something different to you than it does to the Sultan of Brunei (unless you are the Sultan of Brunei), and this difference would be reflected in your respective utility functions. A utility function measures worth by translating values (such as \$1 million) into a measure called utiles or utility values.

Your utility function has the following two properties that, as we will see in a moment, make it useful in establishing a decision criterion:

- You prefer the alternative with the largest expected utility.
- Because utility values have no intrinsic meaning, you can arbitrarily assign the utilities of any two outcomes. This allows you to set a scale that determines the utility value of any other number, much like the way water freezing at 0°C and boiling at 100°C sets the centigrade scale.


These two properties can be seen in the proof of the existence of a utility function in the section “Deriving the Existence of a Utility Function.”

Before it can be used, a utility function must first be created. For example, let us continue with the case of someone who is willing to bid a maximum of \$60 for the chance to flip a coin for a prize of \$200. Using tree notation, we can illustrate a very similar situation where the decision-maker is indifferent between accepting a sure \$60 and accepting the coin flip (Figure 5–6). Another way to state the situation would be to say that the decision-maker’s certain equivalent for the coin flip is \$60—sixty percent of the expected value of \$100.

We can write the utility of an outcome, x , as $u(x)$. Then the (expected) utility

Figure 5–6

Tree Showing Coin Flip Decision



of the second alternative is simply $u(\$60)$. The expected utility of the first alternative is computed in the same manner as the expected value: multiply the probabilities by the utility values and sum the results. Thus, the expected utility of the first alternative is $(.5 \times u(\$200)) + (.5 \times u(\$0))$.

If we always prefer the alternative with the largest expected utility, indifference means that the two alternatives have the same expected utility.

$$u(\$60) = (.5 \times u(\$200)) + (.5 \times u(\$0)) \tag{5-1}$$

Using the second property from above, we can arbitrarily assign two utility values.

$$u(\$0) = 0 \tag{5-2}$$

$$u(\$200) = 2 \tag{5-3}$$


Now we can calculate $u(\$60)$.

$$u(\$60) = (.5 \times 2) + (.5 \times 0) = 1 \tag{5-4}$$

We can then graph these three points and sketch in a smooth curve (Figure 5-7). (A procedure for assessing the curve more carefully is given in the section “Encoding a Utility Function.”)

How can we interpret this curve? In the positive direction, the curve flattens out so that, for instance, going from \$200 to \$220 produces less increase in utility value than going from \$0 to \$20. In other words, going from \$200 to \$220 means less to us than going from \$0 to \$20. In the negative direction, the curve becomes ever steeper, indicating that addi-

Figure 5-7



tional decreases in value become ever more significant.

Utility curves have some general properties that help in understanding them.

- Utility curves should generally be monotonically increasing. A utility function that sometimes decreases would mean that getting more of what you want would sometimes be undesirable!
- It is the curvature of the utility function that is important. The shape of the curve in Figure 5–7 (concave downward) indicates that the certain equivalent is less than the expected value. Upside potential means less (the curve flattens in the positive direction) than downside risk (the curve grows steeper in the negative direction). Having a certain equivalent smaller than the expected value is characteristic of risk-averse behavior. The more risk averse you are, the more curved (concave downward) your utility function is.
- The reader can easily verify that risk-neutral behavior (with the certain equivalent equal to the expected value) is described by a straight-line utility function. Utility curves for expected-value decision-makers are straight lines.
- Risk-seeking behavior is possible. Risk seeking implies certain equivalents greater than expected values. Utility curves for risk-seeking individuals would be concave upward. Significant risk seeking (not just buying a state lottery ticket) is ordinarily found in individuals whose aspirations are blocked by lack of fair markets. For instance, a man who needs \$100,000 to start a small business and cannot find any backing may be willing to risk all his \$50,000 capital in a desperate gamble whose net expected value is much less than \$50,000, but which has a possible outcome of \$100,000. Clearly, this type of behavior is very special and will not be considered further in this book.

Using a Utility Function

How do we use the utility curve? Suppose, for instance, that the person whose utility curve we plotted above is faced with a venture with smaller stakes, such as a coin flip for only \$100. We can use the utility curve to determine this person's certain equivalent (maximum bid) for the new coin flip. As in equation 5–1, the utility of the maximum acceptable bid is equal to the expected utility of the coin flip. We could draw a tree similar to the one for the first flip, but the following steps more directly accomplish the same calculation (finding the certain equivalent):

1. Use the utility curve to find the utility values of the outcomes.
2. Calculate the expected utility.
3. Use the utility curve to find the certain equivalent value corresponding to the expected utility.

This calculation is illustrated in the tree in Figure 5–8 describing a coin flip for a \$100 prize. The values used in this tree are obtained from the utility

Figure 5-8





Figure 5-9

Reading Values from a Utility Curve



curve shown in 5-9.

Note that with an expected value of \$50, the certain equivalent is now 80 percent of the expected value, whereas in the first coin flip the \$60 certain equivalent was only 60 percent of the \$100 expected value. With smaller amounts at stake, the certain equivalent becomes a higher percentage of the expected value. Another result the reader can verify is that the certain equivalent value becomes equal to the expected value as the probability of winning (or losing) approaches 1.0.


After some discussion, the president of Positronics consented to an interview in private by the decision facilitators. In the interview, the facilitators assessed a utility curve for the president. Reluctantly, he allowed his utility curve (Figure 5-10) to be shown the next day at the analysis group meeting that

had by now begun to be scheduled on a regular basis. At first, he was a little hesitant to allow people to see what his attitude to risk was. The facilitators also presented a few examples to show some implications of the utility curve. Upper level management was a bit surprised that the examples showed the president willing to take much greater risks than they had thought he would. They had previously discarded ventures as too risky that they now perceived might be interesting to him. Some rather frank discussion revealed that although the president was willing to take risk, he frequently rejected ventures not because of the risk but because he thought the optimists who presented the decision to him had overstated the probability of success.

By examining the curve (Figure 5-10), we can see that, like most people, the president is risk averse. As mentioned above, this means that his utility values drop off ever more sharply for large negative numbers and flatten out for large positive numbers.

We saw above how to use a utility function in a simple situation. A similar procedure is used for more complicated situations (and more complicated trees). As the proof in the section “Deriving the Existence of a Utility Function”

Figure 5-10



shows, the procedure for using the utility function starts by first converting all the values at the endpoints of the tree into utiles. Then, in the rollback procedure, chance nodes are replaced by their expected utility (instead of expected value) and decision nodes by the alternative with the greatest (expected) utility. The utility curve can then be used in reverse to convert the expected utility into the certain equivalent of the overall tree, or even to convert the expected utility into the certain equivalent at any point in the tree.

Positronics applied this procedure to the decision tree (Figure 5-11). The utility values at the end of the tree were calculated and rolled back through the tree to obtain the expected utility at each node.

The choice is still to bid \$500,000, with a utility of .088. Although it is difficult to tell from the encoded utility function (one reason why we will go to a computerized function), we can estimate that the certain equivalent corresponding to the utility is quite close to the expected value of \$65,000.

Value of Information with an Encoded Utility Function

One thing to remember when using a directly encoded utility function is to be particularly careful about doing value of information calculations. To calculate the value of information, you should use the full value of information tree (Figure 3-17) rather than just reorder the nodes.

To find the value of information, put in different information costs as (negative) rewards on the “Yes” branch until you have found a cost level for which the value of the two alternatives is equal. This cost is the value of information. This iterative procedure is necessary because different costs shift us to different regions of the utility curve where the utility/value trade-off may be different.

However, when using an exponential utility function such as the one discussed below, we can calculate the value of information as we did for expected-value decision-makers—there is no need for this iterative procedure.

An Exponential Utility Function

Within a reasonable range of values, many personal and corporate utility curves can be approximated well by an exponential function:

$$u(x) = a - be^{-x/R} \quad (5-5)$$

where x is the value (such as dollars), R is the risk tolerance, and a and b are parameters set by the choice of two points in the utility curve. The parameter b must be greater than zero. The risk tolerance is the parameter that describes the curvature of the utility curve and is expressed in the same units as the values. One sometimes sees reference to the risk-aversion coefficient, which is defined as $1/R$.

The facilitator graphed exponential utility functions for several different

Figure 5-11

Using the Utility Function in the Positronics Tree





Figure 5-12



values of risk tolerance and compared these functions with the curve encoded from the president of Positronics.* As he saw in Figure 5-12, the utility function encoded from the president was well described by an exponential utility function with a risk tolerance of about \$1 million. Further, within the range over which the curve was encoded, he saw that Positronics was not far from an expected-value decision-maker (which would mean a straight-line utility function and imply infinite risk tolerance).

The facilitator decided to use the exponential utility function with a risk tolerance of \$1 million. Tree calculations could be done more quickly and consistently using this function rather than the encoded curve. In addition, the facilitator decided to show certain equivalents rather than expected utilities in the tree drawing (Figure 5-13).

As the facilitator saw from the tree drawing, the certain equivalents for each alternative were quite close to the expected values. The difference is small because none of the possible outcomes in this decision problem posed a significant risk to the business. The risk premium was only \$65,000 – \$57,600


* The two arbitrary points in the encoded curve were $u(\$0) = 0$ and $u(\$1 \text{ million}) = 1$. The exponential utility function that goes through these points is

$$u(x) = (1 - \exp(-x/R)) / (1 - \exp(-1,000/R))$$

where R is the risk tolerance and x and R are expressed in units of thousands of dollars.

Figure 5-13

Positronics' Decision Tree with a Risk Tolerance of One Million Dollars




= \$7,400.

There are several quick ways of estimating risk tolerance. The proofs of these estimating techniques are in problems 5.14 and 5.15. First, you can consider the uncertain venture illustrated by the chance node in Figure 5-14. To make it more concrete, imagine that an acquaintance needs backing to start a small business and offers to “double your money in a short period.” You judge there are 3 chances in 4 he will be able to pay off.

The question, then, is what is the largest sum of money you would be willing to lend your friend? In other words, what is the largest investment, R , you would consider making in a venture that has a .75 chance of paying

Figure 5-14

Use of a Venture to Estimate Risk Tolerance




$2R$ in return and a .25 chance of paying nothing in return?

An equivalent way of posing this question is shown below. Suppose you had an opportunity to take the chance illustrated in Figure 5-15 or pass. What is the largest value R for which you will accept a .75 chance of winding up R richer and a .25 chance of winding up R poorer? (This question just combines the investment and return from the Figure 5-14.) R is approximate-

Figure 5-15


Would You Take the Chance or Pass?



ly your risk tolerance and should be the same in both cases.

You can also use another venture (Figure 5-16) to estimate your risk tolerance without having to consider the possibility of such severe losses. What is the largest value of R for which you will accept a .5 chance of winding up R richer and a .5 chance of winding up $R/2$ poorer? Once again, R is

Figure 5-16



approximately your risk tolerance.

Value of Corporate Risk Tolerance?

In the relatively rare event of assessing corporate risk tolerances from top management in the corporation, we have used the encoding methods described at the end of this chapter rather than the quick estimates discussed above. However, there are several rules of thumb for estimating risk tolerance. One rule of thumb is to set risk tolerance about equal to net income for companies that generally take moderate risks. Another rule of thumb is to set the risk tolerance equal to one-sixth of equity, or to one-fifth of market value. Compare these rules to the figures in the following table, which shows the risk tolerance obtained in the 1970s from the top management in three real companies.

Measure (\$ millions)	Company		
	A	B	C
Net Sales	2,300	16,000	31,000
Net Income	120	700	1,900
Equity	1,000	6,500	12,000
Market Value	940	4,600	9,900
Risk Tolerance	150	1,000	2,000

These three companies are large capital-intensive oil and chemical companies. Individual companies and different types of companies vary widely in the ratio of risk tolerance to net income, company size, stockholders' equity,

*There is some descriptive evidence that divisions within corporations act as if they had even smaller risk tolerances (are more risk-averse) than described here. See, for instance, "Risk Propensity and Firm Performance: A Study of the Petroleum Exploration Industry", M. R. Walls and J. S. Dyer, *Management Science*, 42, 7 (1966), 1004-1021.

or other measures. One would expect the ratio of risk tolerance to equity or market value to translate best between companies in different industries.*

However, there is a basic question whether publicly traded corporations should have a utility function and exhibit risk aversion. Standard corporate finance argues that corporate officers should maximize share price in their decision-making. Given that the shareholder has diversified holdings, the shareholder sees only the overall market risk and not project risk. Following this logic, projects should be evaluated at their expected value with market risk incorporated in the cost of equity.

How can we reconcile this conclusion with the fact that corporations appear to act in a risk-averse way? One suggestion is that there are hidden costs created by uncertainty—hidden in the sense that they are usually not completely modeled in project analysis. A corporation is the nexus for a set of contractual relationships among many individuals: shareholders, directors and officers, debt holders, employees, suppliers, distributors, customers, and government. Because many of these individuals are exposed to project risk (e.g., officers, employees, suppliers, and distributors), a risk premium may be required in executing risky projects; these risk premiums can effectively create a corporate utility function.*

In practice, most decisions do not pose serious “risk” to the corporation. Also, in practice, most corporations do exhibit risk aversion when faced with “risky” ventures. Clarity in corporate decision making is usually achieved by displaying probability distributions for their expected values.

An Approximation to the Certain Equivalent

Once we have assessed a risk tolerance, the following approximate formula is a quick way to approximate the certain equivalent for a given probability distribution:

$$\text{Certain Equivalent} = \text{Expected Value} - \frac{1}{2} \frac{\text{Variance}}{\text{Risk Tolerance}} \quad (5-6)$$

This approximation comes from the Taylor series expansion of the utility function and of the probability distribution (problem 5.22). The variance characterizes the width (uncertainty) of the probability distribution, and the risk tolerance characterizes the risk attitude embodied in the utility function. The second term on the right is just the risk penalty (Figures 5–1 and 5–2.) To estimate the variance, take the width from the .1 to the .9 points on the cumulative probability curve, square it, and then divide by 6.6. This works because that width is 2.56 times the standard deviation for a normal (Gaussian) distribution. (This value can be seen in the graph of the normal distribution in Figure 10–17.) Since the variance is the square of the standard deviation, when we square the width, we get

* "The Corporate Contractual System and Normative Corporate Risk Attitude," James Eric Bickel, Ph. D. thesis, Engineering-Economic Systems, Stanford University, June 1999.


$(2.56)^2 = 6.6$ times the variance. For distributions that are not normal distributions, this is, of course, only an approximation of the variance.

What is a reasonable range of values for which the exponential function is adequate? Normally, a reasonable range is one for which none of the values (positive or negative) are much larger than the risk tolerance. There are a number of other functional forms that fit well to utility functions with a wider range of values but that are not as convenient to work with. Furthermore, our experience shows that probability distributions should be explicitly examined by the decision-maker when they involve a range of values so great that the exponential utility function is inadequate; decision-makers will not rely solely on the decision implied by a utility function when so much is at stake.

The president of Positronics had tried all these techniques and obtained values for the company risk tolerance that, although well within an order of magnitude of each other, varied considerably. Part of the problem was that he had never worked with an explicit representation of uncertainty, so he was on unfamiliar ground. Furthermore, risk attitude is difficult to express quantitatively because it involves very personal values. To help him with his decision, the president asked the facilitators to evaluate the tree for a number of values of risk tolerance and see if it changed the preferred decision.

The tree was reevaluated using a number of different values for risk tolerance, and the results are shown in Figure 5-17. The horizontal axis is linear in $1/\text{risk tolerance}$, which is the risk-aversion coefficient. The leftmost

Figure 5-17



values are for infinite risk tolerance, which is equivalent to the expected value. The certain equivalents are plotted for each alternative for each value of the risk tolerance.


We see that the \$500,000 bid is optimal for risk tolerances down to around \$120,000. Below that value, the \$700,000 bid becomes preferable.

The risk attitude sensitivity is a very powerful tool. Frequently, decision-makers do not feel they have been consistent in encoding their risk tolerance and, thus, would not want an important decision to hang on such a difficult number. Also, in large companies, it is often difficult to interview decision-makers long enough to encode a utility curve. One approach is to estimate a risk tolerance for the company from the ratios above and then use the sensitivity to find out if a more exact determination of the risk tolerance is called for. The risk tolerance sensitivity frequently shows that, for any reasonable range of risk tolerances, the decision is clear.

Encoding a Utility Function

Encoding a utility function is not particularly complicated for a single continuous attribute such as money. First, arbitrarily set two points on the utility scale. A common practice is to set the utility of \$0 value to be 0: $u(\$0) = 0$. Pick another point (say \$10,000 for a personal utility curve) and let it

Figure 5-18



Simple Two-Branch Venture

be 1: $u(\$10,000) = 1$. Next, construct a simple two-branch venture with these two values as outcomes (Figure 5-18.)

Choose a probability p and ask the decision-maker what his or her certain equivalent is. Suppose the probability p were .7 and the decision-maker responded with a certain equivalent of \$5,000. We would then be able to plot a third point on the utility curve by using the basic utility property: indifference between the sure \$5,000 alternative and the uncertain alternative means their expected utilities are equal.

$$u(\$5,000) = [.7 \times u(\$10,000)] + [.3 \times u(\$0)]$$

$$u(\$5,000) = (.7 \times 1) + (.3 \times 0)$$

$$u(\$5,000) = .7 \tag{5-7}$$

We can now either change the probability, p , or construct a new simple venture from the values for which we have utility values. We then request that the decision-maker furnish a certain equivalent for this venture, which is used to find the utility point of still another value.

The process of encoding the utility points is thus fairly straightforward. However, the process of giving the certain equivalents is very unfamiliar and difficult. The first try at a utility function may produce a set of points that do not fit on a smooth curve.

The person encoding should be alert for discontinuities in slope occurring at \$0. People frequently get unduly alarmed by numbers with a negative sign, no matter how small. One technique to iron out this problem is to ask the decision-maker to consolidate this venture with another (certain) venture. This does not change the problem, but it does shift the place where the zero occurs. By doing this several times, you should obtain smooth continuity through the zero point.

In corporate decision problems, the utility function should be encoded from the top management of the company—often the president or CEO. However, since top management time is valuable, it is often best to estimate an appropriate risk tolerance for the company, perhaps using the ratios mentioned earlier in this chapter. You can then find out if the problem requires going beyond the expected value and risk sensitivity. For many business decisions, expected value turns out to be adequate. Only the largest and most uncertain problems require careful use of the utility function.


Deriving the Existence of a Utility Function _____

The behavioral rules discussed in the section “Toward a Consistent Risk Attitude” can be used to show that a utility function exists that has the desired property: decisions are made to choose the alternative with the greatest expected utility. The following seven-step proof derives this property from the rules. To keep things simple, we will present the proof for a finite number of discrete outcomes. The outcomes need not all be monetary and can represent a mix of different values.

1. The order rule is used to order the outcomes as R_1, R_2, \dots, R_n , where $R_1 > R_2 > \dots > R_n$.
2. Any uncertain venture, P , can be reduced by the probability rule to the form shown in Figure 5-19.

Figure 5-19


Representation of Uncertain Venture P



3. The equivalence rule is used to find, for each R_i , a probability u_i such that R_i is equivalent to the uncertain venture shown in Figure 5-20.


Figure 5-20

Using the Equivalence Rule



- The substitution rule can be used to substitute the results of step 3 of the tree in step 2 (Figure 5-19) and obtain the tree in Figure 5-21.


Figure 5-21



- The probability rule can be used again to simplify this tree, as shown in Figure 5-22.

Figure 5-22

Probability Rule Used to Simplify the Tree



6. Another uncertain venture, Q , can be represented by the tree in Figure 5-23 (as P was in step 2), and, by the same procedure, put in the form shown in Figure 5-24.

Figure 5-23

Representation of Uncertain Venture Q





Figure 5-24

Revised Form



7. If P is preferred to Q , then the choice rule states that

$$\sum_{i=1}^n p_i u_i > \sum_{i=1}^n q_i u_i \tag{5-8}$$

Believe it or not, we have achieved the desired result. If we define u_i to be the utility of R_i , then we prefer the alternative with the larger expected utility. (Remember that p_i and q_i are the probabilities associated with the alternatives P and Q .) We have constructed a utility function with the desired property.

Incidentally, you can easily see that when the utility values are transformed by a linear transformation

$$u'_i = a + bu_i \tag{5-9}$$

(with $b > 0$), the decision criterion that P is preferred to Q if $\sum_i p_i u'_i > \sum_i q_i u'_i$ is still maintained. The two parameters a and b are set by arbitrarily choosing the utility of any two values (provided the utility of the preferred value is greater than the utility of the other value).

Risk-Free Discount Rates

The choice of a discount rate is important in almost any business decision. Making a proper choice involves both the finance department and the decision analysis team and is one of the thornier questions a facilitator faces. It is often difficult to find people who really understand the considerations in choosing a discount rate and who will commit themselves to a particular rate. In this section, we describe some of the considerations and pitfalls in choosing a discount rate.

In Chapter 3, we mentioned that we accounted for time in the value function by discounting future cash flows to obtain the net present value. In this chapter, we have shown how to account for risk. For this reason, a risk-free discount rate should be used in the evaluation. What is an appropriate value for a risk-free discount rate? One risk-free investment opportunity is U.S. Treasury bills. Examining interest rates on these bills over the last 50 years or so has shown that, in real terms (excluding inflation effects), these bills return less than 4 percent (Figure 5-25).

Note that if d is the real interest rate (excluding inflation effects), i is the inflation rate, and d^* is the nominal interest rate (including inflation effects), then $(1 + d^*) = (1 + i)(1 + d)$.

Figure 5-25

Real Average Annual Returns on Common Investments

Real Average* Annual Return (%)

Series	1926-98	1940-89	1960-89	1980-89
Common stocks	—	7.4	5.4	12.4
Large company stocks	7.9	—	—	—
Small company stocks	9.1	—	—	—
Long-term corporate bonds	2.6	0.4	2.0	7.9
U.S. Treasury bills	0.7	-0.2	1.4	3.8

* Geometric mean, net of inflation

Source: Roger G. Ibbotson and Rex A. Sinquefeld, *Stocks, Bonds, Bills, and Inflation*. 1990 Edition, Institute of Chartered Financial Analysts, Charlottesville, Virginia; *Stocks, Bonds, Bills, and Inflation, 1999 Yearbook*, Ibbotson Associates, Chicago, Illinois

However, the situation that most projects face in the corporate world is that funding can not be obtained at the U.S. Treasury bill rate. Rather, capital is available at the Weighted Average Cost of Capital (WACC) rate. For corporate decision-makers, this often represents the “risk-free” corporate time preference. Whether this is appropriate is a matter of ongoing debate.

In business evaluations, an internal rate of return (IRR) criterion is often used. To account for risk, a minimum acceptable value for the IRR is established. In Chapter 3, we showed why the IRR criterion is inappropriate for dealing with situations in which uncertainty is important.

Another commonly used decision criterion is to apply a risk-adjusted discount rate and then to check for a positive present value. This criterion addresses uncertainty and risk indirectly by adding several percent to the discount rate to account for risk. However, as shown below, this approach contains many potential pitfalls, since it mixes considerations of time and risk, which are not necessarily connected.


In contrast to the risk-adjusted discount rate approach, the certain equivalent accounts for risk explicitly in the risk penalty, and the discount rate separately accounts for time with a risk-free rate. With the certain equivalent approach, when the certain equivalent for any uncertain venture falls below zero, the venture becomes undesirable (Figure 5–26). This is determined by the risk tolerance, a number that is determined once and then used for all decisions.

With the risk-adjusted discount rate approach, on the other hand, the venture is desirable for all values of the discount rate for which the NPV of some base case is greater than 0. The adjustment to the discount rate, however, has to be estimated for each alternative, taking into account the uncertainty of each venture. The Capital Asset Pricing Model (CAPM) is one approach to determining the discount rate for risky ventures.

The basic problem is that there is no intrinsic relationship between time and risk; thus, why try to evaluate risk by a discount rate? Several simple cases where blindly applying risk-adjusted discount rates produces spurious results illustrate this problem, as shown below. The certain equivalent method discussed in this chapter always works.

1. For any ventures where the uncertainty is resolved almost immediately, virtually no discount rate is high enough for the risk-adjusted discount rate method to make a risky venture appear undesirable. Ventures of this type include coin tosses, foreign currency exchanges, options, and opportunities arising because of acquisitions and divestitures.
2. For ventures with long time horizons, risk-adjusted discount rates tend to discourage ventures with any amount of uncertainty. This is one of the reasons why some companies spend too little money on long-term activities like research and development. Ironically, much of the uncertainty in research and development is usually resolved in the short term with the technical success or failure of a project.


Figure 5-26



3. For ventures with unconventional cash flows (profits occurring before payment), the higher we adjust the discount rate, the better the risky venture looks. However, higher discount rates are intended to discourage risky ventures! Examples of such “fly now, pay later” ventures occur in litigation (settle now or wait to see if the courts will declare you liable), in long-term financing, and in nuclear power plant construction (initial investment followed by operating profit followed by decommissioning costs).
4. A last example shows how the certain equivalent approach is sensitive to the size of the amounts involved, while the risk-adjusted discount rate approach is not. Consider Venture A and Venture B; B is identical to A except that it has all A’s

Figure 5-27

Sensitivity of Certain Equivalent and Risk-Adjusted Discount Rate Approaches to Size of Amount at Risk



outcomes cut in half. B could be, for instance, a 50/50 joint venture of A. The certain equivalent approach shows that below some value of the risk tolerance, the larger venture A becomes too risky and B is preferred (Figure 5–27). (This is one of the reasons for engaging in joint ventures.) However, the risk-adjusted discount rate approach will always prefer the larger Venture A for values of the discount rate where the NPV is positive. Venture B is only preferred for discount rates for which the NPV of both A and B is negative—the “lesser of two evils.” This characteristic depends only on the time pattern of the cash flow and not at all on the size of the alternatives, even though A is obviously riskier than B because it is bigger. Different risk adjustments must be used for the discount rates for Ventures A and B to make the risk-adjusted discount rate approach work.

Summary

Expected values are often an inadequate criterion for decision-making. People often exhibit an aversion to risk that leads them to different choices than looking at expected values would indicate. This behavior can be accounted for by substituting certain equivalent values for expected values. Certain equivalent values can be assessed directly for each probability distribution, but this procedure is cumbersome and difficult to do consistently.

Certain equivalent values for a decision-maker can be calculated consistently and easily by using a utility function. If the decision-maker is willing to accept a reasonable set of behavioral rules, it can be shown that he or she has a utility function with the desired properties. The utility function can be used to translate a probability distribution on values into the certain equivalent value.

A utility function can be encoded directly. However, an exponential utility function often provides a good approximation of actual behavior for most individuals and corporations and is analytically much easier to use. An exponential utility function can be characterized by a single parameter, the risk tolerance.

Finally, we briefly discussed why it is better to account for risk by using a utility function and a risk-free discount rate rather than by using risk-adjusted discount rates.

Problems and Discussion Topics

- 5.1 Consider several decisions you have made, ranging from minor importance to major importance. Was there implicit or explicit risk aversion in the way you went about making these decisions? Do you think your risk attitude was consistent across these decisions? Give examples

and say why or why not.

- 5.2 Do the behavioral rules adequately describe the way you would like to make certain types of decisions? Are there cases that do not fit the way you would like to make a decision? If so, give an example.
- 5.3 What factors contribute to the difference between using the expected value and using a utility function with risk aversion? Under what circumstances would the expected value and the certain equivalent value be the same?
- 5.4 Suppose you have a certain equivalent, CE, for a venture with probabilities (p_1, p_2, \dots, p_n) and prizes (x_1, x_2, \dots, x_n) . The Delta Property states that if we add some arbitrary amount Δ to all the prizes such that the venture is now for prizes $(x_1 + \Delta, x_2 + \Delta, \dots, x_n + \Delta)$, then your certain equivalent for the new venture will be $CE + \Delta$. Furthermore, if you subscribe to the Delta Property, then your utility function is exponential or linear.

Describe a situation where the Delta Property would not apply to you.

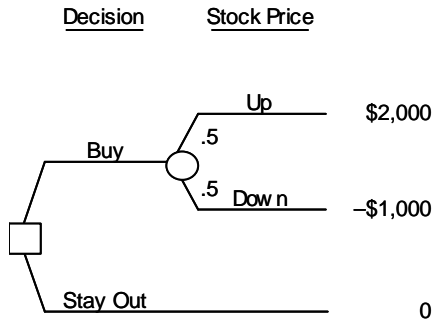
- 5.5 Use one of the two methods described in the text to assess a risk tolerance for yourself. Have the rewards or losses used in this assessment be paid immediately.

Now consider that the money you invest or lose can be paid monthly over a thirty-year period. For instance, at 10 percent you would pay roughly \$1,000 each month for the next thirty years to pay off \$100,000. Reassess your risk tolerance. Is it any different? Why or why not? What if interest were included on the balance?

- 5.6 Use the method described in the text to encode a utility function for a classmate. Then, directly assess his or her certain equivalent (minimum selling price) for a 1 in 5,000 chance of winning \$10,000. Compare the result with the certain equivalent from using the utility function. What does this tell you about your classmate's risk attitude for this kind of opportunity?
- 5.7 Peter Portfolio faces a decision for a short-term investment based on the prospective movement of a stock. The possibilities are shown below.

Peter has an exponential utility function with a risk tolerance of \$5,000.


- a. What is Peter's decision? What is his certain equivalent for the venture?
- b. Peter is not sure if his risk tolerance is exactly \$5,000. For what range of risk tolerance should he "Buy"?



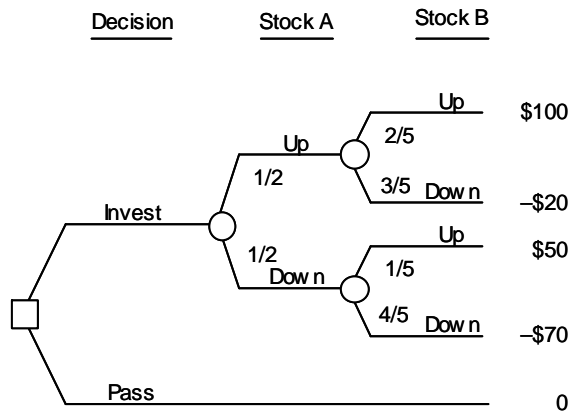
c. Sheba Sisters brokerage firm has investigated the stock and offers Peter perfect information on whether the stock price will go up or down. What is the maximum he should pay for the information?

- 5.8 Your regular morning radio show has awarded you the uncertain venture shown below.

You have an exponential utility function with a risk tolerance of \$2,000. You are indifferent to selling the venture for \$700. What is the probability p ?



- 5.9 J. K. Kay faces a short-term investment decision on stocks A and B whose performance is correlated.



- a. If JK is an expected-value decision-maker, how much should he pay the clairvoyant for perfect information on whether stock B goes up or down?
 - b. If JK is risk averse and has an exponential utility function with a risk tolerance of \$500, what is the most he should pay for the perfect information on stock B?
- 5.10 Compare the certain equivalents from using the exponential utility function for some problems in previous chapters with the certain equivalents obtained from the following approximation:

$$\text{Certain Equivalent} = \text{Expected Value} - \frac{1}{2} \frac{\text{Variance}}{\text{Risk Tolerance}}$$

where


$$\text{Expected Value} = \sum_{i=1}^m x_i p(x_i|S)$$

$$\text{Variance} = \sum_{i=1}^m (x_i - \text{Expected Value})^2 p(x_i|S)$$

- 5.11 Je has the following utility function:

$$u(x) = \ln(1 + x / R)$$

where $R = \$200$. His utility is set so that outcomes x are measured as differences from his present wealth. There is an uncertain venture, L, shown below.




Assume that Je can borrow money at no cost.

- a. Suppose that Je does not own the venture. What is the maximum he should be willing to pay for this venture?
- b. Assume that he owns the venture. What is the lowest price that he should be willing to sell it for?

- 5.12 There are extreme situations in which people may behave differently than normal. Some of these situations can be explained by a special utility function.

Mr. Sam Spade, after a night of partying in San Francisco, suddenly realizes that he has only a \$10 bill left in his pocket. Unfortunately, the train fare home for him is \$15. Then he observes a wild gambler in the nearby corner who offers him the three opportunities shown below, each at a certain cost:



Right now, all Sam cares about is getting home. What is his special utility function for this situation? Which opportunity should Sam choose?


- 5.13 Missouri Tubing, Inc., is a manufacturer of specialty steel and copper tubing. Its production facility is located adjacent to the Missouri River and is protected from the waters by a 20-foot-high dike. Extremely heavy rains in recent weeks have raised the level of the river dangerously close to the top of the dike. Several other areas near the river have already been badly flooded and the weather forecast is for continued rain.

The risk manager of Missouri Tubing estimates that there is a 20 percent chance that the river will top the dike in the coming weeks and flood the factory. If the factory is flooded, there is a 50 percent chance the damage will be heavy, costing about \$20 million to repair, and a 50 percent chance it will be light, costing about \$10 million. Furthermore, flooding of the factory will force it to shut down while repairs are made. If damage is heavy, the factory will be closed for four months. If damage is light, there is a 60 percent chance the shutdown will last four months and a 40 percent chance it will last only two months. For each month the factory is closed, Missouri Tubing will lose \$25 million in profits.


The risk manager now regrets that he recently cancelled the company's insurance policy covering flood damage. However, his insurance agent has decided to help him out by offering him a special emergency flood insurance policy. The policy provides coverage of both property damage and business interruption (i.e., lost profits) from flooding for a period of six months. The premium for the policy is \$30 million. Another policy providing only the business interruption coverage is also available for a premium of \$25 million.

- a. Draw the influence diagram for Missouri Tubing’s problem.
 - b. Structure the decision tree for Missouri Tubing’s problem and calculate the expected value of each alternative.
 - c. Calculate the value of clairvoyance for an expected-value decision-maker on whether or not the Missouri Tubing factory is flooded.
 - d. Calculate the certain equivalent for each alternative assuming that Missouri Tubing’s risk tolerance is \$50 million.
 - e. Calculate the value of clairvoyance with risk aversion on whether the Missouri Tubing factory is flooded.
- 5.14 Suppose that a decision-maker agrees on the five utility rules. We then know that there exists a utility function for the decision-maker. If he or she also agrees on the Delta Property, then we know further that the utility function is exponential and can be characterized by a single number, the risk tolerance. (See problems 5.4 and 5.17.)

One way to approximate the risk tolerance is to find the largest number X for which the venture shown below is still acceptable to the decision-maker. Show that X is within about 4 percent of the true risk tolerance.



- 5.15 Show that the largest value X for which the venture shown below is acceptable is within 10 percent of the risk tolerance (assuming an exponential utility function).



- 5.16 In general, the risk tolerance function can be defined as

$$\rho(x) = -u'(x) / u''(x)$$

where x is the increment to current wealth of the person in question (wealth = x + current wealth), u' is the first derivative of the utility function, and u'' is the second derivative. Some people agree that if

they had more wealth, they would probably tolerate more risks. To a first approximation, this can be captured by:

$$\rho(x) = \rho_0 + Kx \quad \text{where } K \geq 0,$$

which says that risk tolerance would increase if wealth increased. At the current level of wealth, $x = 0$ and the risk tolerance is just ρ_0 . What are the values of K and ρ_0 for the following utility functions?

- a. $u(x) = a - be^{-x/R}$
- b. $u(x) = a - \frac{b}{(1 + cx/R)^{(1-c)/c}}$
- c. $u(x) = a + b \ln(1 + x/R)$
- d. $u(x) = a + bx$ (expected - value decision - maker)

The utility functions in (b) and (c) go to negative infinity when $x = -\rho_0/R$. What does this mean behaviorally? Is it meaningful to use these utility functions for values $x \leq -\rho_0/R$?

- 5.17 If someone owns a venture, then the selling price for that venture should be such that the owner is indifferent between having the venture and having the selling price. The buying price should be such that the person is equally happy (has equal utilities) before and after giving up the buying price and getting the venture.
 - a. Show that for a particular person the selling price of a venture is, by definition, equal to the certain equivalent of that venture.
 - b. Show that if a person agrees to the Delta Property (which implies an exponential or linear utility function), then the selling price and buying price for a venture are the same for that person. (See problems 5.4 and 5.14 for a discussion of the Delta Property.) Does this result make sense to you?
- 5.18 A decision-maker is confronted with an uncertain venture, A , with outcomes x_i and associated probabilities $p(x_i | S)$.
 - a. Justify the statement that if the decision-maker already owns A , then the minimum selling price, S , of the venture is equal to the certain equivalent.
 - b. Justify the statement that the maximum buying price, B , of the venture is a number such that if the decision-maker wishes to acquire A , his certain equivalent for A with outcomes set to $x_i - B$ is zero.
 - c. Show that $S = B$ for a linear utility function $u(x) = a + bx$.

- d. Show that $S = B$ for an exponential utility function $u(x) = a - b e^{-x/R}$.
- e. Explain why if the stakes are large enough, the decision-maker can logically have $S > B$.
- 5.19 Value of information can be viewed as the largest amount the decision-maker would be willing to pay to get the information. It is, therefore, really the buying price of the information. Show that the value of information is given by

$$\text{Value of Information} = \text{Value with Information} \\ - \text{Value Without Information}$$

for a straight line or exponential utility function. Explain behaviorally why this is not generally the case.

- 5.20 Assume that you own two uncertain ventures, A and B , with outcomes x_i and y_j , respectively. There are no synergies between the ventures, so the net outcome to the company is the sum of the outcomes for each venture. The probability distributions for A and B are independent.

$$p(x_i, y_j | S) = p(x_i | S) p(y_j | S)$$

- a. For a straight-line utility function $u(x) = a + bx$, show that the certain equivalent of A and B together is the sum of the certain equivalents for A and B separately.
- b. Repeat (a) for an exponential utility function $u(x) = a - b e^{-x/R}$.
- c. Explain in behavioral terms why the certain equivalent for A and B together can logically be less than the certain equivalents for A and B separately (e.g., for cases when (a) or (b) above do not apply).

It is the property in (a) and (b) above that lets the decision facilitator work on a decision problem without too much concern for the resolution of other decision problems within the company.

- 5.21 You are given an uncertain venture with outcomes $x_i (i = 1, 2, \dots, n)$ and probabilities $p(x_i | S)$.
- a. Show that if the utility function is straight line $u(x) = a + bx$, then the certain equivalent of A is equal to its expected value.
- b. If the utility function is exponential $u(x) = a - b e^{-x/R}$, show that the certain equivalent of A is

$$CE(A) = -R \ln \left[\sum_{i=1}^n p(x_i | S) e^{-x_i/R} \right]$$

and is independent of the choice of values for a and b (provided that b is not equal to 0).

- 5.22 Given the results of the previous problem, prove that for an exponential utility function there is the following approximate expression for the certain equivalent.


$$\text{Certain Equivalent} = \text{Expected Value} - \frac{1}{2} \frac{\text{Variance}}{\text{Risk Tolerance}}$$

Hint: Use the following Taylor Series approximations, which are valid for small x .

$$e^x = 1 + x + x^2 / 2! + x^3 / 3! + \dots$$

$$\ln(1+x) = x - x^2 / 2 + x^3 / 3 - \dots$$

- 5.23 You have the venture shown below and your certain equivalent is zero.



- a. If your utility function is $u(x) = a - b w^x$, show that $w^x = p/(1-p)$. The value of $p/(1-p)$ is often called the odds.
- b. If your utility function is $u(x) = a - b e^{-x/R}$, show that your risk tolerance is $R = x / \ln(p/(1-p))$.

Part II

Dealing with Complex Problems

6

The Complexity of Real-World Problems

Up to now, we have dealt with problems that are relatively simple and easy to structure and analyze. Actual decision problems, however, are usually complicated and thorny. Furthermore, they never come to the decision facilitator in simple form. The professional decision facilitator sees only the problems that are complicated and often poorly described; frequently, it is unclear what the problem is and what decisions need to be made.

Fortunately, the techniques for dealing with simple problems are virtually the same as those for dealing with complex problems. The most important rule for dealing with both kinds of problems is “Keep it simple!” Unnecessary complexity makes the analysis more difficult without offering additional insight.

Throughout this chapter, we use an example to illustrate how decision analysis applies to the complicated problems (and short time frames) typically encountered in practice. This example is based on an actual analysis performed in a week. However, the example has been modified. Medequip and Diagnostics, the companies in the example, are fictitious. The products in the original analysis were *not* medical diagnostic equipment.* And the analysis itself has been modified and simplified.

Diagnostics was a relatively small company in the medical diagnostics equipment business. From the start, it had chosen to supply high quality products at a reasonable price to health care facilities such as hospitals. A well informed sales force and an excellent service and maintenance department had contributed to an excellent reputation for the limited number of products it offered.

Medequip was a large company offering a wide variety of equipment in the medical area. During a strategic review, Medequip's management realized that there was an area in which Medequip was weak: diagnostic equipment. And there was an obvious tie-in between diagnostic equipment and the rest of the

*The products in this example are not real products, and the input and results contained in this example do not describe real products.

Medequip line. With diagnostic equipment, Medequip would have the advantage of being a supplier of a larger fraction of the hospitals' equipment needs. And the attractiveness of Medequip's products could be enhanced by software and hardware links between the diagnostic equipment and other Medequip products.

For many reasons, Medequip chose to acquire a diagnostic equipment company rather than develop these products internally. They acquired Diagnostics to form the core of the new Diagnostics Division. Several months after the acquisition, Medequip's company-wide budgeting process had reached the point of requiring input for the following fiscal year. And the new Diagnostics Division realized that it had to solve a problem that had been postponed during the turbulent times before and after the acquisition.

Diagnostics had two instruments that were used to monitor a certain characteristic of a patient's blood. These instruments were designed to be kept at the patient's bedside and to provide instant readout to doctors and nurses. One instrument, DiagStatic, used a static electric field to separate elements in blood samples. The other instrument, DiagRF, used a radio-frequency electric field to perform a similar function. Some tests could be done by one instrument and not the other, some tests could be done by both.

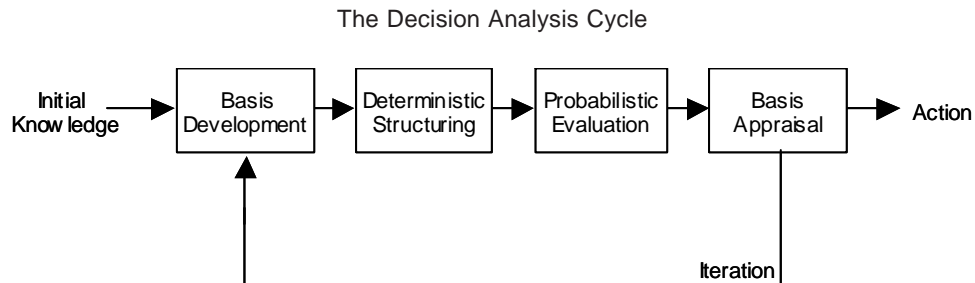
DiagStatic was a low margin, low market share product. It was an old product and both sales and margin were eroding rapidly; hospitals were switching over to a competitor's product at an alarming rate. If nothing were done, the product would be discontinued in five years. With a radically new design, new features would allow enhanced performance and also interconnection with other equipment.

DiagRF, on the other hand, was a high margin, high market share product in a smaller segment of the industry. It was based on a newer design and was selling reasonably well, but was also threatened by competition. An upgrade using the current design would enhance sales and prolong the life of the product.

At the time of purchase by Medequip, Diagnostics had the people and resources to develop only one new product at a time, and development was estimated to take two to three years. Which product development should Diagnostics choose to include in its budget request? Or should it choose to focus its budget requests in other areas? Or should it think big and ask for development for both products?

Budget submission was due in two weeks, but Laura, the head of Diagnostics Division (former CEO of Diagnostics), wanted results in one week so that consensus could be built during the following week. And amid all this turmoil, someone came up with a new alternative: a piece of equipment used in the operating room could be modified so that, when needed, it could fill the function of both DiagStatic and DiagRF at the patient's bedside. And then someone suggested that a small company had a design for a DiagStatic replacement that might be available for purchase. And then somebody mentioned that the consumables used by DiagStatic and DiagRF might be a profitable market for Diagnostics to enter—the current supplier was doing very well.

Figure 6-1



A Cyclical Approach

To deal with complex real-world problems, decision analysis uses a cyclical approach (Figure 6-1). The phases of this approach are illustrated in detail in this chapter. In this section, we give an overview of the process.

To start with, we bring some initial knowledge to a problem. Then, in the basis development phase, we gather data, generate alternatives, and identify values for making the decision. In the deterministic structuring phase, we build a model of the decision, develop base-case input to the model, and perform deterministic sensitivity analysis to find the crucial uncertainties in the problem. (Many uncertainties are relatively unimportant.) Then, in the probabilistic evaluation phase, we encode probability distributions for these variables, construct the decision tree, and evaluate the alternatives. Finally, in the basis appraisal phase, we examine the results obtained from the decision tree, revisit the information, alternatives, and values used in the evaluation, and reexamine the model. Based on this appraisal, we decide whether to act or to return to the beginning and start the cycle again.

Drawing on years of experience using the approach, we have developed several rules of thumb for how many times to go through the cycle and for how long to spend on each phase of the cycle. In general, there should be at least two to three iterations through the decision analysis cycle. Through each phase, the analysis should be successively developed and refined until it reflects the best judgment and expertise of the facilitator and personnel involved, keeping in mind that the level of effort should be commensurate with the size and importance of the problem.

In an analysis lasting several months, a preliminary iteration should take place in the first weeks (possibly even in the first few days) to give a feeling for the overall scope and structure of the analysis. Depending on the situation (and to a certain extent on the decision facilitator), another iteration might be performed before the full-scale analysis (when the major work is done). At the end, there should be time for a final iteration to recheck the problem formulation and results and make any final modifications.

Using the cyclical approach minimizes wasted time and effort. For instance, while the initial pass does not produce definitive results, it does give

us a sense of how much additional information will be required and of how to further structure the model. And making multiple passes minimizes the practical errors that often arise: treating the wrong problem, focusing on the wrong factors, using the wrong model, gathering the wrong information, and using the wrong people.

Most of the hard analytical and data-gathering work is done during the main pass through the cycle, which, of course, varies tremendously in terms of time spent. In a typical analysis, the main pass might consume 50 percent of the time, with 25 percent of the time spent in the initial passes and 25 percent spent in a final pass and in preparing the presentations. During the main pass through the cycle, the effort is usually split fairly evenly between constructing the model and gathering the data to use in it.

Typically, there is one presentation after the main deterministic structuring phase and one presentation at the end of the analysis. A presentation after the basis development phase, however, is often an excellent investment of time and resources. In recent years, the Dialog Decision Process (DDP) described in Chapter 8 has been used as the process of interaction with the decision-maker(s), and the decision analysis cycle has been used for the back-room analysis effort. In DDP terms, a presentation at the end of the basis development phase would be called a Framing and Alternatives Dialog, a presentation after the main deterministic structuring phase would be an Analysis Dialog, and the final presentation would be an Analysis and Decision dialog.

The total amount of time and effort spent in each of the phases varies with the type of problem. Basis development ranges from 10 percent of the effort in a well-defined problem to 50 percent for a strategy project involving many people and complex alternatives. Basis appraisal can vary from 5 percent of the effort for clear problems with clear solutions and a single decision-maker to 25 percent for problems involving many people and many divisions within a company. Finally, deterministic structuring usually takes about twice the effort required for probabilistic evaluation.

Laura, the head of Diagnostics Division, realized that she had to sort through all these possibilities within the week and arrive at a budget request that the division could live with. In addition, Medequip was a rather staid company, not used to taking gambles—and in the DiagStatic/DiagRF area, technology and competition were changing so rapidly that all courses of action (including doing nothing) were a gamble. She needed to make a choice and be able to justify it during this first budget review with the new parent company.

She therefore called in one of her staff who had done decision analyses in the past, explained to him the need for fast turnaround, authorized him to assemble a small team for two or three meetings during the coming week, and sat down with him to start working on the decision.

Basis Development

The basis development phase determines the scope, method of approach, requirements, and objectives of the analysis. This phase is especially critical because all the successive phases depend on it. When something really goes wrong in analyzing a problem, the roots of the difficulty usually lie in the problem structure that came out of the basis development. Similarly, when a problem is exceptionally well analyzed, it is usually because the analysis was well structured from the beginning.

The most common errors in structuring are analyzing the wrong problem or analyzing the correct problem in an overly complicated manner. While there is unfortunately no procedure to follow that guarantees success in structuring a problem—problem structuring is a skill learned by doing—most successful analyses have a number of common elements.

Starting the Process

The process starts with identifying a decision facilitator and establishing the roles of decision-maker, decision facilitator, project team, and subject matter key personnel/experts.

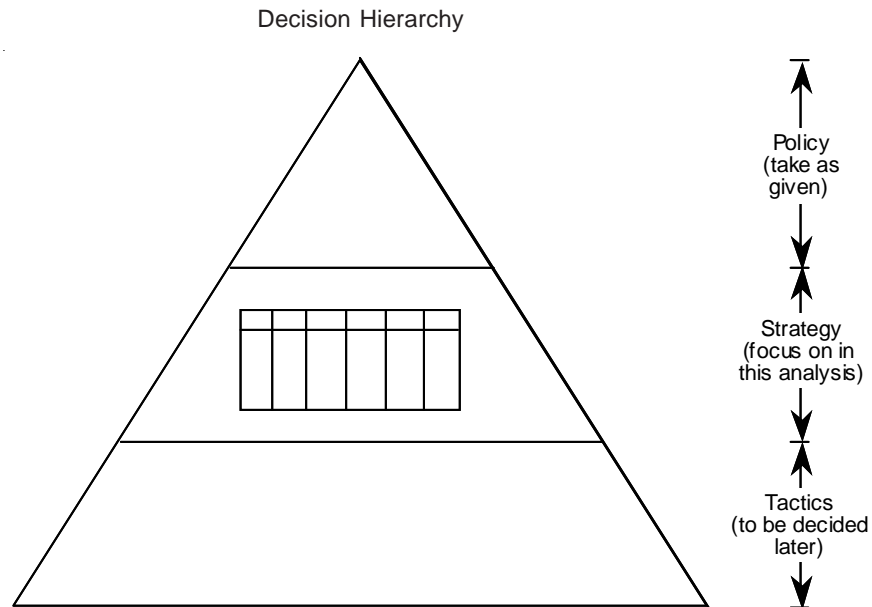
The decision-maker should almost certainly be involved at some point in the basis development phase, because the analysis will be much smoother if from the start it addresses at least the alternatives the decision-maker sees as viable, the uncertainties or worries he or she thinks critical, and any special values that may be important.

The decision facilitator must be given sufficient authority and expertise to direct the analysis effort, though he or she may be relatively unfamiliar with the problem at hand (as is usually the case when acting as a decision consultant). The decision facilitator must have the courage to ask questions about the obvious and the perspective to see the forest, not just the trees. Furthermore, the facilitator must have at least occasional access to the decision-maker. Finally, the decision facilitator must be (and be perceived to be) neutral concerning the choice between alternatives; if not, biases easily arise (or are thought to arise) in the analysis.

It is very difficult for a decision facilitator to work alone. A project team should be identified to assist him or her. The project team does not have to be as neutral as the decision facilitator, but the team should be balanced. How many team members? Enough to represent the principal stakeholders (e.g., design, manufacturing, marketing, sales), though one person can often represent different constituencies. How often need the team meet? Once or twice a week appears to be a natural frequency.

Key personnel and experts are the people who need to be consulted during the decision process. Some of these people have the knowledge that the decision-maker wishes to use in making the decision. Other have a crucial role in implementation and need to be consulted and kept informed.

Figure 6-2



Using Decision Hierarchy

The Decision Hierarchy (Figure 6-2) helps people sort through the different types of decisions. This hierarchy distinguishes between three types of decisions. *Policy* decisions are those decisions made at a “higher” level within the organization; more loosely, they are the “givens” of the problem. *Strategy* is the name given to the decisions that are under consideration in the analysis; the symbol in the middle section of the decision hierarchy stands for the strategy table dealt with below. *Tactics* are the decisions to be taken later, after the strategy has been chosen.

Initial meetings and interviews with key personnel should focus on identifying the key policy and strategy decision or decisions. All too frequently, the initial focus is on operational, tactical, or “how-to-do-it” decisions rather than on the underlying strategic or “what-should-we-be-doing” decisions. One way to identify the key decisions is to elicit a short description of the problem, such as, “Explain to an interested stranger what all the effort is about.”

When the principal policy and strategic decisions have been identified, a number of operational or tactical decisions can be dropped from the analysis. This must be done in a realistic manner. Usually, it should be assumed that these decisions will be made using the expertise available to the company. To avoid losing the goodwill of important personnel, the process of eliminating “how-to-do-it” decisions from the analysis must be done delicately.

Because this was the first time through the budget process with Medequip, Laura was not really sure what conditions corporate headquarters would

impose on the division. She and the facilitator put together the following list of what she thought were the policy decisions, the “givens” of the problem. This list would furnish a starting point for the analysis.

- *Diagnostics Division could not obtain more resources than they already had, so at most one new product could be developed.*
- *Purchase of the DiagStatic replacement from another company was not a possibility. Medequip was not in a position to purchase anything for some time, given its recent acquisition of Diagnostics.*
- *Diagnostics could not enter the market for consumables for these products.*

Laura promised to phone corporate headquarters and get some clarification of these policy statements within a few days. As it turned out, the list was substantially correct.

She was less definite about the division between strategy and tactics, but she stated that she wanted a clear statement of what needed to be developed. The team could work out the level of detail needed to achieve clarity in that decision

Using Strategy Tables

In the decision hierarchy, the policy decisions and the strategy decision areas were identified. Now the decision facilitator needs to develop several strategic alternatives that are significantly different.

The first problem that often arises in finding strategy alternatives is tunnel vision, where companies evaluate only a few fundamentally similar alternatives. By using processes ranging from simple heuristics to extensive group exercises, companies can stimulate the creativity needed to generate these alternatives. One simple exercise that often elicits creative responses is to imagine we are looking back from some future point (perhaps retirement) and critically examining this portion of our life.

The second problem that arises in finding strategy alternatives is often the multiplicity of alternatives. In many problems, there is a bewildering number of alternatives simply because choices must be made in several decision areas, and each decision area has several possibilities to choose from—the number of possible combinations increases rapidly with the number of decision areas. Creativity exercises may have dramatically enlarged the number of possible alternatives. The challenge now is to pick several alternatives (three to five alternatives, so that the analysis is feasible) that are significantly different (so that the analysis is creative). The decision analysis cycle will allow subsequent refinement of alternatives that appear most attractive.

The strategy table is a convenient and powerful technique for dealing with alternatives. The strategy table consists of a column for each decision area, with all the possible choices for that decision area listed in the column. A strategy alternative (a complete set of decisions) for the problem can then be specified by making a choice for each decision area.

In practice, forming a strategy table is often a group exercise. Different people contribute the decision areas they feel are important and then help to list all the choices in that decision's column. Initially, all the decision areas relevant to the problem are included, whether strategic or operational, important or unimportant. The group then discusses which decisions clearly do not belong in the strategy table.

Next, the group identifies a set of strategy alternatives to be developed. One way to begin developing an alternative is to pick a strategy theme or to make a choice in a key decision area. Then the alternative is developed by making one choice in each decision column. One of the alternatives should describe how things would have happened on their own—this might be the “momentum” or “do nothing” alternative, depending on the context.

Tables for major strategic studies can start with as many as 40 or 50 columns. For tables this large, explicitly identifying the elements of major alternatives often shows that what at first appeared to be a good alternative is a poor or unfeasible one. In addition, some decision areas may be perceived to be mainly tactical or operational. Finally, some decision areas may have choices dictated by the choice made for a more important strategic decision area. Through these kinds of considerations, the complexity of even large strategy tables can be reduced to a conceptually manageable level.

Although it is not a quantitative tool, the strategy table is a surprisingly effective way to apply intuition and experience to a highly complex situation. When used correctly, the table can also be used to eliminate most of the possible strategies that are inconsistent, inferior, or undesirable.

The decision facilitator called a meeting of four people who were responsible for most aspects of the products under discussion. Two were development engineers, one was in charge of manufacturing, and one doubled as the head of sales and of marketing.

The facilitator quickly described the problem, the time frame, and the desires of the division head. He then stated the purpose of the meeting was to identify the key decisions that would have to be made in choosing a course of action. As the facilitator listened, he noted all the decision areas and the options in these areas. The team then organized these into a simple strategy table (Figure 6-3) with four decision areas:

- *If they chose a new DiagStatic, should they concentrate on product features for competitive advantage or should they work on adding links to other Medequip devices?*
- *If they chose a new DiagStatic, should they integrate graphic output capabilities or should they continue with the current OEM add-ons?*
- *What target average selling price should be used in designing the devices?*
- *What about that operating room device? Should they design in Static and RF capability so that it could also be used at bedside?*

At this point, they called Laura in and asked her if the strategy table fairly represented the choices the division faced. She agreed that for this quick

Figure 6–3

Initial Diagnostics' Strategy Table

DiagRF		DiagStatic				Other
New DiagRF	Target Selling Price	New DiagStatic	Competitive Advantage	Output Graphics	Target Selling Price	Operating Room Device
None	\$10,000	None	Features	OEM add-on	\$10,000	No change
Upgrade, current design	\$15,000	Upgrade, current design	Features + links to other equipment	Proprietary add-on	\$15,000	Add Stat+RF capability for bedside use
	\$20,000			Integral, dedicated	\$20,000	
	\$25,000	New design		\$25,000		

analysis, they had the right level of decision in the table. She thought that the operating unit device decision was not really relevant—operating unit devices were expensive and dedicated and would never be available for bedside use. But she thought that the column should be left in the table so that the issue could be discussed before being dismissed.

Using this table, Laura and the group discussed the different possible alternatives. The problem was simple enough that the group readily came up with three alternatives (Figure 6–4):

- *Milk Existing Products: Keep both DiagStatic and DiagRF on the market as long as feasible and then discontinue them.*
- *DiagRF Plus: Upgrade the current DiagRF to DiagRF Plus using current product technology and design. Keep DiagStatic on the market in its current state.*
- *New DiagStatic: Develop a completely new design for DiagStatic, DiagStatic New, including links to a number of Medequip devices. Keep DiagRF on the market for as long as feasible, and then abandon the static segment of the market.*

Everyone thought that, once they understood the differences between these three alternatives, they would be able to refine the best alternative and make a choice.

Using Influence Diagrams

As soon as possible, the decision facilitator should develop a list of the uncertainties that will probably be important. Although this list will be revised during the analysis, it lays the groundwork for developing a deterministic model. The model will need to contain as explicit variables the major uncertainties identified and should be suitable for analyzing the alternatives that have been developed.

Earlier, we developed the basic tools and theory of influence diagrams. In this section, we focus on using influence diagrams in fairly complicated

Figure 6-4

Diagnostics' Alternatives

Strategy Theme	DiagRF		DiagStatic			Other	
	New DiagRF	Target Selling Price	New DiagStatic	Competitive Advantage	Output Graphics	Target Selling Price	Operating Room Device
Milk existing products	None	\$10,000	None	Features	OEM add-on	\$10,000	No change
DiagRF Plus	Upgrade, current design	\$15,000	Upgrade, current design	Features + links to other equipment	Proprietary add-on	\$15,000	Add Stat+RF capability for bedside use
DiagStatic New		\$20,000				New design	
	\$25,000	\$25,000					

analyses involving many people in an organization and, later, on how deterministic sensitivity analysis identifies the most crucial of these uncertainties.

There are two ways to begin an introductory session using influence diagrams. One way is to start from a simple influence diagram that the facilitator is fairly sure represents the “backbone” of the problem. This has the advantage of cutting short the initial phases of the discussion and the disadvantage of not allowing the session participants’ understanding to grow while the diagram grows.

The opposite approach is to start with the value function on one side of the page and a decision node on the other and to gradually break the value function uncertainty into its components; near the end of the process, the decision variable becomes connected to the growing diagram. The advantages of this approach is threefold.

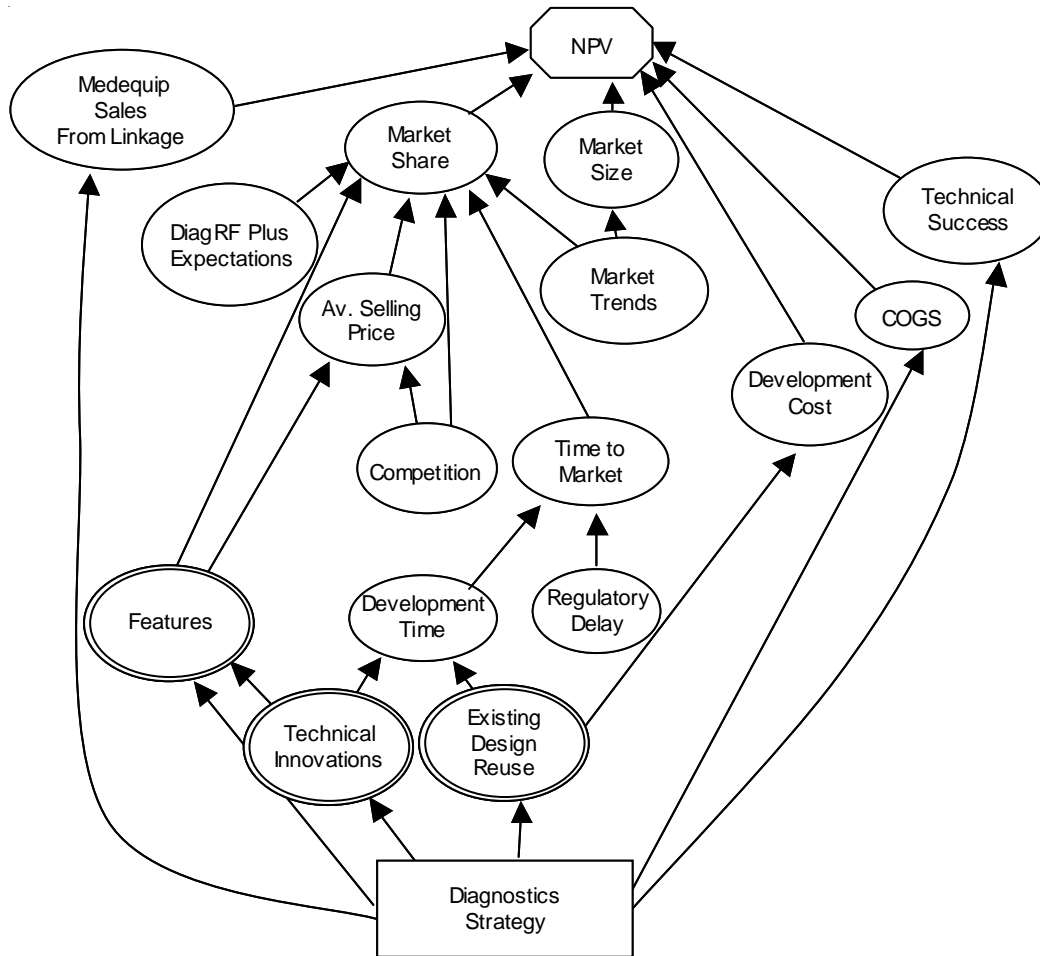
1. It focuses attention and demands agreement on the value function at the outset. It is surprising how much confusion in a decision problem arises because the decision criterion is not explicit and agreed on.
2. It focuses attention on the future and its attendant uncertainty rather than on the immediate concerns of selecting and implementing an alternative.
3. It allows all the participants to join in the thought process. This is valuable for obtaining enthusiastic implementation of the alternative that is finally chosen.

Most often, a modification of the second approach is used. The facilitator begins the process, creating several nodes. Fairly soon, the participants get the idea and begin to contribute. People who are not quantitative often find themselves contributing enthusiastically and productively at this stage of the analysis.

Occasionally, a tree is a more natural tool than an influence diagram for opening the structuring session. This can occur when the sequential nature

Figure 6-5

Influence Diagram for Diagnostics



of the problem is very important. For instance, some litigation and R&D problems are more readily described by trees in the early stages of structuring. (See Chapter 7.)

The group used an influence diagram to structure their problem (Figure 6-5.) They first created the topmost node NPV (net present value of cash flow) as the fundamental uncertainty. If they knew the probability distribution on NPV for each alternative, it would be easy to make the decision. Since this was the value function, it was drawn as an octagon.

Next, the facilitator asked what would the group would most like to know to reduce the uncertainty on NPV, given that they could not determine the

distribution on NPV for each alternative directly? The group immediately replied “Market Share.” The facilitator classified this response as an uncertainty, wrote it down in an oval, and drew an arrow from it to NPV. He mentioned to the group that the oval was really shorthand for two ovals: the Static and the RF market shares.

The facilitator next asked what the group would most like to know to reduce the uncertainty on market share. He got five immediate answers: Average Selling Price, the state of competition, the features in the devices, the date the new device arrived in the market, and market trends (RF might take some of the Static market).

The facilitator started a discussion around Average Selling Price. Target Selling Price appeared in the strategy table (Figure 6-3 and 6-4). After some discussion, some clarity was reached. In the development process, design and marketing reached agreement on the set of features to be included and a target selling price and a target cost to manufacture. When the device appears in the market several years later, the actual selling price is set by market considerations, and it is this actual selling price that is the uncertainty.

Some additional questioning revealed that Features and Competition both influenced Average Selling Price. Further questioning revealed that Features was not really an uncertainty but was determined by the strategy decision, so it was put in a double oval and an arrow drawn to it from the Diagnostics' Strategy decision rectangle. At this point the head of Marketing/Sales added an additional uncertainty called DiagRF Plus Expectations—the upgrade to DiagRF had already been announced, and he was uncertain whether the market would react unfavorably if the upgrade did not take place.

At this point the facilitator thought the Market Share uncertainty was quite well developed. So he asked what other uncertainty was important to NPV. Market Size, Development Cost, Cost of Goods Sold (COGS) were all identified and included in the diagram. One design engineer pointed out that the new design that would be attempted for DiagStatic was not that easy, and there was a possibility that, after some initial development effort, they would be forced to abandon the effort—this resulted in a oval labeled Technical Success. Finally, somebody suggested that if links were established to other Medequip equipment, there might be some additional sales of that Medequip equipment, so one final oval was added to the influence diagram.

Was the decision facilitator correct in allowing each oval to stand for two markets/devices? An alternative approach would be to duplicate each oval, once for Static and once for RF and to ask what arrows should be drawn from Static to RF bubbles. In reality, that is what is going to happen when the time comes to obtain input and assess probabilities. Given that discussions had indicated that interaction between the two markets/devices was weak, the influence diagram in Figure 6-5 is probably a reasonable communication tool. But the facilitator must remember to ask about interactions between Static and RF when obtaining data.

When do we stop adding to the diagram? When we reach a level of detail where we can use intuition and judgment to make meaningful assessments.

In practice, we sometimes add another level of decisions and uncertainties beyond this point, which are discarded later on as too detailed. Typically, an influence diagram for a complex diagram can grow to 50 or 60 decisions and uncertainties, can then be reduced to 10 or 20 by general considerations, and can be finally reduced to 5 to 10 by sensitivity analysis. (See “Sensitivity Analysis” below.)

The chief value of constructing the influence diagram is that it forces people to think hard about the decisions and uncertainties and their interrelationships. Specifically, people can contribute their thoughts and argue with others’ conceptions, areas of investigation can be identified and responsibility established, and probabilistic dependencies can be elicited and examined.

Deterministic Structuring: Modeling the Problem _____

Most corporate decision problems are complex enough to require using a model, usually implemented on a computer, to determine the value of each alternative. The models used today tend to be simpler than those used when decision analysis was first routinely applied to real-world problems. This change has come with the realization that the model does not need to depict the real world with absolute accuracy, but rather mainly needs to differentiate between the possible alternatives. However, in recent years, a rising level of sophistication and expectations appears to have resulted in larger, more complex models than were common a decade ago.

The basic rule of decision analysis is “keep the model as simple as possible—and no simpler.”

Preexistent corporate models are rarely useful for decision analyses. They tend to be too complex, take too long or are too complicated to run, and usually focus on quantities important to running the company but not really relevant to making major decisions. For this reason, decision facilitators almost always create a new “throwaway” model that simply and clearly describes the problem at hand. Once the decision is made, much of the model becomes irrelevant to future operations. What remains can sometimes be used as the core of a model for future business.

The primary criterion for the model is that it be accurate enough to distinguish between the alternatives. The large value swings caused by uncertainty will make the accuracy of the model relatively unimportant. The cyclical decision analysis approach enables us to successively refine the model as we go through iterations of the decision analysis cycle (Figure 6-1), which means the model does not have to be perfect the first time used.

One frequent comment is that “the model has to be complete and detailed enough to convince the decision-maker and others of its accuracy.” Such logic is poor justification for building an overly complex model. To the contrary, decision-makers are often suspicious that very complex models are black boxes hiding all sorts of errors or incorrect assumptions. A well-wrought model is seldom rejected for being too simple. Because they are

easier to understand, simple models also have the advantage of yielding insight into a problem that helps generate new alternatives.

Today's spreadsheet programs (e.g., Microsoft Excel) furnish the capability for simple logic and easy output, making it easy to understand and validate the model and to understand the results. Whatever the modeling language chosen, it should be accepted within the organization, familiar to those involved in the analysis, and compatible with the tools of decision analysis.

To analyze the options, the facilitator built a spreadsheet model. The model had to relate together all the quantities indicated on the influence diagram (as either data inputs or calculation results) and had to be able to value all the different alternatives indicated in the strategy table. Most of these relationships could be expressed as simple financial calculations.

A basic model for market share for the new products (DiagStat New or DiagRF Plus) is shown in Figure 6-6. A five parameter model of this type facilitates discussion and is surprisingly efficient in modeling many different situations.

After much discussion, the project team simplified the influence diagram to the one shown in Figure 6-7. Medequip sales linked to the new products were thought to be small and, given the time available, the team decided to neglect this added value; the new products would have to stand on their own. Other nodes were judged to be useful in thinking about variables, but were not used directly in the model.

Many of the nodes in the original influence diagram do not wind up as explicit variables in the model. Some nodes are eliminated at a qualitative level as being interesting, but not worth modeling because the effect is too small; "Medequip Sales from Linkage" is an example. Other nodes are not explicitly modeled, but encourage systematic thinking about the uncertainties that are modeled; "Competition" is an example. This latter type of node is sometimes called an "evocative node."

Figure 6-6

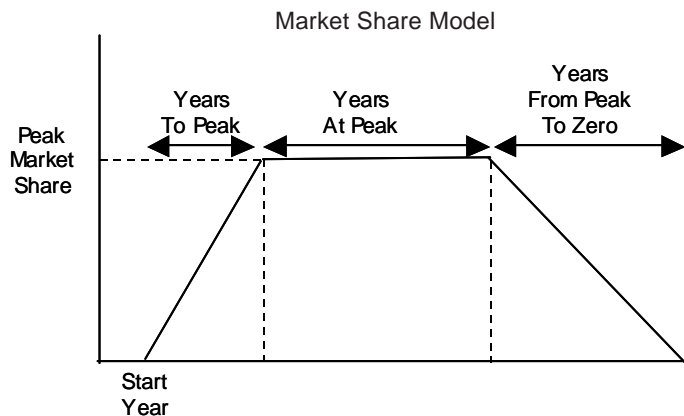
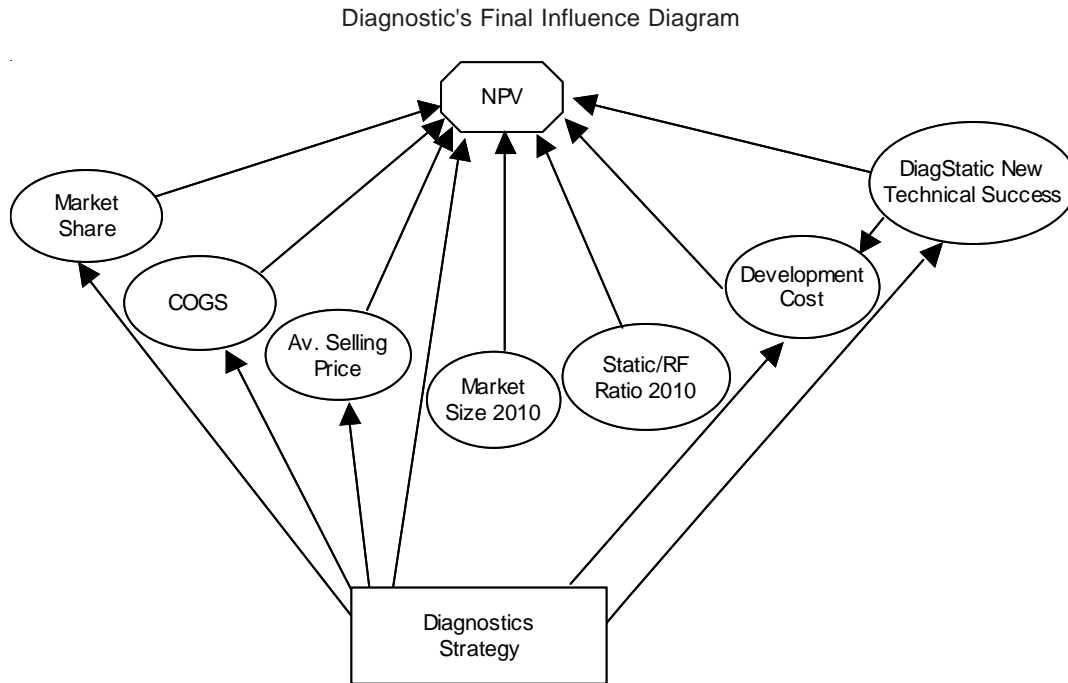


Figure 6-7



Note that the facilitator separated the spreadsheet into a data input section (Figure 6-8) and a calculation section (Figure 6-9.) The data input section contains all the input numbers for the spreadsheet for both constant and uncertain parameters. The input that is actually used in the calculations is contained in the shaded cells, mostly in the “In Use” column.

The calculation section contains only equations, though, of course, the results of these calculations appear as numbers. By setting up the spreadsheet this way, the facilitator has automatically parameterized the inputs for later analysis and made sure that he can easily scan down the complete list of data inputs.

Some programmers, especially those who use spreadsheets, are accustomed to entering a string of numbers to represent a time series, such as estimates of sales volumes for the next 10 or 20 years. To model the uncertainty and to help understand the results, these strings of numbers should be parameterized. For instance, for sales volume, we could enter an initial value and a growth rate instead of 10 or 20 numbers. Thus, decision analysis may require programmers to change their style of programming, especially if spreadsheets are used.

The following paragraphs contain modeling suggestions that may appear obvious to some. However, we have noticed that oversights in these areas are surprisingly frequent.

Figure 6-8

Input Section for Diagnostics' Spreadsheet

Alternative Table		DiagRF		DiagStatic			
#	Name	Device	Target Price	Device	Comp. Adv.	Graphics	Target Price
1	Milk Existing Products	Current	\$20,000	Current	Features	OEM	\$10,000
2	DiagRF Plus	Upgrade	\$25,000	Current	Features	OEM	\$15,000
3	DiagStatic New	Current	\$20,000	New	Link+Features	Integral	\$10,000

INPUT						
Description	In Use	Range Name	Low	Base	High	Units
Alternative	3	Alternative				1, 2, 3
Market						
Total market, 2000	3,500	Mkt_00		3,500		units/year
Total Market, 2010	4,500	Mkt_10	4,000	4,500	5,000	units/year
RF Fraction of Market, 2000	17%	Mkt_RF_fr_00		17%		percent
RF Fraction of Market, 2010	25%	Mkt_RF_fr_10	20%	25%	35%	percent
DiagStatic						
DiagStatic, Market Share, 2000	10%	MS_00_DSt		10%		percent
DiagStatic, Market Share growth rate	-20%	MS_gr_DSt	-25%	-20%	-15%	percent
DiagStatic, Year Discontinued	2005	Yr_DSt_stop		2005		year
If DiagStatic New is introduced before this year, DiagStatic is discontinued at that time						
DiagStatic, Price, 2000	\$9,500	ASP_00_DSt		\$9,500		\$/unit
DiagStatic, Price CAGR	0%	ASP_gr_DSt	-5%	-1%	1%	percent
DiagStatic COGS, 2000	\$4,500	COGS_00_DSt		\$4,500		\$/unit
DiagStatic, COGS CAGR	2%	COGS_gr_DSt	0%	2%	3%	percent
DiagStatic New						
DiagStatic New, Technical Success	Yes	tech_succ_DStn	No	Yes	Yes	yes/no
DiagStatic New, Start Year	2003	Startyr_DStn		2003		year
DiagStatic New, Years to Peak	3	Yrs_to_pk_DStn		3		years
DiagStatic New, Years at Peak	8	Yrs_at_pk_DStn		8		years
DiagStatic New, Years from Peak to 0	7	Yrs_to_0_DStn		7		years
DiagStatic New, Market Share at Peak	20%	MS_pk_DStn	10%	20%	40%	percent
DiagStatic New, Price, start year	\$14,000	ASP_sy_DStn	\$12,000	\$14,000	\$15,000	\$/unit
DiagStatic New, Price CAGR	0%	ASP_gr_DStn	-3%	0%	1%	percent
DiagStatic New, COGS, start year	\$4,800	COGS_sy_DStn	\$4,600	\$4,800	\$5,200	\$/unit
DiagStatic New, COGS CAGR	2%	COGS_gr_DStn	0%	2%	3%	percent
DiagRF						
DiagRF, Market Share, 2000	40%	MS_00_DRF		40%		year
DiagRF, Market Share growth rate	-10%	MS_gr_DRF	-20%	-10%	-5%	percent
DiagRF, Year Discontinued	2008	Yr_DRF_stop		2008		year
If DiagRF Plus is introduced before this year, DiagRF is discontinued at that time						
DiagRF, Price, 2000	\$17,000	ASP_00_DRF		\$17,000		\$/unit
DiagRF, Price CAGR	0%	ASP_gr_DRF	-4%	0%	1%	percent
DiagRF, COGS, 2000	\$4,500	COGS_00_DRF		\$4,500		\$/unit
DiagRF, COGS CAGR	2%	COGS_gr_DRF	0%	2%	3%	percent
DiagRF Plus						
DiagRF Plus, Start Year	2001	Startyr_DRFp		2001		year
DiagRF Plus, Years to Peak	2	Yrs_to_pk_DRFp		2		years
DiagRF Plus, Years at Peak	6	Yrs_at_pk_DRFp		6		years
DiagRF Plus, Years from Peak to 0	5	Yrs_to_0_DRFp		5		years
DiagRF Plus, Market share at Peak	40%	MS_pk_DRFp	25%	40%	45%	percent
DiagRF Plus, Price, start year	\$19,000	ASP_sy_DRFp	\$17,000	\$19,000	\$23,000	\$/unit
DiagRF Plus, Price CAGR	0%	ASP_gr_DRFp	-3%	0%	1%	percent
DiagRF Plus, COGS, start year	\$4,800	COGS_sy_DRFp	\$4,600	4,800	\$5,300	\$/unit
DiagRF Plus, COGS CAGR	2%	COGS_gr_DRFp	0%	2%	3%	percent
GLOBAL VARIABLES						
Base Year	2000	BaseYear		2000		year
Discount Rate	10%	DiscountRate		10%		percent
Tax Rate	38%	TaxRate		38%		percent
Sales, General, Administrative	10%	SGA		10%		% of unit sales
Working Capital, 1999	\$3,000,000	Wcap_99		\$3,000,000		\$
Working Capital	29%	Wcap		29%		% of revenue

R&D Investment (\$000)		2000	2001	2002	2003
DiagRF Plus	RD_DRFp	500	1,500		
DiagStatic New, pre-success	RD_pre_DStn	1,800	1,800		
DiagStatic New, post-success	RD_post_DStn		400	1,900	1,100

Figure 6–9

Calculation Section for Diagnostics' Spreadsheet

Alternative #	Name	DiagRF				DiagStatic							
		Device	Target Price			Device	Comp. Adv.	Graphics	Target Price				
3	DiagStatic New	Current	\$20,000			New	Link+Features	Integral	\$10,000				
		1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010
Market Units Sold													
	Static		2,905	2,950	2,996	3,042	3,088	3,135	3,182	3,230	3,278	3,326	3,375
	RF		595	639	685	732	782	833	887	943	1,001	1,062	1,125
Diagnostics' Units Sold													
	DiagStatic		291	236	192	156	0	0	0	0	0	0	0
	DiagStatic New		0	0	0	0	206	418	636	646	656	665	675
	DiagRF		238	230	222	214	205	197	189	180	172	0	0
	DiagRF Plus		0	0	0	0	0	0	0	0	0	0	0
Revenue (\$000)													
	DiagStatic		2,760	2,242	1,821	1,480	0	0	0	0	0	0	0
	DiagStatic New		0	0	0	0	2,882	5,852	8,911	9,044	9,178	9,314	9,450
	DiagRF		4,046	3,910	3,771	3,630	3,488	3,346	3,206	3,067	2,931	0	0
	DiagRF Plus		0	0	0	0	0	0	0	0	0	0	0
Cost of Goods Sold (\$000)													
	DiagStatic		1,307	1,083	898	744	0	0	0	0	0	0	0
	DiagStatic New		0	0	0	0	1,008	2,088	3,242	3,356	3,474	3,596	3,722
	DiagRF		1,071	1,056	1,038	1,020	999	978	956	933	909	0	0
	DiagRF Plus		0	0	0	0	0	0	0	0	0	0	0
Financial Results (\$000)													
	Revenue		6,806	6,152	5,592	5,109	6,370	9,199	12,117	12,112	12,110	9,314	9,450
	Costs		2,378	2,139	1,936	1,763	2,007	3,066	4,198	4,289	4,384	3,596	3,722
	Gross Margin		4,428	4,013	3,656	3,346	4,363	6,133	7,919	7,822	7,726	5,718	5,728
	SG&A		681	615	559	511	637	920	1,212	1,211	1,211	931	945
	R&D Expense		1,800	2,200	1,900	1,100	0	0	0	0	0	0	0
	EBIT		1,947	1,198	1,197	1,735	3,726	5,213	6,707	6,611	6,515	4,786	4,783
	Tax		740	455	455	659	1,416	1,981	2,549	2,512	2,476	1,819	1,818
	Earnings		1,207	743	742	1,076	2,310	3,232	4,159	4,099	4,039	2,967	2,966
	Working Capital	3,000	1,974	1,784	1,622	1,482	1,847	2,668	3,514	3,512	3,512	2,701	2,741
	Change in Working Capital		-1,026	-190	-162	-140	366	820	846	-2	-1	-811	40
	Cash Flow		2,233	932	904	1,216	1,944	2,412	3,312	4,101	4,040	3,778	2,926
	Cumulative Cash Flow		2,233	3,166	4,070	5,286	7,230	9,642	12,954	17,055	21,095	24,873	27,799
	PV Cash Flow		2,233	847	747	913	1,328	1,498	1,870	2,104	1,885	1,602	1,128
	Cumulative PV Cash Flow		2,233	3,081	3,828	4,742	6,070	7,567	9,437	11,541	13,426	15,028	16,156
NPV (\$ million)													
	Cash Flow 2000-2010		16										
	Terminal Value		10										
	Total		27	NPV									

Modeling Profit

Care must be taken in modeling how revenue and cost change over time. Because profit is a small difference between these two larger numbers, it is extremely sensitive to the details of modeling. For many businesses, it is better to estimate costs and margin (or revenue and margin) rather than revenue and costs directly.

Sunk Costs

It is surprising how much confusion arises if investments have already been made in the area of the decision problem. This money is spent (it is a “sunk

cost”), and the decision problem is what to do in the future, given present assets and experience. One approach is to exclude all sunk costs from the evaluation of the problem; the other approach is to include the sunk costs in all alternatives, including the “do nothing” alternatives.

Shutting Down the Business

Models should have a cutoff in them so that when margins are negative for more than a few years, the business shuts down. Models should also have checks to prevent unrealistic or impossible situations, such as a market share exceeding 100 percent. Remember that the decision tree will evaluate the model for a number of extreme cases and that unexpected values may occur in these cases, thus distorting the overall results.

Inflation

It is usually better to model in constant dollars, because you can then avoid simultaneously modeling both the systematic changes of prices and costs and the change from inflation. Remember, however, to deflate such items as interest payments, undepreciated capital, and working capital to account for the effects of inflation on them.

Terminal Value

Even when the model explicitly considers 10 to 15 years, as much as 50 percent of the net present value of a business may arise from the period beyond that modeled. This occurs especially for new product introductions and for research and development. It is therefore usually important to explicitly include a value for the business beyond the time period of the model. This value might be estimated as a sale value (perhaps expressed as a multiple of cash flow in the final year considered) or as a salvage value (estimated from assets, including working capital). In estimating this value of ongoing business, remember that no business lasts forever and that migrating to new forms of the business may require substantial new capital investments.

Diagnostics decided to focus on the NPV of cash flow. There were no sunk costs to worry about. There would be no cost associated with eliminating old products if DiagStatic New or DiagRF Plus were not developed; the resources devoted to these products could be gradually migrated to other areas of the business. Inflation effects were included in the estimates for price and cost.

The terminal value for the analysis was the cash flow from operations in the final year (2010) divided by the discount rate. This value is the NPV of operations that continue in perpetuity with the same cash flow as 2010. In this case, the terminal value arises only from new products (DiagStatic New or DiagRF Plus) which have sales beyond 2010. Because terminal value in this case is a small fraction of total NPV, this approximation was acceptable.

Deterministic Structuring: Sensitivity Analysis

One of the main tasks in the deterministic structuring phase is to sufficiently develop the model to produce credible base-case results. The base case is a set of input parameters used as a starting point for further analysis. If the decision alternatives are very different, it may be necessary to establish a base case for each alternative. If there is an underlying uncertainty of success or failure (as with a new product introduction), the base case should usually be estimated given success. Too much effort should not be spent in establishing base-case values, since they are just a starting point. On the other hand, the values should be reasonable enough not to stir up unnecessary controversy.

Ideally, the base-case parameters should all represent a median value, which means that in most people's judgment there is a 50/50 chance that the true value will be above or below the value. In practice, however, base-case parameters are often developed from an unduly optimistic or pessimistic business plan. As the analysis progresses, some effort should be spent establishing a set of base-case parameters close to most people's median value.

For all important variables, a range of values should be established in early interviews: the base-case value, a low value (10 percent chance that the value will turn out to be lower), and a high value (10 percent chance that the value will turn out to be higher). The next task is to identify the crucial uncertainties, which is accomplished using deterministic sensitivity analysis. Frequently, this process shows that uncertainties originally thought to be important are relatively unimportant.


Deterministic sensitivity analysis starts by evaluating the model with all variables set to their base-case values, which gives the base-case value of the model. Then, for each of the variables in turn, the model is evaluated at the low value of the variable and then at the high value of the variable (with all other parameters kept at their base value), thus yielding low and high values of the model. The difference between the low and high value of the model is a measure of how sensitive the problem is to the uncertainty in that variable.

The project team developed ranges for the data inputs for the spreadsheet (Figure 6–8) in the columns labeled Low, Base, and High. They then calculated the sensitivity for the Alternative 3, DiagStatic New. The results are shown in Figure 6–10.

The results shown in Figure 6–10 are plotted in what is called the “tornado” chart. The variables are arranged so that those having the largest swing are at the top. The vertical line shows the position of the base case. The horizontal bars show the size of the swing for each variable or joint sensitivity; the swing is the largest of the absolute values of the differences between the high, low and base values of the model for that variation.

Those variables with the greatest swing are the crucial uncertainties that dominate the problem. As such, they will be treated as explicit uncertainties in the decision tree, while the remaining variables will be set to their base-

Figure 6-10



case values (but may change with the decision alternatives).

Sometimes it is not reasonable to vary two or more variables independently—they are correlated. When one is high, the other is high (or low if they are anti-correlated). For these variables, it is appropriate to set up a joint sensitivity. The variables in the joint sensitivity are varied jointly, i.e., at the same time. Even if these variables individually appear low in the sensitivity analysis table, their joint effect may be large enough for them to be included in the tree. These variables may be entered in the tree as separate nodes (with dependent probabilities) or as a single joint node that sets the values for two or more variables on each branch.

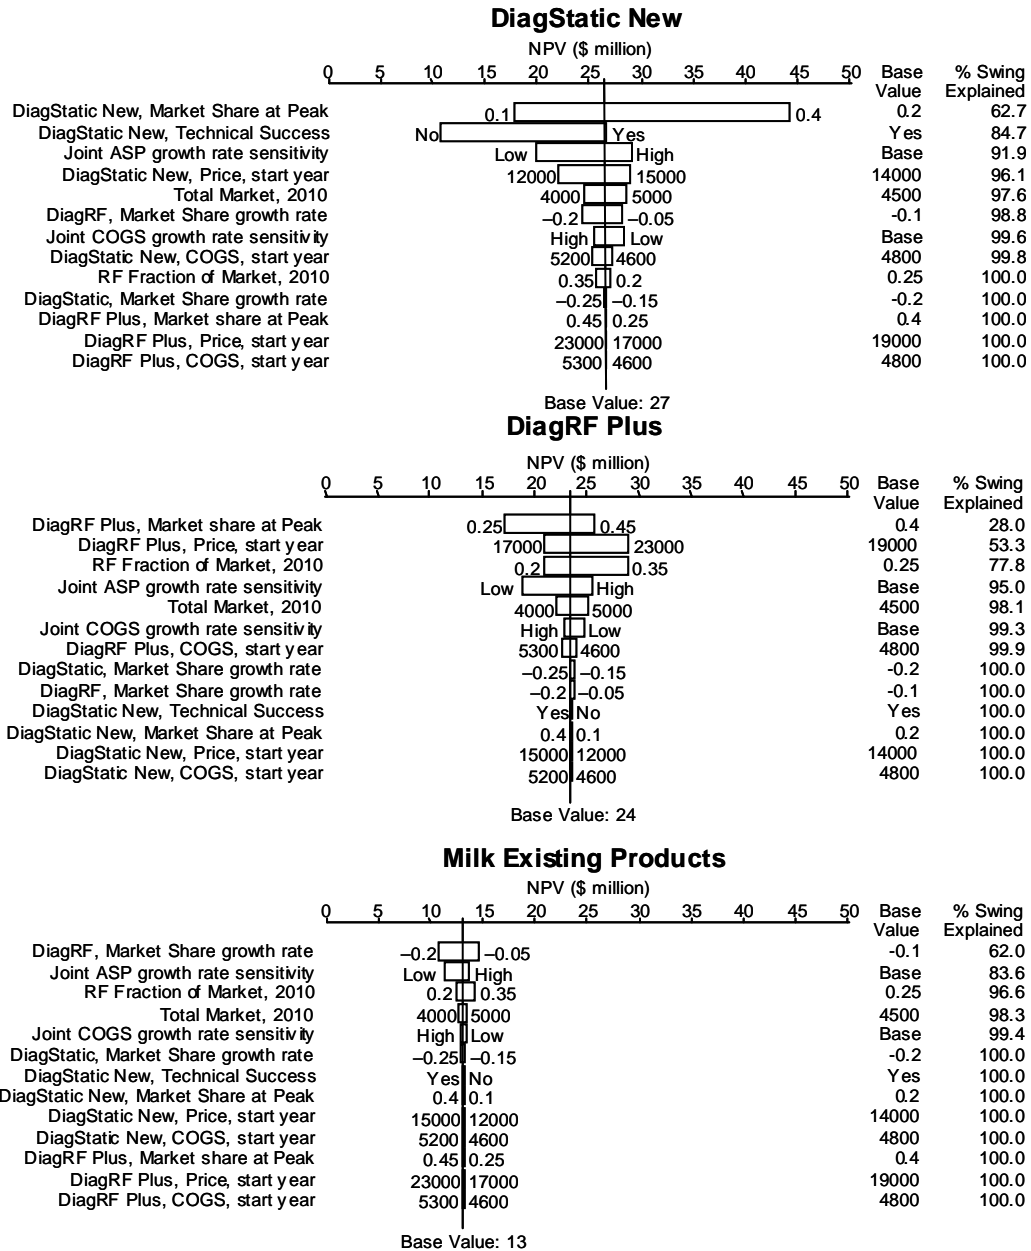
The project team realized that not all these variables were independent. The ovals in the influence diagram represented uncertainties for multiple products, and the uncertainties were linked. For instance, the compound annual growth rate (CAGR) for costs reflected underlying economics of the

Figure 6-11

Joint Sensitivities				
Joint COGS growth rate sensitivity				
Product	Variable	Low	Base	High
DiagStatic	COGS_gr_DSt	0	2%	3%
DiagStatic New	COGS_gr_DStn	0	2%	3%
DiagRF	COGS_gr_DRF	0	2%	3%
DiagRF Plus	COGS_gr_DRFp	0	2%	3%
Joint Average Selling Price growth rate sensitivity				
Product	Variable	Low	Base	High
DiagStatic	ASP_gr_DSt	-5%	-1%	1%
DiagStatic New	ASP_gr_DStn	-3%	0	1%
DiagRF	ASP_gr_DRF	-4%	0	1%
DiagRF Plus	ASP_gr_DRFp	-3%	0	1%

Figure 6-12

Joint Sensitivities



business, and would be correlated for all the products under consideration—if one were at its low (or high) value, all would be at their low (or high) values. Similarly, the CAGR for average selling price reflected the forces pressing for reductions in medical expenses, and would be the same for all the products. So these variables were grouped into two “joint sensitivities” rather than being varied individually, as shown in Figure 6–11. The sensitivities were then recalculated for all three alternatives, and are shown in Figure 6–12.

As mentioned above, when two or more variables are varied jointly, they often appear higher in the tornado than the variables varied independently. However, the effect was small for the DiagStatic New alternative.

Some variables in the sensitivity analysis table may be decision or value variables (such as the discount rate). These variables may have been included in the sensitivity analysis for general interest, but they should not be confused with the critical/– uncertainties that will go into the tree as chance nodes.

Often, the largest uncertainties affect all the alternatives in the same way, causing the model result for each alternative to vary up or down about the same amount. Usually, the uncertainties that affect all the alternatives this way are underlying macroeconomic variables. To conserve tree size and leave space for the uncertainties that distinguish between alternatives, these macroeconomic variables can be combined into one or two nodes. Each branch of the combined node would then represent a scenario in which several macroeconomic variables are set.

Since ordinarily a few uncertainties really do dominate the problem, we should not be afraid to limit the number of uncertainties in the tree to several with the largest swing in the sensitivity analysis. However, we need to examine our assumptions and results a little more closely to understand how to recognize this dominance. The basic assumptions are listed below.

1. The certain equivalent is the measure that takes into account the uncertainty of a venture. We can therefore determine how important a variable is to the overall uncertainty by determining how important it is to the certain equivalent value.
2. As mentioned in Chapter 5, we can approximate the certain equivalent using the following equation:

$$\text{Certain Equivalent} = \text{Expected Value} - \frac{1}{2} \frac{\text{Variance}}{\text{Risk Tolerance}} \quad (6-1)$$

Since the variance accounts for the way uncertainty affects the certain equivalent, the variance is the measure of uncertainty we are looking for.

3. To obtain the variance, we recall the statistical result that, for probabilistically independent variables and for small variations of these variables, the total variance is the sum of the variances for each variable. The variance for each variable is the square of the standard deviation, which is, in turn, proportional to the width of the distribution or swing for each variable. (Recall that

the swing is obtained by evaluating the model for the 10 percent and 90 percent points for each variable and see Figure 10–17.)

4. Thus, the square of the swing for a variable determines how important it is to the overall uncertainty.

Accordingly, one easy way to compare the importance of uncertainty in the variables is to square the swing and express it as a percentage of the sum of squares of all the swings. The final column in Figures 6–10 and 6–12 show this swing. (Technically, the variation for Technical Success for the DiagStatic New alternative is not a 10/50/90 variation and distorts the overall uncertainty figures shown in the final column.)

Thus, we see that, for the DiagRF Plus alternative, the top four variables (peak market share, initial price, RF fraction of market in 2010, and average selling price) account for 95% of the uncertainty. In most problems with more variables and complex models, we find that the top three or four variables capture 80 to 90 percent of the effects of uncertainty.

Remember, though, that for this approximation to be valid, we need to use probabilistically independent variables. With dependent variables, use the joint sensitivity as illustrated above. This joint sensitivity must then be probabilistically independent from the other variables listed.

Probabilistic Evaluation: Building and Pruning the Tree

Even after deterministic sensitivity analysis, it is often difficult to limit the tree to a reasonable size. If there are, for instance, n nodes in a tree and each node has three branches, a symmetric tree will have 3^n paths. For example, for seven nodes, we would have 2,187 paths. If the model takes 1 second to calculate an answer (not untypical of a large spreadsheet model on a personal computer), the evaluation command will take a little over half an hour to execute. If we add several more nodes, change a node from three to four branches, or construct a more elaborate model, we will have a tree that is impractical to evaluate on a personal computer.

How large a tree is reasonable? This depends greatly on the type of problem (and on the opinion of the facilitator). However, we feel that for most problems, 50 to 200 paths per alternative is sufficient. This number of paths allows us to include the three generic uncertainties that often affect an alternative: uncertainty about the growth of the market, uncertainty in competitive action or reaction, and uncertainty in how well we will fare. At three branches a node, we have 27 paths per alternative, leaving room for several other nodes if called for. After the full-scale analysis is complete, we will probably find that around 20 paths per alternative are enough to draw all the conclusions. Reducing the tree to this size is often important for clarifying the results and drawing the tree for the final presentation to the decision-maker.

How can you make your tree small enough to evaluate? By reducing the number of branches at each node, by reducing the number of nodes, and by creating asymmetric trees. Simplifying the model can also help by reducing

the time required to evaluate the tree. We consider each of these options in turn.

What is a reasonable number of branches at a node? Our experience indicates that chance nodes are usually well approximated by three branches. Going to four or more branches seldom perceptibly affects the overall profit distribution. While we can use two branches for uncertainties that deterministic sensitivity analysis shows to be less important, there is no central branch in a two-branch node to trace the effect of the base case. With decision nodes, we can use the preliminary evaluation to eliminate the inferior alternatives and narrow down to the three or four really distinct and most promising alternatives.

What is a reasonable number of nodes in a tree? A good number to aim for is five or six nodes. Normally, there are too many nodes in the initial version of the tree, and the tree has to be pruned before it can be evaluated. In pruning the tree, four types of nodes compete for a place on the final tree: decision nodes, chance nodes with effects common to all alternatives, chance nodes with effects that distinguish between alternatives, and chance (or decision) nodes that are not really important to the tree but should be included for political reasons.

By using the results of the deterministic sensitivity analysis, we can usually manage to discuss and eliminate the politically important nodes from consideration. Indeed, this may be one of the important insights from the analysis. If necessary, chance nodes with effects common to all alternatives can be combined into one or two nodes whose branches represent scenarios (combinations of events). If there are still too many nodes, we must be creative in combining variables, restructuring the tree, and using further sensitivity analysis to narrow down the number of important uncertainties even more. In cases of real necessity, we may have to run a separate tree for each alternative. If we do this, we will have to do some calculations by hand to reorder the tree (i.e., to obtain value of information), but we can still have software do most of the work for us.

Creating asymmetric trees is another way to reduce tree size. For instance, when uncertainty is important for one alternative but not for another, we can have the chance node for this uncertainty follow only the alternative for which it is important. Similarly, an uncertainty may be important only for certain branches of a chance node, in which case you can have the node follow only those branches.

Finally, we can simplify the model to shorten its running time and, thus, the time required to evaluate the tree. Models can be simplified by eliminating calculations of quantities that were initially of interest but that are not necessary for calculating the net present value of cash flow. Models can similarly be simplified by replacing complicated calculations with simplified ones that produce approximately the same results.

In general, beginning with a small version of the tree is better, even if you are using a software program to evaluate the tree. Nothing is as discouraging as starting off by running up against multiple errors from hardware/software

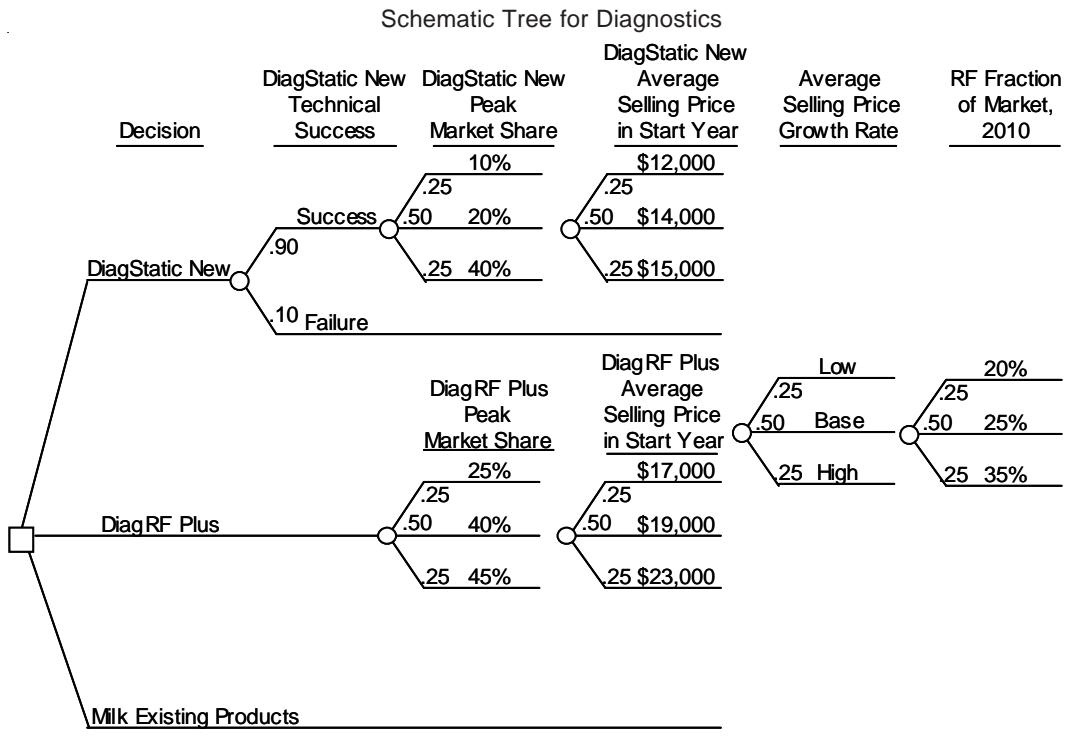
limitations—correcting one error, waiting a long time for the tree to reevaluate, and then finding another error.

After examining the sensitivity analysis results in Figure 6-12, the team drew the schematic decision tree shown in Figure 6-13. These seven uncertain variables account for almost all of the uncertainty in the three alternatives. All other variables were set to their base-case values or determined by model calculations. The decision node occurs before (to the left of) all the uncertainty nodes because information on the resolution of these uncertainties will not be available until after the decision is made.

The team started to more carefully examine the uncertainty on the variables in the tree. The previous ranges were fine for sensitivity analysis, but now the facilitator wanted to assess complete distributions to obtain better quality information for the tree. (See Chapter 12 for more on this process.) After some careful thought, however, the team judged that the information prepared for the sensitivity analysis represented their best state of information. They assigned .25/.50/.25 probabilities to the values used in the 10/50/90 ranges in the sensitivity analysis—see Chapter 2 for a justification of this process.

The one probability they were missing was the probability of technical success for DiagStatic New development. A lengthy discussion of technical hurdles with the design staff established that the probability of technical

Figure 6-13



success was 90 percent.

Basis Appraisal: Obtaining Information from the Tree _____

The facilitator began analyzing the tree by plotting the probability distributions (Figures 6–14 and 6–15). The distributions were shown in two forms.

Figure 6–14 shows the cumulative probability distributions for each of the alternatives. Figure 6–15 shows a plot of the 10-90 percentile range around the expected value for each alternative. For each bar, there is a 10 percent chance the value falls to the left of the bar and a 10 percent chance it falls to the right; the star near the center marks the expected value; the lines on either side of the bar show the full limit of the distributions. This simpler plot contains the essential information needed to compare the risk and return of each alternative. Because they are so simple and intuitive, plots like Figure 6–15 are frequently used in presentations.

The expected values showed the DiagStatic New alternative was preferred, with an expected value of \$25 million. The DiagRF Plus alternative was a close second, with an expected value of \$22 million. The Milk Existing Products alternative had an expected value of only \$13 million.

However, as illustrated by the width of the distributions, The DiagStatic New alternative was also the most uncertain, with potential values ranging from \$9 million (\$2 million lower than \$11 million, the lowest value for Milk Existing Products) and \$55 million (\$16 million higher than \$39 million, the highest value for the DiagRF Plus alternative.)

On an expected value basis, Diagnostics should go with the DiagStatic New alternative. In terms of the budget request, it seemed clear that creating a new product (DiagStat New or DiagRF Plus) was \$9-13 million better than doing nothing (Milk Existing Products.)

Besides getting expected values (certain equivalents) and probability distributions on profit for the alternatives, we can get a wealth of information to help us glean all possible insights from the tree.

Conditional Distributions

First, we can look at the distributions and expected values (certain equivalents) at points other than the first node in the tree and see if these conditional values hold any surprises.

To better understand where the potential risks and profits came from, the facilitator decided to look at what the distributions for the alternatives would be, given technical success or failure of DiagStatic New (Figure 6–16). For technical success, DiagStatic New is clearly preferred to the other alternatives; for technical failure, DiagStatic New reverts to the Milk Existing Products alternative less the \$3.6 million R&D investment incurred before failure is known.

He then plotted the conditional distributions for DiagStatic New for Peak Market Share (Figure 6–17), which shows dramatic differences between the three values; note that the lowest 10 percent of the distribution shows no difference, corresponding to technical failure. Figure 6-18 shows DiagStatic

Figure 6-14

Cumulative Probability Distributions for Diagnostics' Alternatives





Figure 6-15

10-90 Bars for Probability Distributions for Diagnostics' Alternatives



New for the uncertainty on the compound annual growth rate on Average Selling Price for all the products. Figure 6-19 shows DiagStatic New for the different values of the Average Selling Price in DiagStatic New start year, 1003. Figure 6-20 shows DiagStatic New for the different values of RF fraction of total market in 2010.

Value of Perfect Information

Besides looking at conditional distributions, we can calculate the value of perfect information for each uncertainty. If possible information-gathering efforts have costs comparable to these values, we can formulate the decision for gathering imperfect information and see if it is worthwhile.

The facilitator rolled back the tree with the chance nodes brought one by one in front of the decision node and then took the difference between the resulting expected value and the original expected value to obtain the values of information (Figure 6-21).

Peak Market Share for DiagStatic New has the largest value of information, even more than Technical Success. RF Fraction of Total Market, 2010, has a larger value of information than one would have expected from Figure 6-20; although DiagStatic New is not very sensitive to this variable, DiagStatic Plus is quite sensitive (see the sensitivity bars in Figure 6-12), creating the relatively large value of information.

The value of perfect information on all uncertainties together is \$3.42

Figure 6-16

Diagnostics' Alternatives, Given Technical Success or Failure

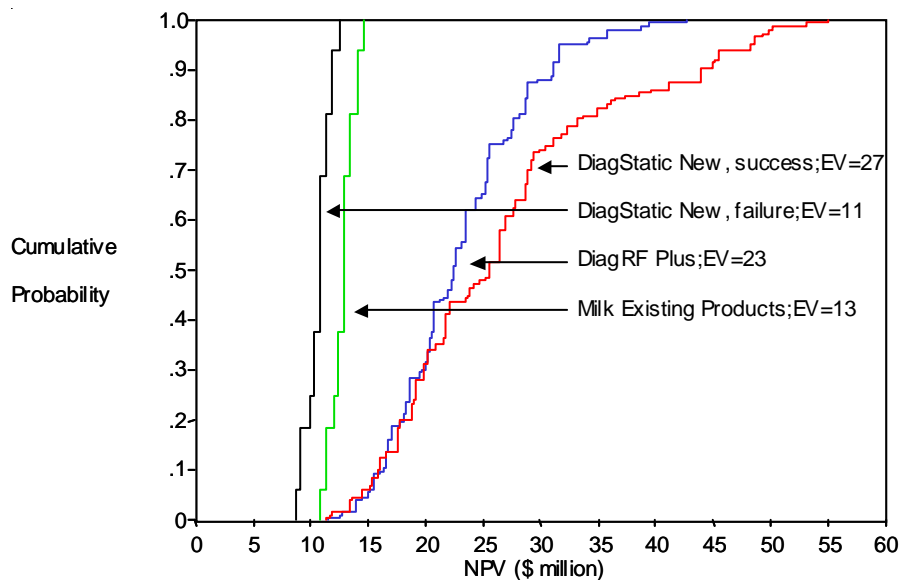


Figure 6-17

DiagStatic New, Given Different Peak Market Share




Figure 6-18

DiagStatic New, Given Different Joint Average Selling Price Growth Rate




Figure 6-19

DiagStatic New, Given Different Average Selling Price in the Start Year, 2003




Figure 6-20

DiagStatic New, Given Different RF Fraction of Total Market, 2010

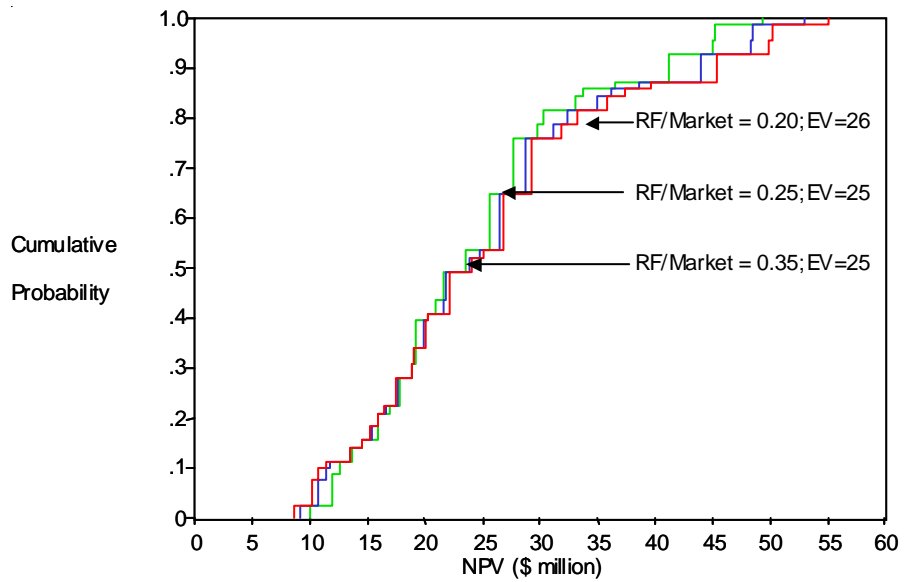


Figure 6–21

The Values of Perfect Information and Perfect Control for Diagnostics

Uncertainty	Value of Perfect Information (\$ million)	Value of Perfect Control (\$ million)
Technical Success, DiagStatic New	1.22	1.60
Peak Market Share, DiagStatic New	1.70	12.61
Peak Market Share, DiagRF Plus	.28	1.10
Average Selling Price CAGR	0	3.39
Average Selling Price 2003, DiagStatic New	.43	2.78
Average Selling Price 2001, DiagRF Plus	.57	2.26
RF Fraction of Total Market 2010	.70	2.12
All seven uncertainties	3.42	29.88

million, less that the sum of the value of information for each uncertainty separately, \$4.96 million.

Given the very short time allotted to the analysis, no realistic information gathering activities were possible concerning the market. However, it was noted for the record that discussions with potential customers might shed some light on the market uncertainties.

When the design staff saw the value of information for technical success in developing DiagStatic New, they commented that they could clear up most of this uncertainty with a few weeks of concerted effort at a very small cost.

Laura could not wait the few weeks requested by the design staff before submitting the budget request. However, she planned to include this information in the request for DiagStatic New along with a contingency request if design staff indicated that DiagStatic New was technically infeasible.

Value of Perfect Control

Another insight we can glean from the tree is the value of perfect control. If there are feasible control procedures of reasonable cost, we can generate the new alternative of using these control procedures and reevaluate the tree.

To obtain the values with perfect control, the facilitator displayed the tree with the uncertain nodes brought to the front of the tree one by one; then he picked the best expected value for that node and took the difference between it and the original expected value (Figure 6–21).

Peak Market Share for DiagStatic New has the largest value of control, even larger than control of Technical Success. The value of control on all seven uncertainties was \$55.07 million, considerably greater than the sum of the value of control for each individual uncertainty, \$25.89 million.

As for controlling these uncertainties, the facilitator suggested long-term

favorable contracts with customers might achieve the higher market share levels with modest price adjustments.

Sensitivity Analysis to Probabilities, Risk Attitude, and Value Trade-offs

We can use the tree to perform sensitivity analysis to probability, risk attitude, and value trade-offs (such as the time value of money captured in the discount rate). We can also form new sensitivity displays by picking two values for the axes of a graph, such as the probabilities for two different uncertainties, and by showing which alternative is preferred in which region of the graph.

Sensitivity analysis results like these can suggest areas for refining estimates; show the consequences of using different, conflicting estimates from experts; and avoid potential “Monday morning quarterbacking.”

The facilitator had already done a form of probabilistic sensitivity by looking at the conditional probability distributions in Figures 6–16 to 6–20. Now, however, he systematically varied probabilities to see if the preferred decision would change with changes in the probabilities for a particular uncertainty.

If there are more than two branches at a node, there are many ways to vary probabilities for a sensitivity analysis. One method is to put a probability p on the top branch and the rest of the probability ($1-p$) on the bottom branch, and then systematically vary p . If the top and bottom branches are the “best” and “worst” outcomes, you are varying in a systematic way between best and worst.

This is the method used in Figures 6–22 to 6–28. The original state of information had the probability symmetrically distributed around the middle branch; when we go to the two-branch approximation, this translates into assigning 50 percent probability for the top and bottom branches. The one exception is Technical Success, which was originally a two-branch node with 90 percent probability on the top “Success” branch and 10 percent on the bottom “Failure” branch.

The probability sensitivities show that all of the uncertainties (except for Average Selling Price CAGR) have the potential to switch the best alternative from DiagStatic New to DiagRF Plus.

For Technical Success, Figure 6–22 shows that for probabilities of technical success over 70 percent, DiagStatic New is the preferred alternative. Everyone agreed that the probability was in the 90 percent range, so there was no concern about this probability.

The remaining probabilities (Figures 6–23 to 6–28) had fairly general agreement that, in the two-branch approximation, each branch had about a 50 percent probability.

So the conclusions appeared to be quite robust against differences of opinion concerning probabilities.

The facilitator put together several sensitivity plots against pairs of probabilities, such as the one shown in Figure 6–29. In this plot, the vertical axis is the probability that the peak market share for DiagStatic New is Low (10%) and the horizontal axis is the probability that the peak market share for DiagRF Plus is

Figure 6-22

Sensitivity to the Probability of Technical Success for DiagStatic New




Figure 6-23

Sensitivity to the Probability that DiagStatic New Peak Market Share is Low




Figure 6-24

Sensitivity to the Probability that DiagRF Plus Peak Market Share is Low




Figure 6-25

Sensitivity to the Probability that Average Selling Price Growth Rate is Low




Figure 6-26

Sensitivity to the Probability that DiagStatic Price in 2003 is Low

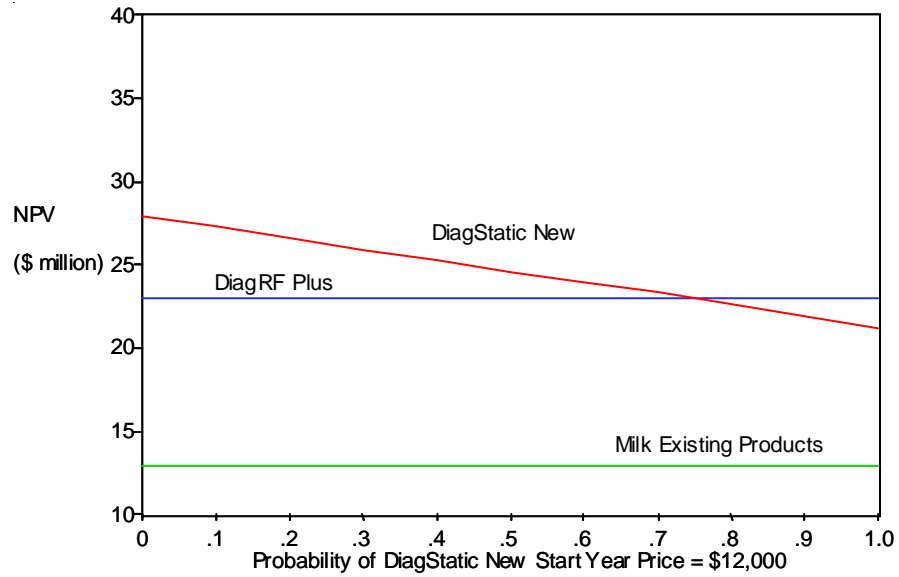


Figure 6-27

Sensitivity to the Probability that DiagRF Price in 2001 is Low

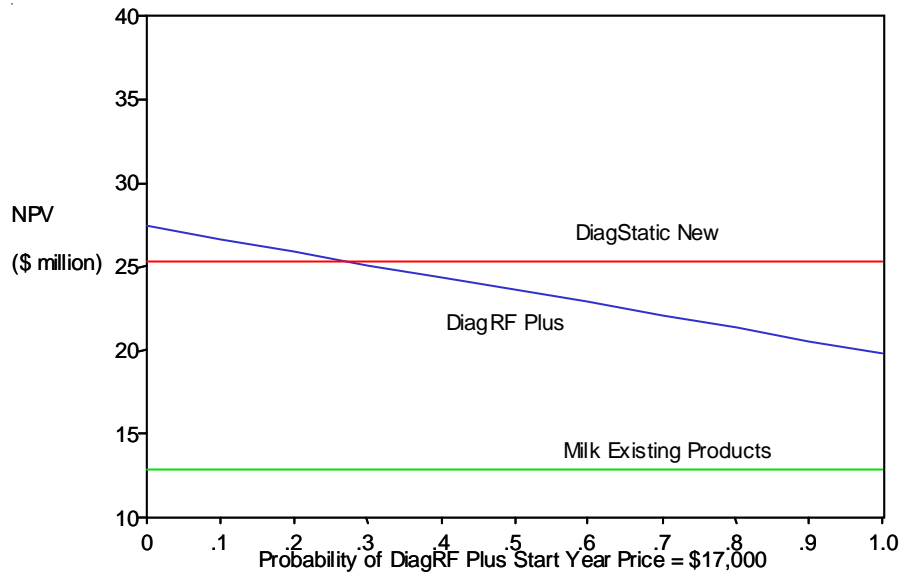


Figure 6-28





Figure 6-29



Low (25%). In the two-branch approximation, the current state of information is that both of these are 50 percent, the point marked X in the plot. If someone were to choose a different set of probabilities, they could refer to this plot. If the point is above the line, DiagRF Plus is preferred, and if it is below the line, DiagStatic New is preferred.


Finally, the facilitator heard someone refer to DiagStatic New as “risky.” He immediately did a sensitivity to risk attitude, as shown in Figure 6–30. Unless the risk tolerance of Medequip is less than \$3 million, DiagStatic New is preferred. Using the rules of thumb described in Chapter 5, the facilitator estimated that Medequip’s risk tolerance should be on the order of \$1,000 million. Clearly, this was not a “risky” choice from Medequip’s point of view.

Policy Matrix

If there are decisions at other than the first node in the tree, we can generate a policy matrix to show what they would be. A policy matrix lists all paths through the tree that lead to the final decision node and the alternatives that would be chosen for that path. The policy matrix shows the importance (or lack of importance) of explicitly considering the decision in the analysis and may help in establishing a monitoring program and contingency plans. The simplest way to obtain the policy matrix is to display the tree for all the nodes up to and including the rightmost decision node.

For Diagnostics, the only decision node was in the front of the tree. However, if Diagnostics did choose to seek out imperfect information on the uncertainty with the highest value of perfect information (peak market share

Figure 6–30



for *DiagStatic New*), they might get the policy matrix of Figure 6–31.

The analysis had reached the point in the decision analysis cycle (Figure 6–1) at which Basis Appraisal is complete. Now it was time to decide whether to act (decide on a budget request) or iterate the analysis.

Laura and the facilitator pondered further improvements to the analysis. The most glaring omission was not explicitly considering the alternative of developing two products for both the Static and RF market. They knew that no amount of sensitivity analysis would reveal the value of an alternative that was not included in the analysis.

An alternative with two products appeared quite attractive, but developing both would put quite a strain on the development staff. Both products could not be developed on the time schedule allotted for each when developed singly. In addition, there would be effects on pricing and market share if products were introduced in both the Static and RF market. Perhaps a platform could be developed which would support both markets.

These new alternatives violated the policy that had been developed at the beginning of the cycle—but they both knew that one of the most valuable outputs of an analysis can be results that suggest that policy should be violated.

The facilitator also suggested they might include the profit from the service part of the business. Diagnostics had a service force that kept its products in working order and made a modest profit. This source of profit was deliberately excluded from the analysis because Medequip was not certain that it wanted to keep the service organization or whether it wanted to sell it.

While Laura could see how added refinement might provide more insights into the decision, she was rather more worried about preparing a good presentation for the budget request. If the budget request succeeded, there would be ample time to refine the analysis and

Figure 6–31

Hypothetical Policy Matrix for Diagnostics

Seek Imperfect Information on Peak Market Share for <i>DiagStatic New</i>	Peak Market Share for <i>DiagStatic New</i>	Decision
No	None	<i>DiagStatic New</i>
Yes	10% 20% 40%	<i>DiagRF Plus</i> <i>DiagStatic New</i> <i>DiagStatic New</i>

optimize the decision.

Time to Prepare and Present _____

A final, equally important phase of the analysis takes place after the last pass through the cycle and the last tree analysis. The facilitator and project team members should set time aside to explore the information contained in the analysis, obtain insight into the real-world problems and questions the analysis was intended to answer, and make sure the analysis addresses all the decision-maker's concerns. The insights and results then need to be packaged for effective communication to the decision-maker. (See the end of Chapter 8 for details.)

An all-too-frequent mistake is to continue the analysis until the last possible moment. There is not sufficient time for the careful preparation necessary for synthesizing, summarizing, and presenting the conclusions. As a result, the report or presentation can be a haphazard one that may omit crucial points, may contain analytical errors, and may (worse yet) be difficult to understand. Such a report may fail to cogently address the decision-maker's concerns and thus fail to provide motivation for action—which relegates the results to a file rather than establishing them as a basis for action.

Laura presented the results at the budget meeting. Because she knew the board would want to see some of the nuts and bolts that went into the numbers, she presented a base-case run of the spreadsheet model for each alternative and then the tree results. The board was very impressed with the thoroughness of the analysis and with the way Laura produced a reasonable picture of Diagnostics' opportunities.

Not surprisingly, the board asked Laura to have the design group do some work on the technical feasibility of DiagStatic New before they gave final approval. Rather surprisingly, however, the board also requested that she look into alternatives that would capitalize on the opportunity in both the Static and RF markets. The two opportunities looked attractive, and the board was willing to stretch the budget rather than leaving profits sitting on the table.

Summary _____

The decision analysis cycle makes it possible to tackle complex, real-world problems by breaking the decision analysis into four stages: basis development, deterministic structuring, probabilistic evaluation, and basis appraisal.

In the basis development stage, the problem is described, a value measure is chosen, the important decisions and uncertainties are identified, and the relationships between these decisions and uncertainties and the value measure are sketched out.

In the deterministic structuring stage, these decision and chance

variables are related together in a deterministic model that calculates a value for any combination of these variables. Deterministic sensitivity analysis is then used to identify the key decisions and uncertainties.

In the probabilistic evaluation stage, these crucial decisions and uncertainties are structured into a decision tree, which is then analyzed.

In the basis appraisal stage, the results of the previous three stages are used to review the analysis for relevance and insight and to formulate recommendations (that may be for action or for further study).

When applied correctly, this cyclical process makes most decision problems tractable. In the case of Diagnostics, the two-week limit allowed only one real pass through the cycle. But the structured approach ensured that at the end of the week, the best possible results were available for decision-making.

Problems and Discussion Topics

- 6.1 A friend has come to you for help in deciding how to maximize his grade point average. What steps would you go through in developing a basis for the decision? What would you have at the end of the basis development? What are the next steps?
- 6.2 In the decision analysis cycle, there is a deterministic structuring phase. This phase often includes deterministic sensitivity analysis. In Chapter 3, there is a discussion of probabilistic sensitivity analysis. What are the differences between the two kinds of sensitivity analysis? What effects (or functions) does deterministic sensitivity analysis have in dealing with complex problems? Compare these two sensitivities with the sensitivity to risk tolerance seen in Chapter 5.
- 6.3 The complexity of a real-world problem is also reflected in its dynamic nature. The process of analyzing a decision problem can create new decision problems and add to the complexity of the original problem. A typical decision faced by the decision-maker (company) after the preliminary or pilot analysis is whether to proceed with the recommendation from the analysis or gather further information. Fortunately, the decision analysis cycle provides a framework with which to make this decision. Frequently, there are even preliminary numerical results available, such as the values of information.

What other complexities can arise in the course of analyzing the decision problem? (Hint: consider the elements of the decision basis.)
- 6.4 The ways of making decisions can be divided into normative methods and descriptive methods. Normative methods describe what people should do in a given situation. Descriptive methods focus on what people typically do.

For instance, if you face a decision on whether to hold on to a stock or

sell it, a decision analysis (normative method) would tell you what you should do. Descriptively, many people make this kind of decision by asking their spouse, broker, or friends to effectively make the decision for them.

Is it possible to reconcile the two methods of decision-making? Provide an argument and example to support your judgment.

- 6.5 The framing of a decision problem describes how the decision is stated. An example is describing the effects of a new, dangerous, and relatively untested drug in terms of either net lives saved or net lives lost. Decision-makers are often affected by the framing of the problem. For example, the information they provide and preferences they express will vary with the framing of the problem. Many of these framing issues are believed to be psychological in origin. What could you do to avoid these problems?
- 6.6 Innovative Foods Corporation (IFC) is a wholly owned subsidiary of Universal Foods Corporation (UFC). UFC is a major Fortune 500 company in the food processing industry. IFC is in the market of supplying specialized processed foods for human consumption. In 1979, the total market for specialized processed foods amounted to \$600 million and is growing.

IFC's leading products are dehydrated and processed foods targeted at two consumer groups: people on a special diet and people who are recovering after a serious illness. IFC has been using a well-known additive in its food processing, Divit, which is a recognized food additive in the food processing industry.

By carefully researching and testing its products, IFC has established a secure but small market share. IFC has a reputation as a good and reliable company and anticipates a \$2 million net yearly cash flow after taxes for the next 20 years.

In the last year, IFC management has become acquainted with some troublesome experiments carried out by its research division. The results of these experiments indicate a high probability that Divit is carcinogenic when applied in very high concentrations to the skin of mice. The carcinogenic properties when Divit is ingested by humans are by no means certain. IFC believes its competitors may be on the same track. If Divit turns out to be carcinogenic, the FDA will surely ban it, thus bringing about a major decrease in IFC's earnings and probably the loss of most of IFC's hard-earned market share.

However, IFC management has also been informed about another option. Its research division has developed another additive, Biovit. The research director believes Biovit is an excellent food additive that will have none of the problems of Divit. At present, the manufacturing costs

of Biovit are uncertain. To process its food with Biovit, IFC will have to invest around \$150 million. IFC management must make an important decision: should it continue to use Divit and face its potential banning or should it invest in new facilities and start using Biovit?

Assume you are a member of IFC's executive committee. How would you structure your thinking about this problem? Are there ethical considerations?

- 6.7 Plastic Co. is a fairly large company that manufactures bulk plastics for a large variety of uses. It has an extensive network of customers—companies that turn the bulk plastic into items that are then sold to the end-user.

Tech Co. is a European company that owns a process for formulating a special plastic that is useful for making bearings for high-speed centrifuges. There are several other potential high-tech applications for this special plastic. The same production equipment can also be used for making a common, low-margin type of plastic, which, it turns out, is not a plastic that Plastic Co. currently makes.

Tech Co. does not want to enter the U.S. market and is offering Plastic Co. an exclusive license for the manufacturing technology. The asking price is a \$500,000 license fee plus 5 percent of sales for 10 years.


Plastic Co. has determined that the equipment could come in a small size (3 million pounds per year) or a large size (10 million pounds per year). The respective costs for the equipment are \$3 million and \$7 million. Tech Co.'s experience has been that production costs are \$0.20 per pound.

- a. Begin to develop the decision basis for this decision. What are the alternatives? What are the uncertainties? What is the value? What information do you have? What information do you still need?
 - b. Designate someone to be the decision-maker. She will be the president of Plastic Co. Review with her the work you have done so far and finish structuring the decision. Assess any further information you may need.
 - c. Analyze the decision and produce a recommended course of action. Review the recommendation with the decision-maker. (Does she need to know the details of the analysis?) Is your recommendation useful to her? Does she believe and understand it? Why or why not?
- 6.8 Form a group to analyze a decision about whether or not to add a salad bar to a pizza parlor. Designate at least one person to be the client and one person to be the facilitator. (You may have more than one of each.) Make a pass through the decision analysis cycle as described below.

- a. **Background.** Develop an image of the pizza parlor that is as realistic as possible for whoever is playing the client. (This will greatly aid in the assessments.) How large is the place? How old is it? Who owns it? Who runs it? What kind of an area is it located in? What kind of clientele does it have? What is currently on the menu? What is the monthly sales volume? How many customers does it have daily? What are the peak hours? What kind of decor does it have?
 - b. **Basis Development.** Develop the basis for the decision. If you like, you may use an influence diagram for this step. What decisions must be made? What are the significant uncertainties? How do they relate to one another? What are the values on which the decision will be made? Try to keep the problem description simple.
 - c. **Deterministic Structuring.** Develop a model to determine the value for any scenario the tree might generate. The model may be assessed values, a Basic endpoint expression in Supertree, or an external spreadsheet model. Use sensitivity analysis, if necessary, to reduce the number of variables in the tree. Focus on modeling to help your understanding of the problem and to distinguish between alternatives.
 - d. **Probabilistic Evaluation.** Build and analyze the decision tree. For simplicity, try to keep the number of nodes down to four or five. You may start with a larger tree and then eliminate the nodes that do not distinguish between alternatives. Examine profit distributions, expected values, tree drawings, probability sensitivities, etc. Check that the results are consistent with your understanding of the problem.
 - e. **Basis Appraisal.** What is the preferred decision? How do its expected value and risk compare with those of the other alternatives? Is the preferred decision sensitive to changes in probabilities or risk attitude. What are the values of information and control? Would the client feel comfortable acting at this point, or would further study be advisable?
 - f. **Action.** Prepare a list of requirements for implementing the recommended alternative. These may include allocating funding, hiring personnel, hiring contractors, etc., or there may be no requirements if the recommendation was to do nothing. Has the analysis shed any light on the steps required for implementation? Is there any value to updating the analysis periodically to provide further guidance?
- 6.9 Form a small group to perform a decision analysis of a case study you have previously worked on. Assign roles. You will need at least one facilitator and one client who can supply structure and probabilities. Complete at least one pass through the cycle, perhaps limiting the exercise to two or three hours. Spend most of the time structuring the problem and preparing a final report.

- 6.10 Insitu Corp., an energy company, had developed a new technology for oil drilling in cold climates. The technology involved injecting a heated chemical solution into the well field at one location and waiting for the solution to percolate through the oil-bearing formation. Then, the solution was pumped out of the well field at another location and the oil was extracted in a processing plant.

Insitu had proven this technology on a pilot scale and was considering whether to build a full-scale project on its Whalebone property in Alaska. One of the major uncertainties was the capital cost of constructing the complex, consisting of the plant, pipeline, and well field. An engineering and design firm had estimated a base cost of \$320 million. To obtain financial backing for the project, Insitu felt it needed to verify this cost.



A team assembled in September to review the cost estimate identified two major risks in the estimate. First was the question of the efficiency of the new technology. In the pilot plant, a flow rate of 500 gallons per minute had produced a solution 30 percent saturated with oil. However, if the full-scale process were less efficient and produced, say, only 10 percent saturation at a flow rate of 1,000 gallons per minute, additional equipment would be required for the volume of oil produced to remain constant. Engineering estimates the additional equipment would add 15 percent to the base cost.

The second major risk was the productivity of the union workers. The largest influence on productivity was the unemployment rate in the

area. If unemployment were high, then the workers would be less likely to strike and would work harder. Unemployment, in turn, depended on the number of large pipeline and energy projects competing for workers and, ultimately, on energy prices. Changes in energy prices in recent years had been correlated with productivity variations as large as +30 percent. The team decided to include in its estimate a contingency to reflect these risks in the capital cost.

As the team was about to adjourn, someone asked if there were any other reasons the base cost could be exceeded. An inexperienced staffer, Ms. Pessi Mist, asked whether they were sure the construction would be finished on time. Since wage and materials rates were escalating at almost 25 percent a year, she felt a late construction schedule would increase costs. Her question was met with disbelief. The venture manager explained that most of the construction had to be completed before the spring thaw date, because heavy equipment could not be operated on the muskeg once it thawed in June. The EPA was very unlikely to allow summer construction on the fragile muskeg. In addition, a June 1 expected completion date had already been announced publicly by the president of Insitu. No one had to mention the company's unblemished record of completing projects within the allotted time once construction was under way. Because of the cost of interest on funds expended during construction, Insitu had made this its trademark.

Undaunted, Ms. Mist pursued the question of what remained to be done before the three-month construction schedule could begin. A cost engineer explained that the board of directors had taken the position that it would not meet to review the project unless the native claims issues were settled for the pipeline route. Without board approval, a contract could not be let. If the contractor did not arrange for materials delivery to the site by March 1, the start of the project would be delayed. In addition, the board required a minimum of one month for deliberation, and two months each were required to let contracts or deliver materials.

Sparked by the mention of the EPA, another young staffer, Enviro Mann, asked what would happen if the EPA did not allow the spent solution to be pumped back into the mine shaft as planned. The environmental engineer assured him that a waste pond would cost only \$5 million to build. The possibility that recycling of the solution would be required was very remote.


- a. Draw an influence diagram for the total capital cost of the complex in current dollars.

Slightly unsettled by the questions of Ms. Mist and Mr. Mann, the team assigned them the job of developing a better picture of the risks in the capital cost. Mist and Mann interviewed a number of people in the corporation. From the engineering manager, who knew the most about

scaling up chemical processes, they assessed a 75 percent chance that the full-scale plant would work as efficiently as the pilot plant. From the regulatory affairs department, they assessed a 50 percent chance that the spent solution could be put in the mine shaft. There was only a 20 percent chance that recycling would be required. However, if this additional step were required in the process, \$120 million of equipment would be added. The manager of regulatory affairs was uncomfortable about whether the EPA would allow summer construction on the muskeg. He could remember only five winter Alaskan projects that had been delayed until summer. Of these, only one had been allowed to proceed before the September 1 freeze date. For the four events necessary to begin construction, the following probabilities were assessed.

Event	Probability
Native claims issues settled for pipeline route by October 1	.70
Board of Directors approval by November 1	.50
Contract let by January 1	.90
Materials on site by March 1	.25

The probability distribution below was assessed for worker productivity. Ms. Mist noted that labor costs were only 35 percent of the total construction costs.



- b. Draw a probability tree for the total capital cost of the project in current dollars. Label the branches and put in the probabilities.
- c. Write an equation for the cost model to calculate the total capital cost of the project.
- d. Calculate a probability distribution on the total capital cost of the project in current dollars. How do you explain its shape?
- e. Perform a sensitivity analysis to determine the most important risks in the total cost. Calculate the following quantities.
 - The change in total expected cost when each individual variable changes from its lowest to its highest value. This answers the question, “How much difference does this variable make in the expected cost?”
 - The expected change in the standard deviation of the total cost if perfect information were available on each variable. This answers the question, “How much does this variable contribute to the risk?” Why can’t we do ordinary value of information calculations?
- f. What conclusions, insights, and recommendations would you make for risk management and cost control?

7

Corporate Applications of Decision Analysis

In the previous chapters of this book, we used simple examples to illustrate decision analysis techniques. In this chapter, we discuss several of the most important corporate applications of decision analysis—examples that provide some general outlines and points to consider when approaching a problem for the first time.

Because of the tremendous variations encountered in practice, these examples should not be viewed as templates, much less as cookbook formulas for analysis. A decision tree formulation is seldom appropriate for any problem other than the one it was designed for. Thus, the reader should concentrate on the process of understanding and structuring the problem that leads to the tree formulation.

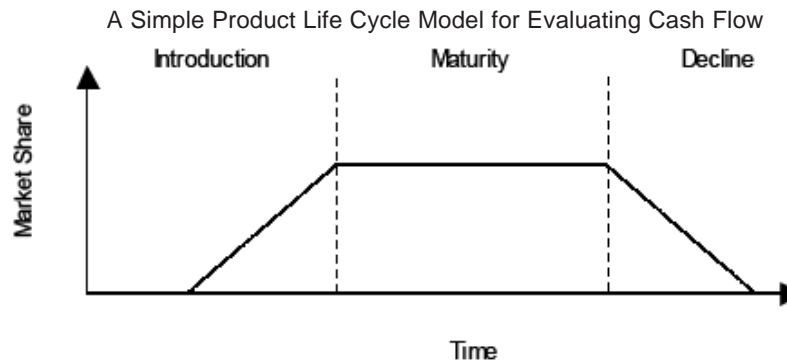
New Product Introduction _____

One of the classic applications of decision analysis is introducing new products to the marketplace. Such problems have a large amount of uncertainty. Will the market accept the new product? How will the competition respond? How long will the product last? What will the margins be? If the product introduction requires large amounts of investment (such as capital for production equipment, promotion, or advertising), the decision may be a source of substantial risk to the company.

For many products, the model to evaluate the cash flow can be based on a simple product life cycle model (Figure 7-1).

This simple model requires only five parameters (year of introduction, length of introduction phase, maturity phase, decline phase, and market share at maturity) and, given the large uncertainties involved, is usually a sufficiently accurate description.

Figure 7-1



To complete the calculation of cash flow, we also need the size and growth rate of the market, fixed and variable costs, capital investment, the cost of ongoing research and development, the margin realized, and the value of possible follow-on business.

One of the most difficult aspects of analyzing new products is determining precisely how to model revenues and costs. Many products are sold like commodities in that prices are determined by some markup or margin over cost, usually the cost of the highest-cost producer. As a new entry to the market, our cost may be high if we have not moved down the learning curve or low if we have the latest, most efficient means of production. The first possibility can sometimes be handled by including a start-up cost. In the second case, we should consider the possibility that our margins decline over time. Specifically, either the competition strives to match our cost or new competitors enter the market and drive out the old high-cost producers, thus reducing prices.

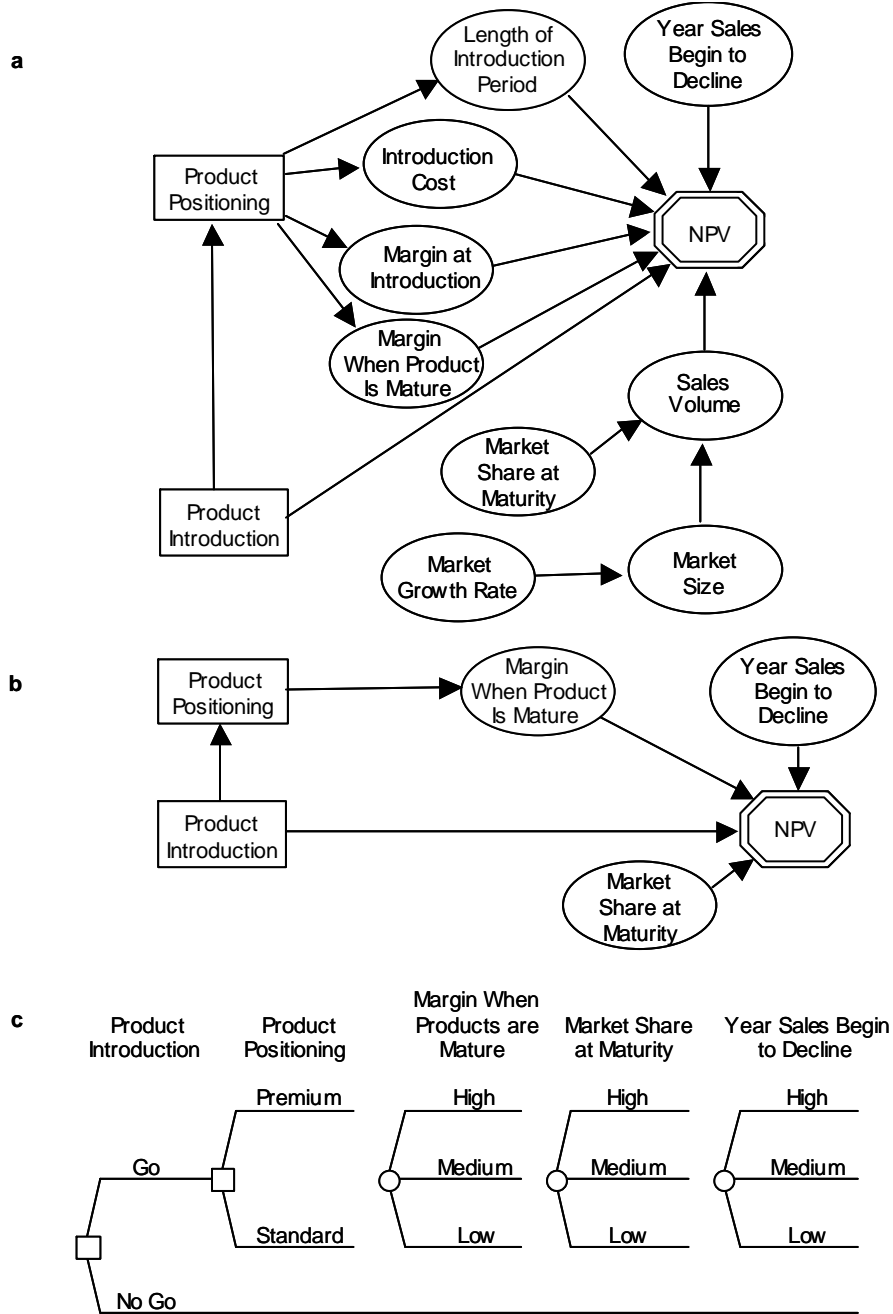
Other types of products have some differentiation from competing products, at least at the time of introduction. In this case, modeling price over time will be quite difficult. One possibility is to make the absolute price (e.g., \$7 each) or relative price (e.g., 20 percent over competition) a decision variable and let market share be an uncertainty, with probabilities depending on the pricing decision. Alternatively, we could make market share a decision (e.g., hold 30 percent of market) and let realized price be an uncertainty, similarly depending on the share decision. In any case, the model should check that margins do not become unreasonably high or low over any extended period.

Two aspects of competition should be considered: (1) What will the competition do over time, perhaps in response to our entry into the market and our pricing? and (2) Is there a possibility of some competitive breakthrough that will completely change the market? We might include these considerations by assessing the price necessary to hold a particular market share given different competitor actions or product introductions.

The influence diagram (Figure 7-2a) shows one way of relating all these factors. This influence diagram is a disguised and simplified

Figure 7-2

Critical Uncertainties and Decisions in a New Product Introduction



version of one used in an actual application. We see that the product introduction affects sales volume (because there are no sales if the product is not introduced) and product positioning. Product positioning in turn influences margins, costs, and some aspects of the product life cycle, which, together with market size, share, and sales volume, determine the net present value (NPV) for the product.

A model was used to capture the relationships in this influence diagram, and sensitivity analysis was then applied to determine the critical uncertainties and decisions. Those uncertainties and decisions are shown in influence diagram form (Figure 7-2b) and in tree form (Figure 7-2c).

The first node is a decision about whether to introduce the product, with the Go option requiring a major capital investment. The second node represents a decision on how the product is positioned relative to existing competitive products, which affects the size of the introduction costs, the length of the growth period, and the initial margins. The third node represents the margins achieved during the maturity period of the life cycle. These margins depend on the product positioning decision. Thus, probabilities for different levels of margin vary according to the product positioning decision; in this case, the company was a market leader, and its initial actions would help set later prices (and margins).

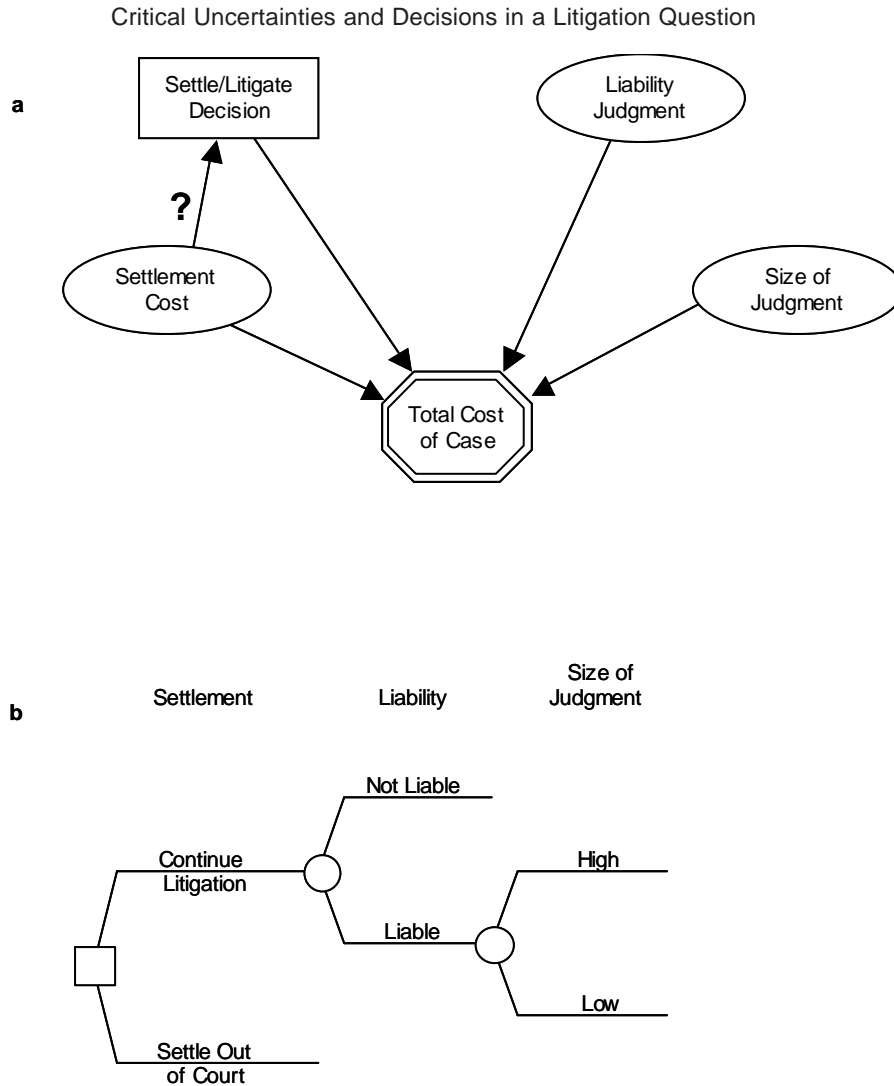
The fourth node represents uncertainty in market share during the maturity phase. The probabilities of this node were independent of the other nodes, since the company and the product were judged important enough that the initial product positioning would not affect the mature market share. The final node is the year the decline phase begins in the product life cycle. The probabilities for this node were also judged probabilistically independent of the previous nodes, implying that in no case is the product so successful (or unsuccessful) as to accelerate (or delay) competitors' introduction of the next generation product.

Litigation Decision Analysis

Litigation is an area in which uncertainty is extremely important and where the stakes may be large enough to pose a serious risk to a company if it loses a suit or provide a significant opportunity if it wins one.

In one common kind of litigation decision, a company is being sued and has the opportunity to settle out of court. Thus, it has to decide on whether to settle the case or to continue in litigation (Figure 7-3). If the company chooses to litigate, there is an uncertainty concerning what the judgment will be and the size of the judgment if the company is found liable. If the company chooses to settle, there is some uncertainty about the final settlement cost. Presumably, however, the plaintiff's counsel has indicated the size of potential settlements. There is a question mark on the arrow between the Settlement Cost node and the Settle/Litigate Decision node to indicate that the arrow should be there if the settlement cost is known and should not be there if it is still an uncertainty.

Figure 7-3



The decision tree formulation for this problem begins with a decision node on whether to litigate or settle. The litigation branch is followed by uncertainties on the verdict and on the size of the judgment if the company is found liable. The Settle Out of Court alternative branch will be followed by an uncertainty in Settlement Cost if the questionable arrow in the influence diagram does not exist.

One implicit assumption in this tree is that this is the company's last opportunity to settle out of court. If this were not the case, there should be another decision about whether to settle out of court later, perhaps after

something has been learned to change the probabilities for liability or size of judgment.

Modeling in litigation tends to be simple. However, it is important to discount future values and to include the cost of the proceedings since large cases tend to last a long time. It is also important to include in all the outcome values the cost or value of setting a precedent.

Other litigation decision problems can be examined with decision analysis. For instance, the value of pretrial work and investigation can be analyzed. The value of this work is either value of information (the trial strategy can be chosen better) or value of control (the probability of a favorable outcome can be increased).

Litigation problems can be difficult for the facilitator. Lawyers are often highly resistant to the notion that someone (the facilitator) is telling them how to do their job. Lawyers, like doctors or polymer chemists or nuclear physicists, are quite right in thinking that their years of training and practice have made them facile with complexities and language to which the ordinary facilitator is not privy.

This impression is perhaps reinforced for lawyers by the special legal monopoly accorded their profession and by their traditional independence from clients in all decisions except those decisive to the outcome of a claim (such as whether to accept a particular settlement offer).

Therefore, it is especially important that a legal decision analysis meet two important goals. First, it should reflect the counsel's best judgment about the law and precedent that bear on the client's prospects. Second, it should forcefully direct that judgment to the value measure of importance to the client.

To meet the first goal, analyses will often be very detailed at the influence diagram stage, reflecting alternative legal theories and perhaps even evidentiary issues bearing on what the judge or jury will see.

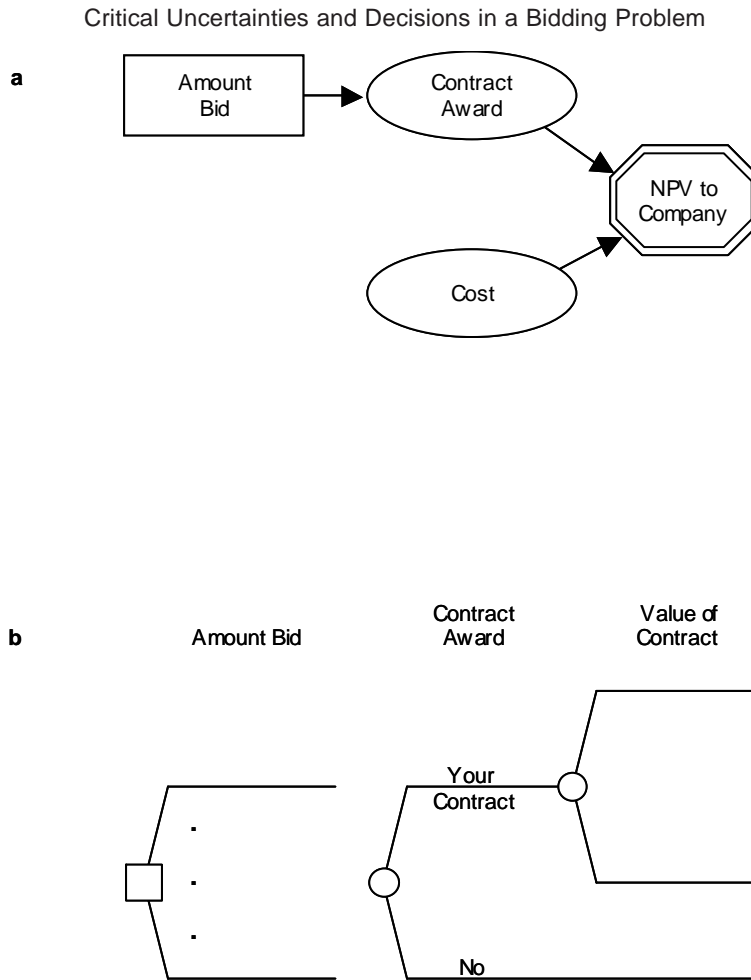
The second goal of forcefully directing the analysis toward the value measure is accomplished by framing the analysis in terms of what the judge or jury will actually consider and base the decision on.

Decision analysis can dramatically improve the quality of communication about legal questions. Without it, the difficulties in communicating about legal complexities and uncertainties often remove effective control from the hands of the president or CEO of the company—a serious condition since the suit may have major implications for the company's future. In addition, explicitly considering risk attitude through decision analysis can enable management to make litigation decisions in a manner consistent with regular business decisions.

Bidding Strategies

Bidding is another area in which uncertainty plays an important role. The difficulties in bidding problems are twofold. First, we are uncertain about how our competitors will bid and thus about whether we will win the bid or

Figure 7-4




not. Second, as the number of serious competitors increases, competition becomes keener and the likelihood of winning decreases. In these highly competitive bidding situations, the winning bid can yield the winner little profit (and even a net loss).

Addressing all the costs and uncertainties associated with losing or winning a particular contract is essential. Assuming there is no cost associated with losing a contract, Figure 7-4 is a good representation of a typical bidding problem. We can look at how some of the quantities described by this tree vary by plotting them (Figure 7-5).

We have superimposed three graphs, with the horizontal axis for each being the size of the bid (the first node in the decision tree). The first line plotted is the probability of winning the contract as a function of the bid size. It

Figure 7-5



obviously decreases as the size of the bid increases. The second line plotted is the expected value of the business, given we win the bid. This line rises, because, as we bid more for the same job, our profit margin increases. The third line is the expected value of the business multiplied by the probability of winning the bid, which is the expected value of the bid. The optimal value of the bid for an expected-value decision-maker is the peak of this third line. This optimal bid maximizes the expected value of the bid. (Note that because of differing vertical scales, there is no relation between the position of the "Optimal Bid" and the bid value where the "Probability of Winning Contract" and "Expected Value Given Bid Is Won" lines cross.)

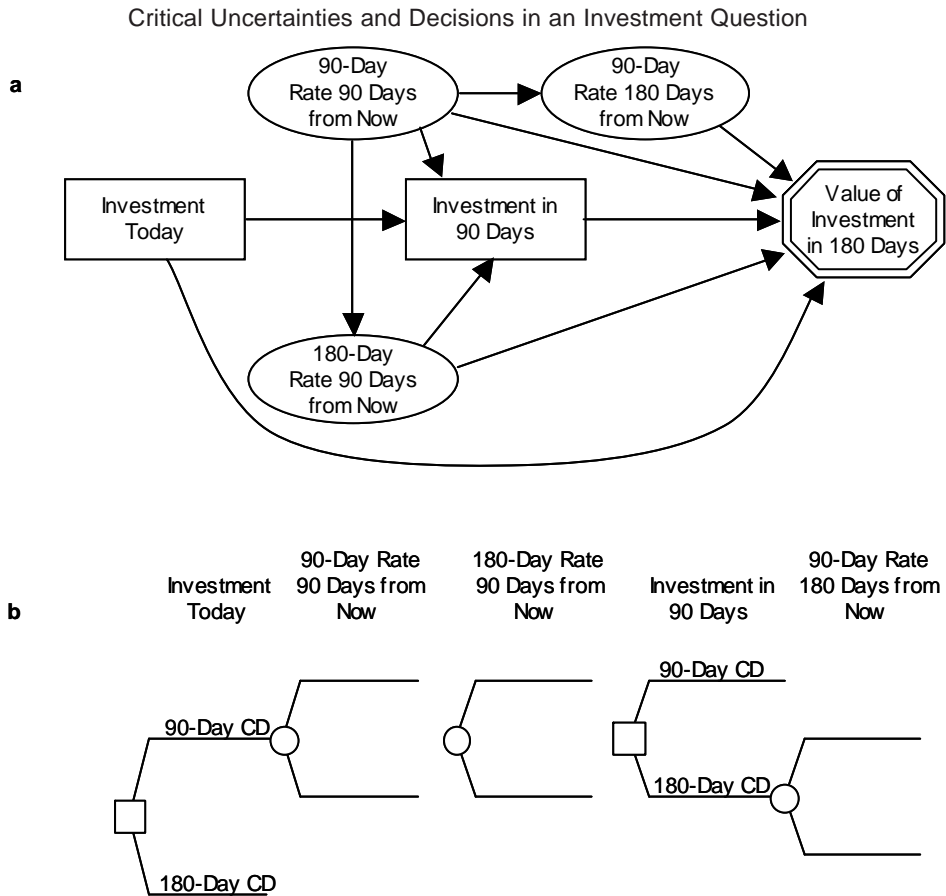
For this type of problem, the bid that maximizes the certain equivalent is frequently somewhat lower than the bid that maximizes the expected value. Risk avoiders prefer to bid lower and increase the probability of getting the business. The higher chances of loss with a lower bid may counteract this tendency, however, and push the risk avoider's bid higher.*

Investment and Investment Rollover Decisions

Many investment decisions are extremely difficult to analyze because of the variety and liquidity of available investment instruments and because of the many points when a decision can be made. As a result, attempts to model

*For an interesting analysis of oil company bids on leases, see E.C. Capen, R.V. Clapp, and W.M. Campbell, "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology* (June 1971): 641-653. Reprinted in *Readings on the Principles and Applications of Decision Analysis*, ed. R.A. Howard and J.E. Matheson, 2 vols. (Menlo Park, California: Strategic Decisions Group, 1984), 2: 805-819.

Figure 7-6



investment decisions in their most general form can produce an unmanageable decision tree. However, once the choice of investments and the time frame have been restricted, decision analysis formulations can be applied.

For instance, banks frequently sell 90- or 180-day certificates of deposit (CDs) to maintain some of their funds in relatively short-term, liquid form. In the CD investment decision, the first decision is whether to buy 90- or 180-day CDs. If we choose 180-day CDs, that decision will take us to the end of the time period considered (Figure 7-6).

However, if 90-day CDs are chosen, we have to decide what to do with the money in 90 days when the CD matures. At that time, we will know what the 90- and 180-day CD rates have gone to, though this is an uncertainty now.

Finally, we need to be able to value the investment decisions made in each scenario. For the initial choice of a 180-day CD and for any choice of a 90-day CD, the value is simply the face value at maturity. However, if we

choose 180-day CDs 90 days from now, they will only be halfway to maturity at the end of the 180-day time horizon. At that point, the 180-day CDs will have 90 days to go and, therefore, will be valued as if they were 90-day CDs. Thus, to value the 180-day CDs bought 90 days from now, we need to know what the 90-day CD rates will be 180 days from now (Figure 7–6).

Options

The stock market would appear to be a natural place to use the tools of decision analysis. Decisions, uncertainty, risk are all there. We will demonstrate how decision analysis can be applied to a very simple choice between buying a stock, buying an option on a stock, and not investing. Then we will give the reasons why analysis of financial investments of this type can not be described in one section of a book.

A stock option is a contract that gives the owner the right, but not the obligation, to buy (“call”) or sell (“put”) the underlying stock at an agreed-on price, the “strike” price. The option can be exercised at any time up to the expiration date (an “American” option) or only on the expiration date itself (a “European” option). In the example below, we will consider a European call option.

When the time comes to exercise this option, the owner knows the price of the stock. If the price of the stock is greater than the strike price, he or she will purchase (exercise the option) at the strike price and take the difference between actual stock price and the strike price as profit—the stock could, in principle, be bought at the strike price and sold immediately at the actual market price. Of course, if the price of the stock is less than the strike price, he or she will refuse to exercise the option.

As an example, take the case in which an investor is interested in buying a stock today at purchase price S and holding it for a year (Figure 7–7a); the uncertainty is in S' , the price of the stock one year from now (Figure 7–7b). Another alternative is to purchase an option today with strike price \$30 at purchase price C ; the value of this option is shown in Figure 7–7b, given that the option would not be exercised if stock price were less than the strike price of \$30. Using the expected values (or certain equivalents), the investor is faced with the tree in Figure 7–7c; now it just a matter of comparing the present value of the outcomes (including transaction and tax costs) with the purchase prices S and C . Note that funds not used in the investments in this simple example are kept in accounts with no return—under the mattress!

To pursue this example a little further, we might adjust the alternatives so that they are approximately of equal magnitude. To do this, we can change the options alternative so that it has an expected value equal to that of the alternative to purchase the stock. The stock purchase alternative had an expected value of \$35.0; the option on one stock with strike price \$30 had an expected value of \$6.6. We change the alternative to an option for 5.27 ($35.0 / 6.6$) stocks; stocks being traded in large lots, there is no problem in having a fractional multiple. This venture is shown in Figure 7–8.

Figure 7-7


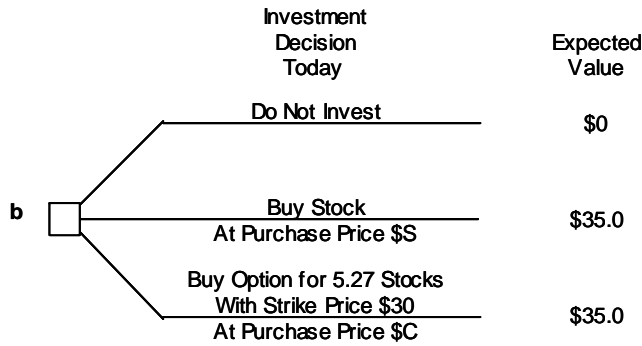



Figure 7-8



What about the purchase prices? For the stock, let us assume that the investor and the market share a common estimate for the future value of the stock and that, for the purposes of this example, the investor and the market both want a 15 percent expected profit during the year. The purchase price of the stock is then $S = \$30.4 = \$35 / 1.15$.


What about the purchase price of the option? In 1973, Black, Scholes, and Merton developed a theory for valuation of this type of asset in an efficient market; for this work, Merton and Scholes won the Nobel prize in 1997. They recognized that portfolios of investments could be constructed which perfectly replicate the payoffs of the option. The value of this replicating portfolio determines a fair market price of the option.

A common application of this methodology to place a market value on a European call option is given by the Black-Scholes formula. This formula has two major assumptions: that the markets are efficient and that stock prices perform "geometric Brownian motion," namely, a random walk superim-

*For an excellent introduction to investment topics, see David G. Luenberger, *Investment Science*, New York: Oxford University Press, 1998.

Figure 7-9

An Investment Decision with an Option for 5.27 Stocks Net of Purchase Price



Investment Decision Today	Net Expected Value
Do Not Invest	\$0
Buy Stock At Purchase Price \$30.4	\$4.6
Buy Option for 5.27 Stocks With Strike Price \$30 At Purchase Price \$24.3	\$10.7

posed on a constant exponential rise. This implies that the uncertainty in S_t , the stock price a year from now, is lognormal (see Problem 7-5.)

Inputs to the Black-Scholes formula for a European call option are current stock price (S), strike price (K), time to expiration of the option (t), the stock's volatility (σ), and the risk-free discount rate (r). The option price is C , where:

$$C = N(d_1)S - N(d_2)Ke^{-rt} \tag{7-1}$$

The "probabilities" $N(d_i)$ are the cumulative normal distribution (in Microsoft Excel, you can calculate $N(d_i)$ as $NORMDIST(d_i, 0, 1, TRUE)$), with

$$d_1 = \left(\ln(S/K) + (r + \sigma^2/2)t \right) / \sigma\sqrt{t} \tag{7-2}$$

$$d_2 = d_1 - \sigma\sqrt{t} \tag{7-3}$$

The example supplies the values $S = \$30.4$, $K = \$30$, and $t = 1$ year. For the

purposes of this illustration, take the risk free discount rate, r , to be 7 percent.

How can we estimate the volatility? The volatility is the standard deviation for the probability distribution for $\ln(S')$, the logarithm of the stock price one year from now. It is a parameter that you can estimate from the past performance of this stock or similar stocks, assuming the future will be like the past. But for this example, we can estimate it from the data given. As shown in Problem 7-5, the volatility is approximately the standard deviation of S' divided by the expected value of S' ; the standard deviation of S' can be approximated as the 10-90 fractile width of the distribution for S' divided by 2.56 (see Chapter 10, Figure 10-17). Inspection of the distribution in Figure 7-7b gives a 10-90 width of approximately 25 and a standard deviation of 9.8 = 25 / 2.56. The expected value is \$35. Thus the volatility is approximately .28 = 9.8 / 35.

Putting this all together, we find the Black-Scholes formula yields an option price of $C = \$4.61$ (70 percent of the expected value of \$6.6) for one stock or a purchase price for an option for 5.27 stocks of $\$24.3 = 5.27 \times \4.61 . The results are shown in Figure 7-9 net of purchase price. The options alternative is much "riskier" (more uncertainty for approximately the same expected value) than the simple stock purchase.

Why is this simple picture too simple? Because there are many decision alternatives that have not been considered in Figure 7-7. The most obvious alternatives concern the timing, size, and nature of investment: the investor can sell the stock or the options during the course of the year or acquire more of the stock or options during the course of the year. In addition, the investor need not restrict investments to this particular stock and can invest in any of myriad other investment possibilities.

Real Options

The idea of real options is that many business decisions create not only a stream of cash flows, but also future investment opportunities that management can later choose whether or not to exercise. As such, these opportunities are valuable and must be considered part of the value delivered by the immediate investment.

Clearly such real options are analogous to stock options, which once purchased by an investor, create a future opportunity to make another investment (namely, to buy the underlying stock).

Introducing this terminology into a decision-analytic framework can stimulate ideas and help uncover hidden issues, possibilities, and sources of value. More than that, it can point to deficiencies in many analyses. Decision analyses all too often tend to focus on uncertainties and up-front decisions; the value of a rich set of downstream decisions (decisions to exercise options) is sometimes slighted. By explicitly thinking about options, we can add an element of quality control to the process.

A topic beyond the scope of this book: From options valuation and from a number of other sources comes the idea of "risk neutral valuation" or "options

valuation.”* In this valuation, future cash flows are always discounted at the *risk-free borrowing rate* rather than the actual borrowing rate (e.g., weighted average cost of capital, WACC) or any other risk-adjusted rate. To the extent that there is a securities market whose risks are correlated with the risks associated with the project, the “probabilities” for future scenarios are deduced from the current market price of these securities. These probabilities are called *risk neutral probabilities*. While they are not the subjective probabilities discussed in this book, they are used as probabilities and take into account the non-diversifiable risks associated with the overall market. This topic is too deep and controversial for further discussion in this book. At the moment, the approach appears to be mostly theoretical and academic. However, the decision facilitator should be able to recognize the topic if it should occur and know when to seek assistance.

A problematic application of the Black-Scholes formula: Sometimes people try to take the analogy to stock options a step further and use the powerful methods developed for stock-option pricing in the real option setting. The essential idea is to apply the Black-Scholes formula to evaluate capital investment and long-term business ventures. Consider a situation in which we can make a modest initial investment, acquiring the right to making a major investment several years from now to enter a new business; R&D is an example of this type of situation. The probability distribution on the value of the new business gives the current value S (expected value) and the volatility σ ; the major investment is K , the strike price; the time until the major decision is made is t , and the risk free discount rate is r . *Voilà*, the value of the venture, to be compared to the modest initial investment.

There are a number of profound problems with this application of the Black-Scholes formula. The formula is formally applicable only under fairly stringent assumptions about the stochastic nature of the underlying stock; these assumptions are almost certainly not met by most business ventures. Furthermore, it implies that the value of the stock is accurately known on the date of exercise, which is definitely not the case for the value of a new business opportunity. It also assumes we have an efficient market that will trade any position on the underlying stock and provide low-interest loans for leveraging positions; again, this is not easy to do for most new business opportunities. From a decision-analysis perspective, the use of the Black-Scholes approach to place a value on real options should be taken as a helpful *qualitative* tool to bring to light the drivers of value in a real option.

In summary concerning real options: Make sure to think carefully and creatively about “downstream” decisions and incorporate them, as appropriate, in the analysis. Sometimes downstream decisions will be built into the model, sometimes into explicit tree/influence diagram structure. If their use is required, use stock option tools (e.g., Black-Scholes formula) with caution and only as an indicator for defining or refining an analysis.

*For an introduction to real options, see Martha Amram and Nalin Kulatilaka, *Real Options: Managing Strategic Investment in an Uncertain World*, Boston: Harvard Business School Press, 1999.

R&D Decisions

Research and development (R&D) is one of the most obvious areas for applying decision analysis in the corporate sector, because R&D decisions usually have large and unavoidable uncertainty. Any evaluation of whether a project justifies its cost must deal with this uncertainty. Though decisions on individual R&D projects usually do not involve large enough costs to pose significant risk to a company, the existence of an adequate portfolio of R&D projects is frequently critically important to the company's continued well-being.

General research without any specific application (also called "blue-sky" research) is difficult to evaluate, no matter what methodology is used. It is usually best to restrict analysis of research to laboratory work with an identified application or applications. An example of research with an identified application could be work on a coating both to reduce corrosion in chemical plant pipes and also to coat nonstick cookware.

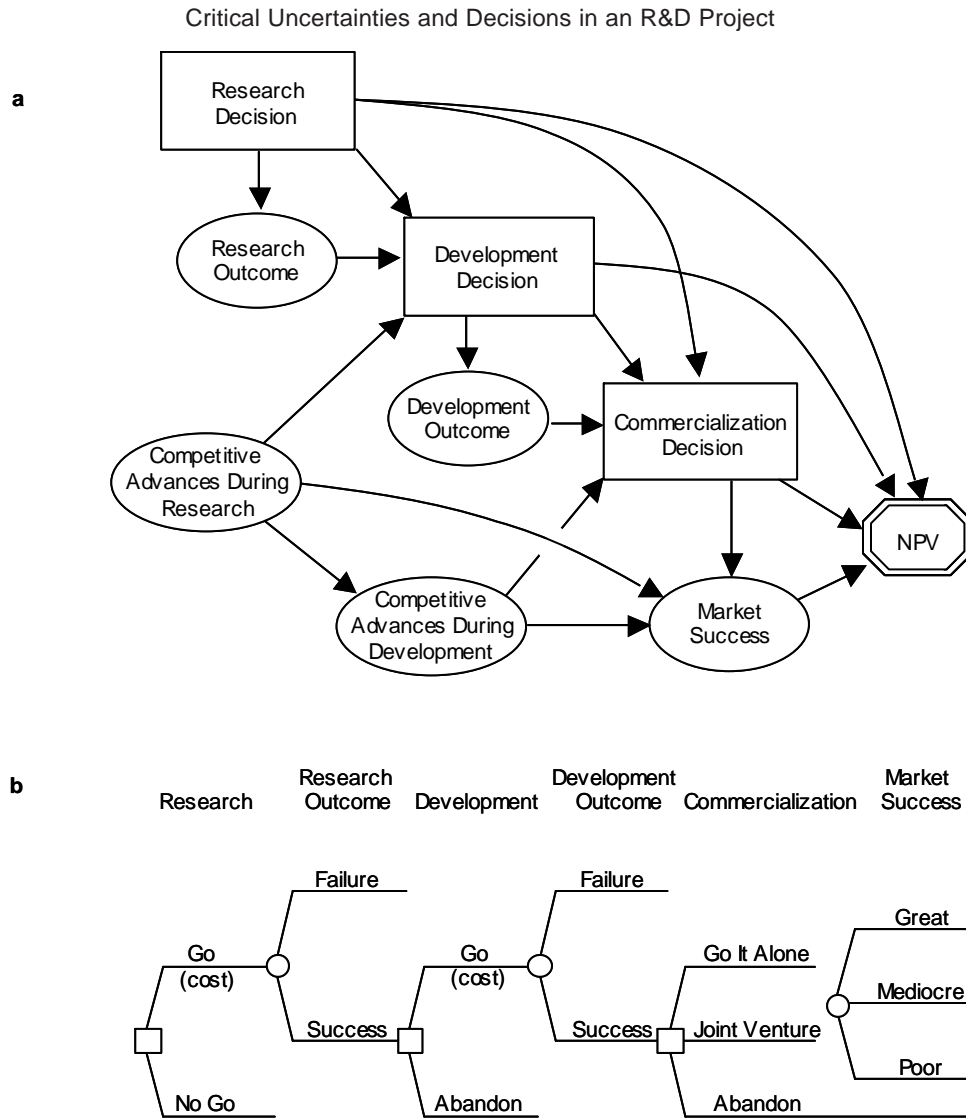
Staged Decisions: In an R&D project where an application has been identified, there are a series of decisions: research, development, scaleup, and commercialization. Luckily, not all these decisions are relevant for every R&D project. For research, we can start or not start the research (or continue or terminate it if it is already in progress). For development, we can abandon development or continue engineering to develop a practical product or process. For scaleup, we can abandon scaleup or build a pilot plant. For commercializing a product, we can abandon commercialization, introduce the product, replace an existing product, or wait until the competition introduces a similar product. For commercializing a process, we can abandon the process, use it ourselves, license it to others, or sell the rights to the process outright.

Pharmaceutical R&D: Pharmaceutical R&D* lends itself well to this type of analysis. Regulation requires a staged set of clinical trials to prove the safety and efficacy of new compounds. These stages provide a natural setting for reviewing the project and making decisions concerning the next stage. Because only 20% of the new compounds that begin clinical trials make it through to registration, it is essential to deal with uncertainty adequately in decision making.

Products and Processes: Products and processes require quite different modeling techniques. Products have a commercial value modeled identically to the new product introductions described at the beginning of this chapter. Process R&D usually targets cost reduction through new catalysts, material substitution, simplification of production steps, and the like. New processes have value when they replace an existing process and lead to more efficient operations; their value usually disappears when they themselves are replaced. So modeling of process improvements is normally modeling cost reductions in an existing business.

*For an example of pharmaceutical R&D analysis, see Paul Sharpe and Tom Keelin, "How SmithKline Beecham Makes Better Resource-Allocation Decisions," *Harvard Business Review* (March-April, 1998): 45-57.


Figure 7-10



Technical Success: A difficulty often arises in defining exactly what it means for research to be technically successful. It may be necessary to model several different levels of technical success, with branches in a chance node corresponding to each possible level. Furthermore, it may be desirable to separate into distinct nodes the various technical hurdles that must be surpassed for overall technical success. These nodes might then be collapsed into a single composite node for entry into the tree.

A typical influence diagram for an R&D problem (Figure 7-10a) shows the different decisions that must be made and the uncertainties that must

Figure 7-11



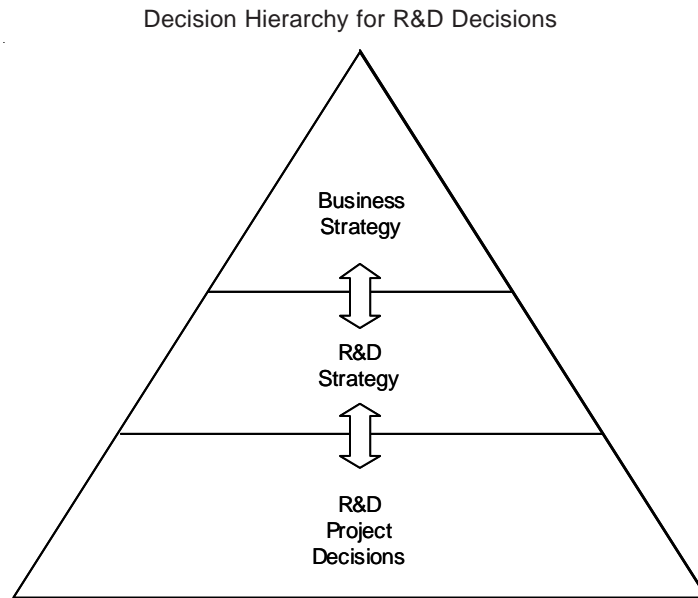
be resolved before an R&D project can be valued. First is the decision of whether or not to conduct the research, followed by an uncertainty on the success or failure of the research. Then follows the decision of whether or not to undertake development of the product, followed by an uncertainty on the success or failure of development. The final decision is of whether or not to commercialize the product, which, in turn, leads to an uncertainty on market success if the product is commercialized. The arrows from each decision to the final net present value show how all the decisions and the market success must be known before the research can be valued.

The influence diagram might have more arrows than shown if, for instance, there were different potential levels of market success depending on the development program chosen (which would mean an arrow from Development Decision to Market Success) or if the level of research affected market success, perhaps because of timing changes in product introduction (which would similarly mean an arrow from Research Decision to Market Success).

In the decision tree for this influence diagram (Figure 7-10b), an initial research decision is followed by a node for the success or failure of the research. (The competitive advances nodes were dropped for simplicity in showing the tree, but should be included in the analysis.) Given research success, there is a development (engineering) decision node and a chance node for the outcome of development. Finally, there is a decision on commercialization, followed by the uncertainty on the market value of the product. At each decision point, there is a cost associated with proceeding.

A typical probability distribution from an R&D decision tree has a vertical line corresponding to the probability of R&D failure, which occurs at a negative value—the cost of performing the unsuccessful research (Figure 7-11). The rest of the distribution corresponds to various outcomes given research success. The distribution shows some probability of outcomes worse

Figure 7-12



than research failure: research success followed by market failure so great that market entry costs are not covered.

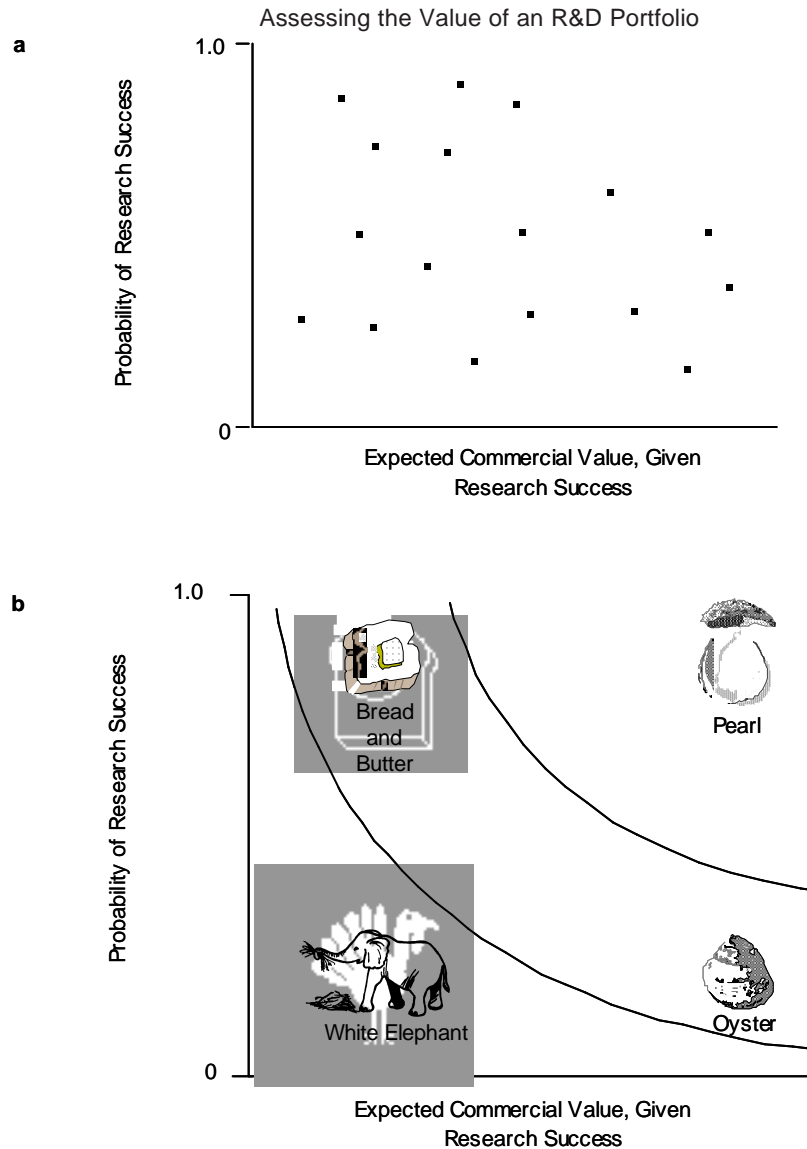
Frequently, one of the main benefits of using decision analysis on R&D projects is improved communication between the research and business departments. In particular, research can use the precise language of probability to communicate its hopes and fears about the technical success of different research projects, and the business side can use probability distributions to communicate its knowledge about future markets in a nonthreatening manner. Without this kind of communication, research could work on projects with little commercial potential, and business could have overly optimistic or pessimistic expectations for research results.

R&D Portfolio

Decision analysis can also help the company manage the portfolio of research projects. The decision hierarchy for R&D decisions (Figure 7-12) involves communication across organizational boundaries. Ideally, R&D strategy should be made in the context of the policy contained in the overall business strategy. However, not all things are possible, and the R&D director needs to work with the business director to make sure that business goals are consistent with the feasibility of the R&D required to achieve those goals.

R&D strategy and R&D project decisions are similarly coupled. Projects should be chosen to further the R&D strategic goals. The R&D goals must be set in light of the probabilities, cost, and time to success of the R&D projects.

Figure 7-13




For the R&D director, the most important project metrics are the probability of research success and the expected value given research success. The product of these two numbers is the expected commercial value of the R&D project, which should be greater than the expected cost of beginning or continuing the research. We can use a graph to display how each research project fits into the portfolio (Figure 7-13a). The graph shows probability of research success versus expected commercial value given research success; each dot represents an R&D project.

The regions of the graph can be characterized according to the value of the projects that fall within them. The curved lines in Figure 7–13b are lines of equal expected project value. (Expected project value is the probability of research success multiplied by the expected commercial value given research success.) The white elephants are low-probability, low-value projects that consume resources but do not contribute much to the portfolio—white elephants were treated with veneration and not used for work. Bread-and-butter projects are low-value but high-probability projects; often routine product or process improvements, they are desirable because of their high probability. However, a portfolio unduly weighted in this area can lead to long-term strategic weakness for a company—nothing new and great is on the way. Oysters are the high-value, low-probability projects that typically represent early research into potential product or process breakthroughs—there is probably not a pearl in the oyster, but if there is...! While these projects are desirable, a portfolio unduly weighted in this area tends to

Figure 7–14

Maximizing Value of R&D Portfolio for Limited R&D Cost Budget

Project	Probability of Technical Success	Expected Commercial Value	Expected Project Value	R&D Cost	Productivity
A	90%	\$60	\$54.0	\$5	10.8
B	90%	\$15	\$13.5	\$10	1.4
C	50%	\$35	\$17.5	\$25	0.7
D	35%	\$300	\$105.0	\$20	5.3
E	35%	\$200	\$70.0	\$25	2.8
F	40%	\$100	\$40.0	\$20	2.0



perform erratically. Pearls (high in value and probability) are the things every R&D director would love to find. Unfortunately, they are rare.

Characterizing R&D projects as in Figure 7–13 may help the R&D director balance the R&D portfolio in terms of consistency of output (Bread and Butter projects give a regular pattern of success, while Oysters succeed only rarely) and long-term strategy (Bread and Butter projects support today's business, Oysters provide for tomorrow's business.)

The output shown in Figure 7–14 can help the R&D director in the difficult and contentious task of allocating limited resources (budget and people). Imagine that R&D projects were analyzed and projects A through F were characterized by Probability of Technical Success, Expected Commercial Value (if technically successful), and expected R&D Cost to perform each project. From these values, we can calculate the Expected Project Value (Probability of Technical Success multiplied by Expected Commercial Value) and Productivity (Expected Project Value divided by R&D Cost)—a “bang for the buck” measure.

If we fund projects in the order of decreasing productivity, we obtain the line in Figure 7–14b showing the most efficient use of R&D funds. If the budget for R&D were \$70 million, we would fund Projects A, D, E, and F; if the budget were increased to \$80 million, we would add B. When the incremental productivity drops below 1, no more projects should be funded (benefit is less than cost), and so Project C should not be funded.

Of course, the R&D director will not blindly use the results of either Figure 7–13 or 7–14. There may be conflicts between meeting the strategic needs of the company and maximizing productivity. There may be constraints imposed by the availability of qualified researchers. Decision analysis can provide tools that help make the portfolio decisions more aligned with the goals of the company and more understandable by those affected by the decisions.

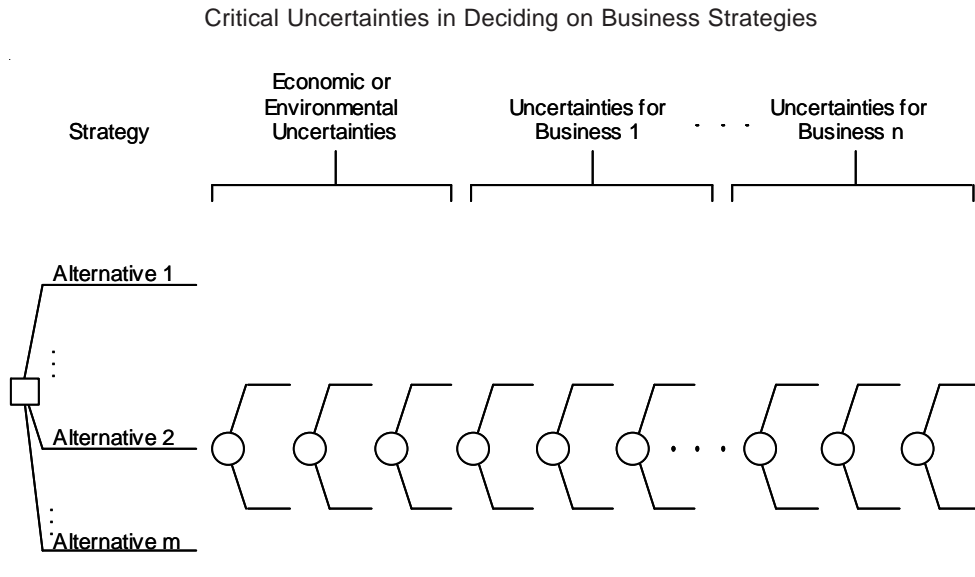
Corporate Strategies and Business Portfolios

Large businesses must often evaluate strategic options that affect many or all of their component businesses or business units. Since these strategies commonly involve decisions for the company as a whole, decisions in the individual business units, uncertainties that affect all the business units, and uncertainties that affect only certain business units, the problem formulation may be very large. Business Portfolio evaluation involves a number of elements that require special attention.*

Business Assessment: Business portfolio decision often take the decision-makers out of their “comfort zone” in terms of the underlying state of knowledge required for a decision. Mergers, acquisitions, and divestitures in a rapidly changing world require deep and current assessments of potential candidates and where they will fit in tomorrow's world. One of the most important (and sometimes lengthy) tasks in the Basis Development steps of

*See, for example, Michael S. Allen, *Business Portfolio Management: Valuation, Risk Assessment, and EVA Strategies*, New York: John Wiley & Sons, 2000.

Figure 7–15



the decision analysis cycle (Figure 6–1) is a thorough assessment of the business environment.

Alternative Generation: Creating decision alternatives for business portfolios is another task that is very important and time-consuming, given the size of the strategy table—there are many fields the company could be in, and many candidates in each field. More important, however, is the need for qualitative thought of high quality in developing alternatives. Before any evaluation, you have to make sure that the alternatives are compelling, coherent, and complete. Will the cultures, core competencies, strategic goals and processes of the business units fit together? How will suppliers, customers, and competitors react? What synergies (or dis-synergies) will occur?

Strategic Flexibility: There is a tendency to think of a strategy is something chosen now and followed forever. The typical decision tree has one decision point at the beginning of the tree. The value of strategic flexibility must be incorporated in the definition of the alternatives, the spirit of the evaluation, and the possibility of inserting a “downstream” decision.

Uncertainty: While decision analysis tools and concepts are exceedingly valuable in evaluating business portfolios, the resulting influence diagram may be very large and the actual decision tree may be difficult to manage. But the tree must be analyzed because the uncertainty characterizing business portfolios may be significant to the corporation.

To construct a decision tree for evaluating strategic alternatives for multiple businesses, we put the decision node for the different alternatives first and then the chance nodes for the uncertainties that affect all

Figure 7-16

Combining Uncertainties of Business Units

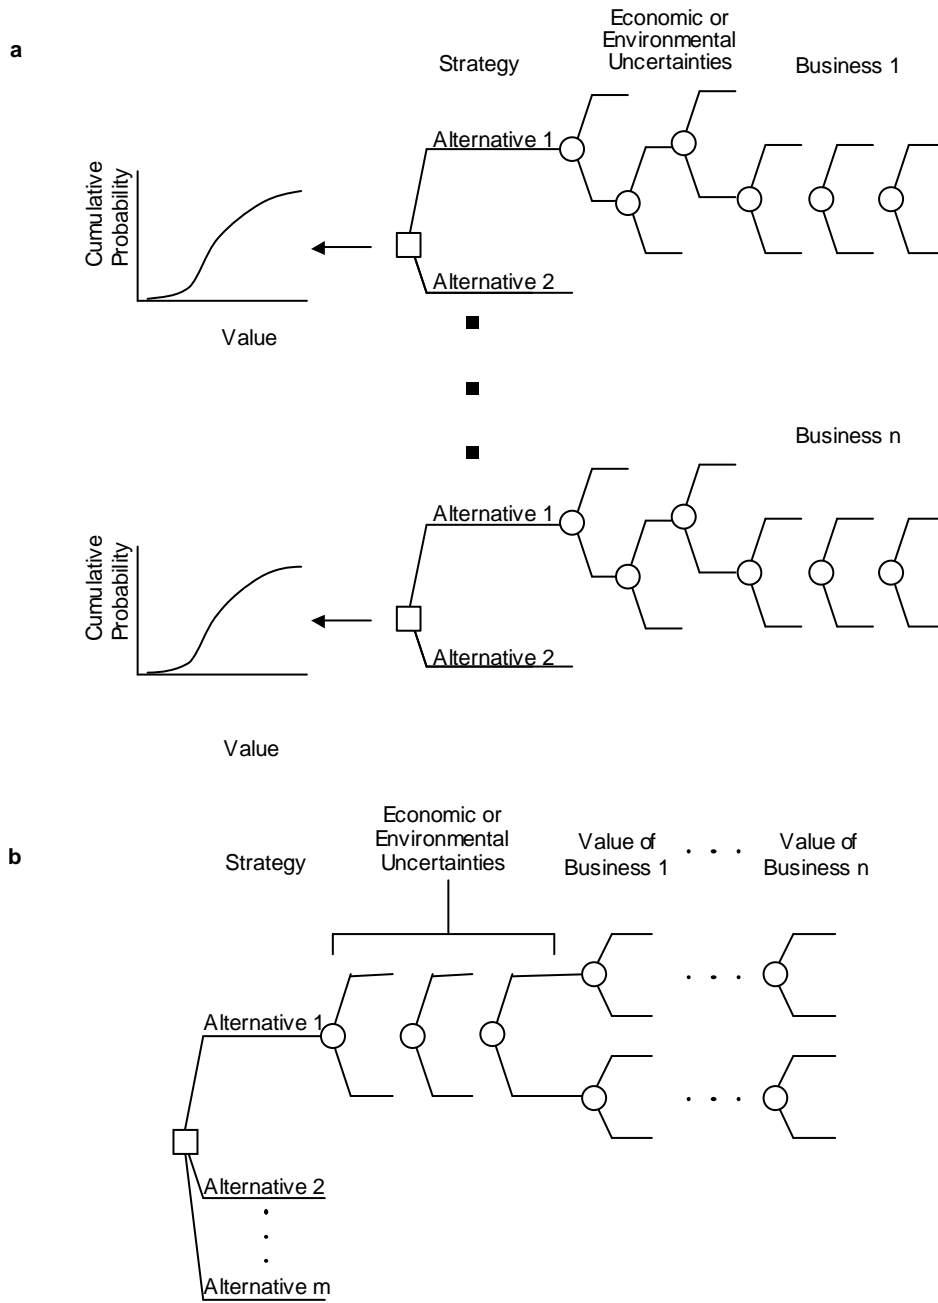
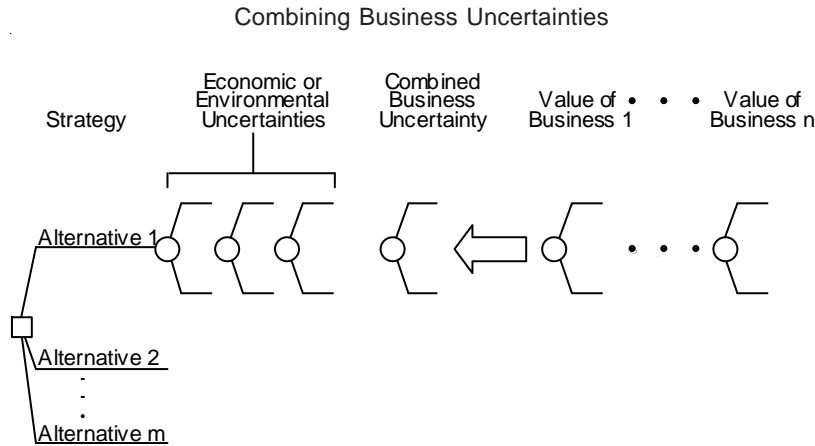


Figure 7-17



businesses—uncertainties that are usually economic or environmental factors, such as energy prices or interest rates (Figure 7-15). Then we have the uncertainties specific to each of the business units. All business-unit uncertainties must appear on each path through the tree so that a model at the end of the tree can calculate the value of each business unit and combine these values to obtain the value of the whole company for each scenario.

As we can see, this tree is impossibly large. However, by using a three-step process, we can reduce the tree to a more manageable size.

1. We construct a tree for each business unit that has the uncertainties particular to that unit. This tree allows us to calculate the probability distribution for each business unit given a particular alternative and set of economic uncertainties (Figure 7-16a). The result of this step is a set of probability distributions. For each business unit, there is a probability distribution:
 - for each alternative
 - given each alternative, for each set of economic or environmental uncertainties.
2. Discretize each of these probability distributions into a two or three branch tree and restructure the tree by stringing these nodes together as in Figure 7-16b.
3. Evaluate this much smaller tree.

While this tree is significantly smaller than the original tree, it is still probably too large to evaluate if the portfolio has more than several business units or if there are many alternatives and economic or environmental uncertainties. Therefore, we can reduce the tree size again by adding a step between step 2 and step 3:

- 2.5 Combine the distributions for the business units for each alternative/environmental scenario to produce a combined

distribution for the business uncertainty for each scenario (Figure 7-17). We can do this by tree evaluation.

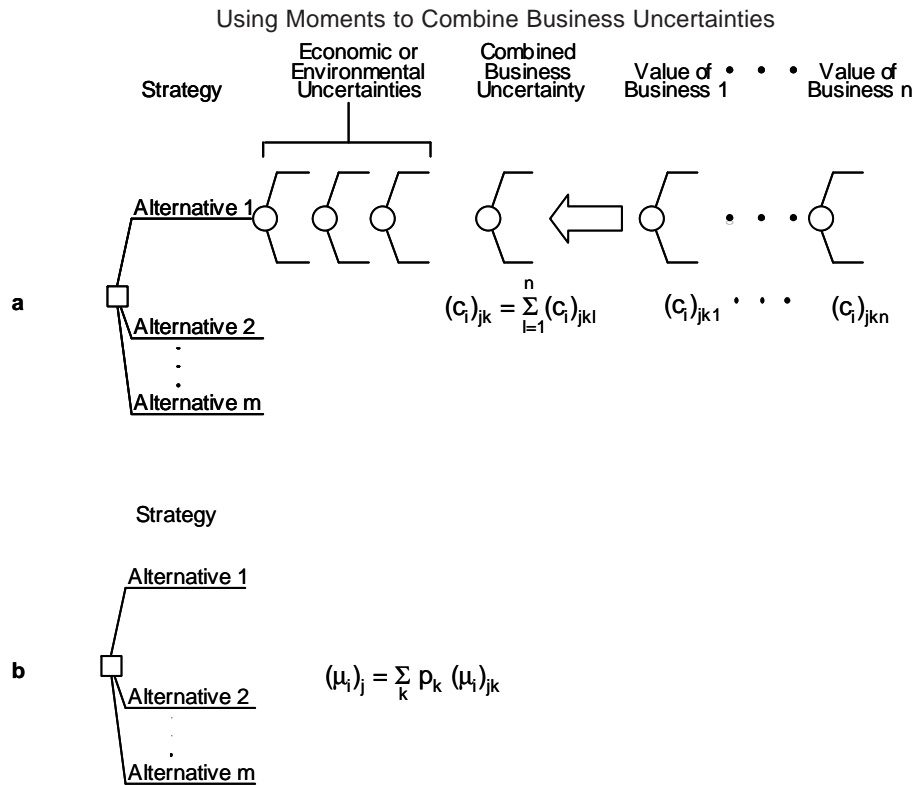
The non-technically inclined may wish to skip the rest of this section!

There is a different way to accomplish this task using moments (defined in Chapter 10) to manipulate probability distributions. This section assumes that there is some easy way to obtain the cumulants of probability distributions, such as a computer program for analyzing decision trees.

The basic tool for the method is to use moments and cumulants to characterize and combine probability distributions. Cumulants are combinations of moments. For example, the first three cumulants are the mean (expected value, first raw moment), the variance (second central moment), and skewness (third central moment.) The two fundamental results we need are:

- If two or more ventures are probabilistically independent and if their values are additive, we can add the cumulants of their probability distributions to obtain the cumulants of the combined venture. For the first two cumulants this reduces to the familiar statement that we can add means and variances if probability distributions are independent and additive.

Figure 7-18



- At an uncertainty node, we can take the expectation of the raw moments describing the probability distribution at the end of each branch to obtain the raw moments of the probability distribution at the node.

For the moment method, the steps are:

1. We construct a tree for each business unit that has the uncertainties particular to that unit. This tree allows us to calculate the probability distribution for each business unit given a particular strategy and set of economic uncertainties (Figure 7–16a). The result of this step is a set of probability distributions. For each business unit, there is a probability distribution:
 - for each alternative
 - given each alternative, for each set of economic or environmental uncertainties.
2. Obtain the cumulants for each of these probability distributions. How many cumulants? Two would give a mean and variance approximation, which may capture most of the problem. Three is probably the practical limit: mean, variance, and skewness. Restructure the tree by replacing the probability distributions with these cumulants, as in Figure 7–18a. Here, j is the alternative, k the environmental scenario, i the order of the cumulant, and l the business unit.
3. Add the cumulants (sum over l) to obtain the cumulants of the combined business distribution, as in Figure 7-18a. The reason we can do this is because the business-unit distributions are independent of each other. We formulated the problem that way.
4. Convert the cumulants c to the raw moments μ , using Equations 10–48.
5. Now take the expectation over the economic or environmental uncertainties, as in Figure 7–18b: Expectation $(\mu_{ij}) = \sum_k p_k (\mu_{ijk})$, where j is the alternative, k is the environmental scenario, and p_k is the probability of that scenario, and i is the order of the raw moment.
6. Now reconstruct a probability distribution that has these raw moments. This is not trivial. The best way is to use a named distribution with the moments as parameters. For instance the normal (Gaussian) probability distribution (Equation 10–50) has two parameters which are the first two moments. The lognormal distribution can be used to fit a distribution for three moments: mean, variance, and skewness—see Problem 10–24.

Though it is cumbersome to evaluate a portfolio of businesses using either of these methods, few other options exist for problems this large. While we could simply truncate the tree to a manageable size, we would lose much of the accuracy, believability, and insight necessary for handling decisions

this large and this important. We could also use a Monte Carlo simulation to estimate the value of the full tree, but again at the cost of some lost insight.

Fortunately, strategic portfolio evaluations do not need to be done often and can be updated fairly easily once they have been done the first time. Considering that the direction of a multibillion-dollar business may be at stake, the effort is well worth it.

Summary

In this chapter, we discussed several of the most common applications of decision analysis to business problems. These examples are not intended as models for other problems, but rather to illustrate the considerations and process that go into structuring an influence diagram and decision tree. Most problems should be individually structured and analyzed to ensure that the analysis is relevant, useful, and insightful.

Problems and Discussion Topics

- 7.1 Huntenpeck Company manufactures typewriters, and Huge company has requested a bid on a contract for 10,000 new typewriters, which cost \$1,000 to manufacture. Huntenpeck would very much like the contract, especially since the publicity would lead to additional sales to Huge's subsidiaries. On the other side of the coin, losing this large contract would introduce a competitor with the potential to cut into Huntenpeck's sales of dictaphones to Huge Company. The secondary impact of winning or losing the contract would be felt for about five years.

- a. Begin structuring the problem and decide what information you need. Make sure you draw an influence diagram.

Huntenpeck decided to make its decision on purely economic factors and to be an expected-value decision-maker.

After heated discussion, Huntenpeck estimated that if it won the contract, it would sell between 1,000 and 1,400 extra typewriters a year to Huge's subsidiaries over the next five years, with equal probabilities of the high and low figures. These typewriters would be sold at \$2,000 (in 1986 dollars). If it lost the contract, it would lose about \$1 million profit a year in dictaphone sales over the next five years. After even more heated discussion, Huntenpeck decided to use a discount rate (time value of money) of 10 percent for constant dollar analysis.

Concerning the contract itself, there seemed to be two principal competitors: Carboncopy and Misprint. No one was sure what their bids would be. However, their typewriters were similar enough to Huntenpeck's that it was certain that Huge would accept the lowest bid. Probability distributions were encoded for both competitors' bids.

Bid Level (\$ per typewriter)	Probability of Competitor Bidding Lower	
	Carboncopy	Misprint
1,200	.01	.00
1,300	.05	.00
1,400	.20	.05
1,500	.50	.45
1,600	.80	1.00
1,700	1.00	1.00

- b. Finish structuring the problem and draw the decision tree for it.
 - c. What is Huntenpeck’s optimal bid (\$ per typewriter)?
 - d. What are the values of information for the crucial uncertainties? Are there circumstances under which the optimal bid would be different given perfect information?
 - e. If Huntenpeck had a risk tolerance of \$10 million, what would its optimal bid be?
- 7.2 The ABC Construction Corporation is being sued for damages as a result of an accident in which the plaintiff fell from a second-floor balcony. The open side of the balcony had been closed off only by a pair of chains held in place by hooks at each end. The plaintiff was leaning against the chains when one of the hooks snapped. He incurred serious injuries in the subsequent fall. In his complaint, the plaintiff charges ABC Construction Corporation with negligence in designing the balcony and the Wisconsin Hook Company with negligence in manufacturing the hooks and asks for \$2 million in damages, including pain and suffering.

ABC has a \$3 million policy covering this kind of claim with United Insurance Company. United’s attorneys have told the claims supervisor that a jury might find ABC negligent in this case. Furthermore, they have said that independent of the jury’s finding about ABC’s negligence, the hook manufacturer could be found negligent. If both defendants are found negligent, each would have to pay 50 percent of the total damages awarded to the plaintiff. If only one is found negligent, that defendant would have to pay the full award. It is believed that the jury would probably award only \$500,000 or \$1 million, but there is some chance of the full \$2 million award.

United Insurance Company has been approached by the plaintiff’s counsel and been given the opportunity to settle out of court for \$500,000. Should the offer be accepted?

- a. Structure the litigate/settle problem as an influence diagram and then draw the decision tree.

b. Assign dollar outcomes to all endpoints of the tree.

Noticing that litigation could produce losses substantially in excess of the \$500,000 settlement offer, the claims supervisor realized that he should quantify the likelihood of the various outcomes. After a long, rigorous discussion of the legal and damage issues, the lawyers provided the claims supervisor with the following probabilities.

There is a 50 percent chance of a jury finding ABC negligent, but only a 30 percent chance of it finding the hook manufacturer negligent. If both defendants are found negligent, there is 1 chance in 5 the full \$2 million will be awarded. The \$500,000 and \$1 million awards are equally likely. If only one defendant is found negligent, on the other hand, the probability of the full \$2 million award is cut in half, and there is a 60 percent chance of the \$500,000 award and a 30 percent chance of the \$1 million award.

c. Find the expected value of each of United Insurance Company's alternatives.

d. Draw the profit distributions for United Insurance Company's alternatives.

The attorneys were surprised to hear that the claims supervisor was going to reject the settlement offer. Because they felt that a little more pretrial discovery could greatly reduce their uncertainty about whether the hook manufacturer would be found negligent, they urged the claims supervisor to briefly delay his decision until they could complete some additional pretrial work.

e. Determine the expected value of perfect information about whether the hook manufacturer will be found negligent.

- 7.3 Hony Pharmaceuticals is a manufacturer engaged in developing and marketing new drugs. The chief research chemist at Hony, Dr. Bing, has informed the president, Mr. Hony, that recent research results have indicated a possible breakthrough to a new drug with wide medical use. Dr. Bing urged an extensive research program to develop the new drug. He estimated that with expenditures of \$100,000 the new drug could be developed at the end of a year's work. When queried by Mr. Hony, Dr. Bing stated that he thought the chances were excellent, "about 8 to 2 odds," that the research group could in fact develop the drug.

Dr. Bing further stated that he had found out that High Drugs, Hony's only competitor for the type of product in question, had recently started developing essentially the same drug. He felt that working independently, there was a 7 out of 10 chance that High would succeed. Mr. Hony was concerned about the possibility that High would be able to develop the drug faster, thus obtaining an advantage in the market, but Dr. Bing assured him that Hony's superior research capability made it certain that by starting development immediately, Hony would succeed in

developing the drug before High. However, Dr. Bing pointed out that if Hony launched its development, succeeded, and marketed the drug, then High, by copying, would get its drug on the market at least as fast as if it had succeeded in an independent development.

Worried about the sales prospects of a drug so costly to develop, Mr. Hony talked to his marketing manager, Mr. Margin, who said the market for the potential new drug depended on the acceptance of the drug by the medical profession and the share of the market Hony could capture. Mr. Hony asked Margin to make future market estimates for different situations, including estimates of future profits (assuming High entered the market shortly after Hony). Margin made the estimates shown below.

Market Condition	Probability	Present Value of Profits (\$)
Large market potential	.1	500,000
Moderate market potential	.6	250,000
Low market potential	.3	80,000

Mr. Hony was somewhat concerned about spending the \$100,000 to develop the drug given such an uncertain market. He returned to Dr. Bing and asked if there was some way to develop the drug more cheaply or to postpone development until the market position was clearer. Dr. Bing said he would prefer his previous suggestion—an orderly research program costing \$100,000—but that an alternative was indeed possible. The alternative plan called for a two-phase research program: an eight-month “low-level” phase costing \$40,000, followed by a four-month “crash” phase costing \$80,000. Dr. Bing did not think this program would change the chances of a successful product development. One advantage of this approach, Dr. Bing added, was that the company would know whether the drug could be developed successfully at the end of the eight-month period. The decision would then be made whether to undertake the crash program. In either case, the cost of introduction and marketing would be \$50,000.

Mr. Hony further consulted Mr. Margin about the possibility that more market information would be available before making the decision to introduce the drug or to complete the crash development program if that strategy were adopted. Mr. Margin stated that without a very expensive market research program, he would have no new information until well after the drug was introduced.

Mr. Hony inquired about the possibility of waiting until High’s drug was on the market and then developing a drug based on a chemical analysis of it. Dr. Bing said this was indeed possible and that such a drug could be developed for \$50,000. However, Mr. Margin was dubious of the value of such an approach, noting that the first drugs out usually got the greater share of the market. He estimated the returns would only

be about 50 percent of those given in the table, but that the cost of introducing the drug would be only \$20,000; for the One-Phase or Two-Phase strategy, the cost of introducing the drug would be \$50,000. Mr. Hony thought briefly about the possibility of going ahead with development after High had failed, but quickly realized that the chance of Hony's success under such circumstances would be much too low to make the investment worthwhile. The chances of a successful development by High after Hony had failed in its development attempt were considered so remote (1 percent) that the "imitate-and-market" strategy was not considered once Hony had failed.

- a. Draw the influence diagram for this case.
- b. Draw the decision tree for this case.
- c. Determine the best decision, assuming risk indifference.
- d. Draw the profit distribution for each alternative.
- e. Determine the best decision, assuming the payoff on each terminal node is certain and that Mr. Hony's risk tolerance is \$100,000.

Consider a clairvoyant so specialized that he could perfectly predict market outcomes but not development or competitive outcomes. Distinguish this case of partial perfect information (perfect information on only one of the uncertain variables) from the case of imperfect information.

- f. Determine the value of market clairvoyance for a risk-indifferent (expected-value) decision-maker.
- g. Determine the value of market clairvoyance for a risk-averse decision-maker using Mr. Hony's utility. Assume that the cost of information is zero.

- 7.4 Mr. Able is the president of Blackgold, a petroleum distribution and marketing company that supplies refined products to a number of customers under long-term contracts at guaranteed prices. Recently, the price Blackgold must pay for petroleum has risen sharply. As a result, Blackgold is faced with a loss of \$480,000 this year because of its long-term contract with a particular customer.

Able has consulted his legal advisers to see if this supply contract might be relaxed in any way, and they have advised him that it contains a clause stating that Blackgold may refuse to supply up to 10 percent of the promised amount because of circumstances beyond its control (a force majeure clause). Able's marketing staff estimates that invoking the clause and selling the contested 10 percent at prevailing market prices would turn a loss of \$480,000 into a net profit of \$900,000.

However, the lawyers caution that the customer's response to Blackgold's invoking the clause is far from certain. The marketing staff claims there

is a small chance that the customer will accept the invocation and agree to pay the higher price for the 10 percent. If the client does not agree to pay the higher price, the lawyers feel it might sue for damages or it might simply decline to press the issue. In either case, Blackgold could then immediately sell the 10 percent on the open market at prevailing prices. A lawsuit would result in one of three possible outcomes: Blackgold loses and pays normal damages of \$1.8 million, Blackgold loses and pays double damages of \$3.6 million, and Blackgold wins. If it loses, it must also pay court costs of \$100,000, but it need not deliver the oil.

- a. Draw the influence diagram for Mr. Able's problem.
- b. Structure the decision tree for Mr. Able's problem.
- c. Assign dollar outcomes to all endpoints of the tree.

Noting that invoking the clause could lead to a profitable outcome, Able has asked his staff to assess the likelihood of the various outcomes. After a great deal of discussion, they report that their best judgment indicates 1 chance in 5 the customer will agree to pay the market price for the contested 10 percent and, if it does not agree, a 50/50 chance it will sue Blackgold for damages. Based on past experience with cases of this type, the lawyers believe there is only a 10 percent chance of Blackgold's winning the lawsuit, an 80 percent chance of losing normal damages, and a 10 percent chance of losing double damages.

- d. Find the expected value of each of Blackgold's alternatives.
- e. Draw the profit distribution for each of Blackgold's alternatives.

Distressed by Mr. Able's persistence in pursuing what they regard was an extremely risky course, the legal and marketing people propose two investigations with the hope of delaying any "reckless" actions. These include \$2,000 for wining and dining the customer's executives to sound out whether they might accept a revocation and \$10,000 to have an objective outside survey team gather information on the possibilities of a lawsuit.

- f. Determine the expected value of perfect information about whether Blackgold's customers would agree to a price increase under the 10 percent clause and about whether they would sue if they did not agree to the increase. In each of these determinations, assume that only one of these two uncertainties can be resolved. How would you advise Able about the proposed studies?

One further investigation that could be conducted is a \$15,000 study by an outside legal firm on the likelihood of the possible outcomes of the lawsuit.

- g. Determine the expected value of perfect information about the outcome of a lawsuit (assuming this is the only uncertainty that can be resolved). How would you advise Able about the above study?
- h. Determine the expected value of perfect information (simultaneously) about all three uncertainties facing Blackgold.
- 7.5 The assumption behind the Black-Scholes formula is that stock prices perform “geometric Brownian motion.” This implies that the probability distribution for the stock price S at some future time, t , is lognormal:

$$f(S) = \frac{1}{\sqrt{2\pi sS}} e^{-\frac{1}{2} \left(\frac{\ln(S)-m}{s} \right)^2}$$

In this formula, m and s are *not* the mean and standard deviation of S , but rather of $\ln(S)$. For stock prices, s is $\sigma\sqrt{t}$, where σ is the stock’s volatility; in this problem, we reserve σ for the definitions below.

The mean, μ , and variance, v (square of standard deviation σ) of S are given by the formulas (see Problem 10–24)

$$\mu = e^{m + \frac{1}{2}s^2}$$

$$v = e^{2m + 2s^2} - \mu^2$$

Show that the volatility, s , is approximately

$$s = \frac{\sigma}{\mu} \left(1 - \frac{1}{4} \left(\frac{\sigma}{\mu} \right)^2 + \dots \right)$$

Hint: Use the approximations for $-1 < x < 1$:

$$\ln(1+x) = x - \frac{1}{2}x^2 + \dots$$

$$(1 \pm x)^n = 1 \pm nx + \dots$$

Part III

Corporate Decision Making

8

A Decision Making Process

In Chapters 6 and 7, we have seen some of the techniques, considerations, and philosophy necessary to conduct an insightful and successful decision analysis of a business problem. In this development, we have used a simplifying frame:

- There is a single decision-maker.
- The decision problem is fairly well defined.
- The decision maker can rely on trusted sources of information.
- There is adequate time to iterate the analysis.
- The decision-maker controls the resources to implement the decision.

For many important decisions, this frame does not describe the reality of modern corporate life.

- Decision making is often delegated to the extent that effective decisions require the cooperation of several decision-makers. Effectively, there are multiple decision-makers.
- The decision problem is not well defined. Each group in the organization sees the problem differently. Not infrequently, values and goals differ among the groups involved.
- Important decisions need cross-organizational teams to address all the facets of the problem. Team members often do not fully understand or trust information developed by team members from other organizations.
- A team approach usually calls for a linear work plan, with little possibility for an open-ended, iterative approach.

- When a decision is made, it is not a foregone conclusion that the chosen alternative will be implemented. Needed resources may lie outside the decision-maker's control. Lack of understanding and commitment by lower-level management may lead to inadequate or postponed implementation.

Needless to say, solutions to all these problems have been developed. Otherwise, decision analysis would have remained an academic curiosity applicable only to a small number of problems in the modern world. This chapter and Chapter 9 will describe techniques for dealing effectively with decision making in the corporate environment.

Decision Projects

Within a large organization, work is commonly organized as a *project*. The project is intended to accomplish a particular goal and then to terminate. Our focus is on designing projects intended to arrive at a decision and lay the groundwork for implementing that decision.

There are three essential processes in a project (Figure 8–1) that are beyond the scope of this text:


- **Creating the Project:** The need for a decision must be identified, communicated, and “sold” within the organization.
- **Managing the Project:** People, budgets, and timelines must be managed.
- **Implementing the Decision:** Implementation of decisions is a topic of vast proportions. Implementation can be simple, but often it involves profound changes and efforts in the organization.

The other two processes, Project Design and Conducting the Decision Process will be discussed below.

Should There be a Project?

Before designing the project, the decision facilitator should first ask whether there should be any decision project at all! The first question is whether the decision-makers are comfortable with decision-making in an open, collaborative forum. Some organizations prefer decision-making behind closed

Figure 8–1



doors—perhaps to preserve secrecy, perhaps as a management tool, perhaps to engage in back-room politics. A group process might not be acceptable for decisions in this type of organization.

Given that an open decision process is acceptable, further questions like the following should be asked:

- Is there a decision to be made?
- Has the decision already been made?
- Is the project a political ploy?

If the answer to any of these questions is “yes,” the decision facilitator would be wise to decline to commit to a decision project. At best, it would be a waste of valuable time and resources. At worst, the dynamics designed into decision processes would lead to “project meltdown” and organizational dissatisfaction with the whole process—and perhaps with the process leaders.

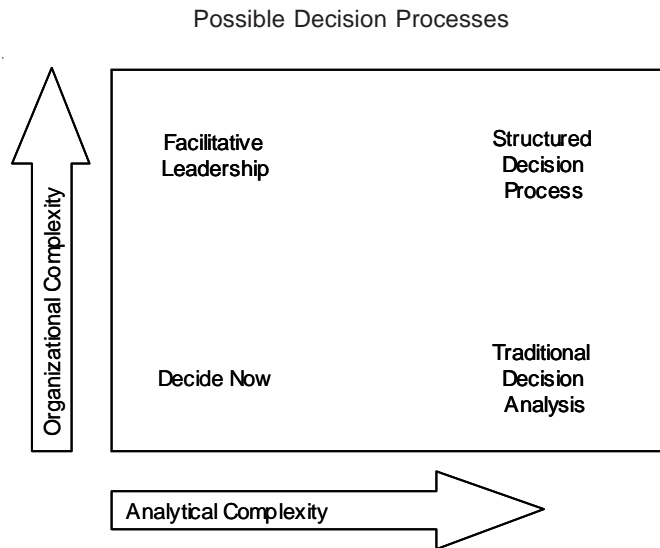
Instead, the decision facilitator might endeavor to discover the underlying needs of the person requesting the process and to suggest ways of meeting that need. Suggestions might range from information-gathering studies to group processes leading to alignment.

Choosing the Decision Process

Given the desirability of a decision process, we need to design the process. Experience has shown that the most important factors in choosing a decision process are organizational complexity and analytical complexity, as shown in Figure 8-2.

Organizational complexity ordinarily rises with the number of people and organizations that have a stake in the decision. Measurement is

Figure 8-2



subjective, but organizational complexity is usually proportional to the number of the following which are true:

- There are many parties involved in the decision.
- They have differences in values, desires, and motivation.
- They have differences in initial convictions.
- They have fundamentally different frames.
- They have very different personalities and competencies.
- They have different degrees of power and resources.

Analytical complexity arises from factors which should be familiar to the reader:

- Uncertainty is important to the decision.
- Many interrelated factors need to be considered.
- Dynamic relationships exist among uncertainties and decisions.
- Multiple alternatives need to be considered.
- Multiple interrelated decision criteria have been proposed.

Possible decision processes are suggested in the four quadrants of Figure 8–2. The decision processes that are appropriate to the four quadrants are quite different, as seen below.

Low organizational and low analytical complexity: Make the decision using the tools ordinarily used by the organization. Most organizations have developed a variety of methods for dealing with routine decisions.

High organizational and low analytical complexity: The methods presented in this book are probably not required to sort out which is the best course of action. In this quadrant, the need is to develop “UAS”—Understanding of the situation, Acceptance of the conclusion, and Support of the decision—among all the parties involved. One way to accomplish this goal is through the process of facilitative leadership: a group of interested parties is created, and a skilled facilitator helps them work through their issues and concerns and arrive at consensus and alignment. Facilitative skills are beyond the scope of this book.

Low organizational and high analytical complexity: The single-decision maker perspective of traditional decision analysis works nicely for this case. The major task is discovering the best alternative; once this is known, it should not be hard to get everyone “on board” and proceed to implementation. Chapter 6 describes a case of this type.

High organizational and high analytical complexity: This is, of course, the subject of this chapter. Analytical complexity requires the tools and skills described in previous chapters of this book. But these tools and skills must be incorporated into a structured decision process which includes elements to address and resolve the people issues.

The remainder of this chapter will deal with the structured decision processes which are appropriate for situations of high organizational and analytical complexity.

Structured Decision Process

A structured decision process is required to achieve success in making a decision in situations of high organizational and high analytical complexity.

Experience has shown that some elements that should be part of the process are:

- The process should be a group process.
- Decision-makers (and potential decision-blockers) should play a role in the process.
- Representatives of the all the parties with a stake in the decision should play a role in the process.
- Someone with analytical skills should play a role in the process.
- Constant communication is essential both in the group(s) and to the rest of the organization.
- Predetermined deliverables and milestones are needed to help the group through unfamiliar territory.

The form of the structured decision process will vary. For instance, the process used to decide on a major investment in an unfamiliar technology will be different from the process used for R&D portfolio management. For the remainder of this chapter, we will focus on the Dialog Decision Process, a process that has been developed to deal with major, one-of-a-kind decisions that cut across organizations within a company. Many other decision processes can be developed by selecting and modifying tools from this process.

Dialog Decision Process ---

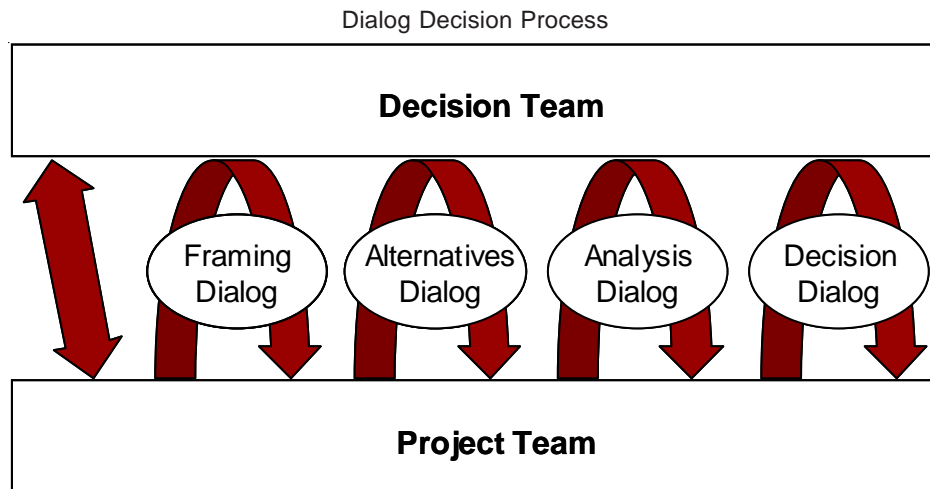
Many decisions in the modern corporation are made in a team context. A cross-organizational or cross-functional team is formed to bring varied expertise together to address opportunities or problems that face the corporation. This team develops a solution that capitalizes on the opportunity or solves the problem. The solution is then presented to a management team for approval.

One of the problems with this type of team-based decision making is that there is not enough formal communication. The decision-makers commission a team to work on the problem. The team analyzes the situation (sometimes in parallel subteams), creates a solution, and at the end presents this solution to the decision-makers. There is no formal communication between decision-makers and project team during the process, only at the beginning and the end. All too often, the project team presents only a single alternative as “the answer” to the problem. At this point, the decision-makers can approve, disapprove, or send the problem back for more work.

What can go wrong in this approach? Working on the wrong problem, choosing alternatives because they are not controversial, and proposing a solution that no one outside the team is committed to.

The process encourages advocacy rather than collaboration—“the answer” needs to be right the first time, otherwise the project team has its work disapproved or sent back for repairs. This encourages developing conventional solutions and hiding any problems with the solution. The decision-

Figure 8-3



makers have only a critical role (“prove it”) and little opportunity to use their collective wisdom and experience to collaborate in creating a great solution.

A process has been developed to bring the power of decision analysis to team-based decision making (Figure 8-3). It has been named the Dialog Decision Process because of its emphasis on systematic dialog between the two groups: the team that makes the decisions, and the team that develops the basis for the decision. This process grew out of the marriage of strategic planning and decision analysis first at SRI International and later at Strategic Decisions Group (SDG) in the late 1970s and early 1980s. A version of the process was christened “Dialog Decision Process” at GM during the late 1980s and early 1990s and was adopted as GM’s decision making process.*

An essential part of the Dialog Decision Process is four formal dialogs during the course of the project. The dialogs are in response to the following questions:

- Framing: do we all see the same problem and see it in the same way?
- Alternatives: have we identified a good set of alternatives to choose from?
- Analysis: what have we learned about the alternatives?
- Decision: what do we choose to do?

The decision dialog does not always occur as a formal dialog, especially if the final decision needs to be made by decision-makers higher in the organization than the decision team.

The decision team is composed of the people who allocate the resources and who represent those who have a stake in the decision. Decision team

*The role of this process at GM has been described by Vincent P. Barabba, *Meeting of the Minds: Creating the Market-Based Enterprise*, Boston: Harvard Business School Press, 1995.

members also tend to be people who have valuable experience, broad understanding, and little time.

The project team is composed of people with profound knowledge and a stake in implementing the decisions. They tend to be people with more time for thought, gathering information, and formulating strategy.

“Dialog” may appear to be a pretentious word for a meeting, but the emphasis really is on sharing of information, not just presentation of material. Part of the dialog is between the decision team and the project team: The project team presents the information it has developed and the decision team offers direction on what more needs to be done. Equally important is the dialog within the decision team itself: The decision-makers share their insights and concerns long before the moment the decision need be made.

To illustrate the concepts of this chapter, we will take the fictitious example of FoodMachines, a rather small company producing food processing equipment. FoodMachines was considering producing and marketing a new line of can opener. This was a new product line for the company, and involved a substantial investment. Stakeholders in the product were North American and European marketing, distribution, and sales, product design, engineering, and manufacturing. The company decided to use the Dialog Decision Process to deal with the organizational complexity of the problem as well as the uncertainties surrounding the introduction of a new product.

This example combines elements from many different cases. The level of detail has been kept at a very simple level. The example is intended to illustrate the process, not the actual case.

Framing Dialog

What precisely is framing? Framing is making sure the right people are treating the right problem from the right perspective. There are three key dimensions to framing a problem well: purpose, scoping and perspective.

When a team begins to work on a decision, its members rarely agree on what has to be done. A clear purpose must be established and agreed to before the team begins to work.

The scoping dimension of framing is usually straightforward and easily understood. It establishes the boundaries of the problem, the dimensions of the solution, the sources of information that will help determine the solution, and the criterion to be used in the choice.

The perspective dimension of framing is much harder to describe. We all have a method of dealing with the flood of information presented to us. One analogy is that each person has established a filter: some information is classified as irrelevant and discarded early in the perception process. Another analogy is a picture frame: we place a frame around a problem, excluding some things from the picture, including others. The engineer focuses on the features of a new product, with little thought for the benefits the marketer may see; neither engineer nor marketer may be sensitive to the interplay of investment, price, variable cost, and volume that is important to the financial analyst. When a cross-functional, multidisciplinary problem

arises, the filter must be revised and the frame enlarged. It is impossible for a team to think strategically or to commit to implementation unless all of the members share the same frame.

Success in framing has proved to be critical for decision quality. Studies and analyses that have failed to achieve acceptance or implementation often prove to have failed in framing the problem correctly. Perhaps it is easiest to describe the purpose of framing in terms of the following challenges to the team:

- Develop a shared understanding of the opportunity or challenge being addressed.
- Create an awareness of the different perspectives of the group members and expand the thinking of each individual in the group.
- Create a respect for the legitimacy and importance of others' perspectives.
- Surface unstated assumptions that could affect the project.
- Explicitly formulate and communicate the problem to be solved.

Warning: During the preparation for the framing dialog, there is an “opening” of perspectives, a creation of new visions, a discovery of new possibilities. This can be difficult for members of the team who feel “the answer is obvious” and “let’s get on with it!” These people must be led along carefully.

The principal processes and deliverables of the framing dialog are:

Issue Raising: Perhaps the most straightforward way to start the framing task is free-form issue raising from the project team, the decision team, and from any other important stakeholders. Issue raising is usually a group process in which team members state whatever comes to mind as being important or of concern to the decision. Issue raising in a group has three purposes: it surfaces much of the material to be dealt with during the process; it exposes the members of the team to each other’s frames; it develops a sense of team ownership of the decision to be addressed.

The list of issues is usually not a deliverable of the process. Rather, once the issues have been recorded, they can be sorted into categories by their underlying content: decisions, uncertainties, values, or process issues. This information can be used in developing the deliverables for the framing dialog. The issues can also be used later to check that the analysis treats everything that was thought important.

Project Vision Statement: The project vision statement is developed from the answer to four simple questions about the project: What are we going to do? Why are we doing this? How will we know if we are successful? How could we fail? When, in the course of the project, the team begins to wander off course, this statement will help refocus efforts. (Remember that the vision statement applies to the decision process, not to the outcome of the decision that will be made at the end of the process.)

The FoodMachines project team developed a vision statement that incorporated the following points:

- *The team was going to investigate the attractiveness of the new product (can opener) using existing manufacturing and distribu-*

tion channels.

- The current product line was growing weak and business projections did not look good. The new product would stimulate distribution channels and would help rejuvenate the product line.
- Project success could be measured by the endorsement of the new product (assuming the financial projections looked good) by the sales and distribution managers.
- Failure in the project could occur if sales, distribution and manufacturing were not closely involved in decisions concerning features and pricing.

Decision Hierarchy: Nothing is quite as wasteful as finding an elegant solution to the wrong problem. The decision hierarchy (Figure 8–4) is a clear way to establish the boundaries of the problem.

During framing, a list of policy decisions should be identified. Strategy decisions are usually dealt with by the strategy tables described in Chapter 6. Although tactics are normally not dealt with explicitly, the concept is very important when simplifying the strategy table—when a column is dropped from the table, it is organizationally important to acknowledge that, although it is not a decision that will be dealt with at the moment, it will be important later on.

The FoodMachines project team identified several policy decisions that were relevant to the new product decision:

- Existing distribution channels must be used.
- Product manufacturing should utilize only existing, idle capacity.
- FoodMachines’ cost of capital is the discount rate for the analysis.

Figure 8–4

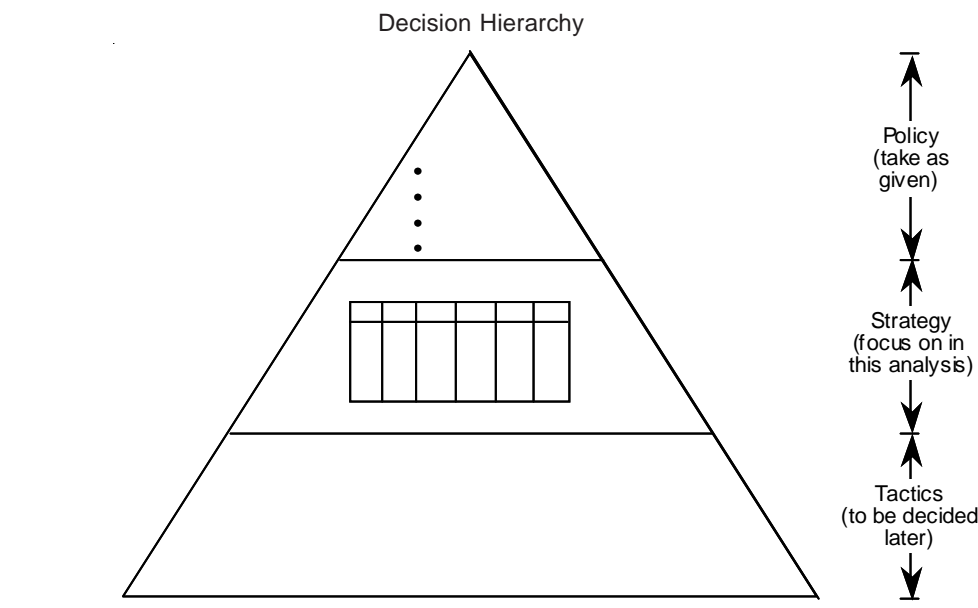


Figure 8-5

FoodMachines' Strategy Table

Product	Distribution	Price
A	N. America	\$12/unit
B	Europe	\$15/unit
A and B	N. America and Europe	\$17/unit

- *The product should not be marketed outside North America or Europe.*
- *Investment should not exceed \$40 million.*

The project team was not sure about the last two policy statements, but decided to write them down and test them with the decision team at the framing dialog.

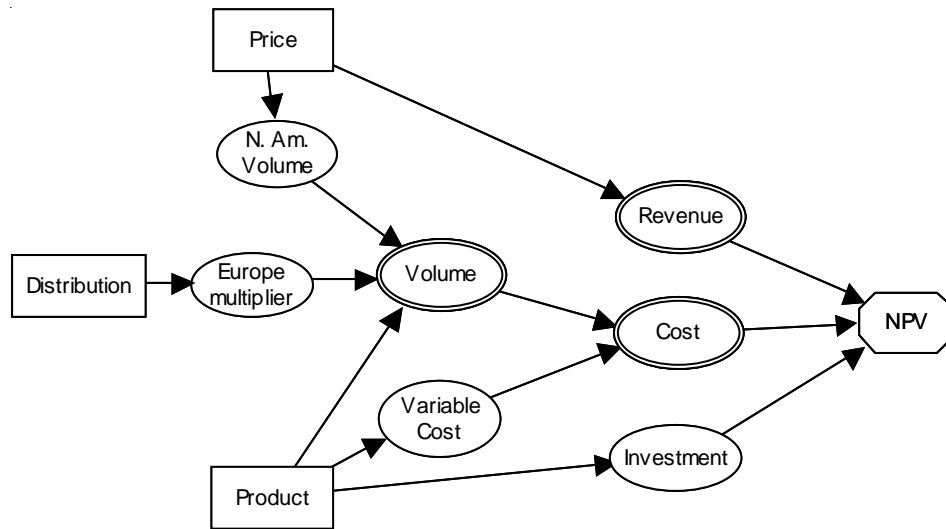
Strategy Table: The strategy table is a tool for sorting through a complex set of decisions (Chapter 6). It lets the decision-maker keep track of many decisions simultaneously. It helps separate the important from the unimportant and the strategic from the tactical. It also provides a framework for coordinating a set of decisions into a coherent strategy. During framing, the task is to develop the skeleton of the strategy table: the decision areas (column headers) and some of the principal choices in these areas (entries in the columns). Strategy tables for complex corporate decisions may have ten to one hundred columns at this stage.

The project team developed a strategy table (Figure 8-5) to describe the possible courses of action open to the company. First, during issue raising, the team discovered that two models of can opener were possible, A and B; B was an upscale version of A and required considerable additional investment. Second, distribution could be in North America or Europe, or both. Pricing could be set at various levels and were effectively fixed over the short life cycle of the product.

Influence Diagram: Uncertainty makes decision-making difficult. In most cases, a good outcome cannot be guaranteed. The quality of decision-making depends in part on the treatment of uncertainty. The influence diagram is a tool that keeps track of decisions, uncertainties, and their interrelationships. Influence diagrams for complex corporate decisions will be

Figure 8-6

FoodMachines' Influence Diagram



much larger than the one shown in Figure 8-6. The influence diagram is very important for developing the project team's understanding of the problem and in managing tasks. However, it is rare that anyone outside the project team will need to invest the time to understand a large, complex influence diagram.

The influence diagram developed (Figure 8-6) showed some modeling decisions based on the team's understanding of the market and of the company. The principal market was judged to be in North America, and North American Volume was entered as an uncertainty. European distribution system was smaller and less developed, and people were most comfortable estimating European Volume as a multiple of North American Volume.

Decision Criterion: Before a decision can be made, the criterion by which the decision will be made must be established and agreed to. As discussed in Chapter 3, the principal criterion for most decisions should be net present value (NPV) of cash flow over the time horizon of the project. Part of the task of the framing phase is to surface and discuss the inadequacy of management goals such as market share or production throughput as criteria for decision-making.

The project team had decided to use NPV over the five year life of the product. Manufacturing representatives on the team had been thinking that their task was to utilize idle capacity. Distribution representatives had been concerned about the health of the distribution network. But the team came to realize that first it had to find the right, profitable product.

Typical framing dialogs tend to center around the deliverables. Are the policy decisions correct? What is the boundary line between policy and strategy? Can the project team challenge any of the policies? Has the project team taken on too broad (or too narrow) a challenge in the strategy table? Is the decision criterion good? Are there relevant constraints, either organizationally or financially?

The decision team agreed that the project team had framed the project well. They agreed to let the project team violate some of the policy decisions in trying out alternatives. The decision team had a long discussion on the health of the product line and the distribution system—some thought things were going well, some thought that there were growing problems. The information shared did not solve any problems, but it did make some of the decision-makers aware of the importance of this opportunity.

Alternatives Dialog

The framing dialog defines the dimensions of the problem; the alternatives dialog defines the dimensions of the solution to the problem. The framing dialog creates a team with a shared frame; the alternatives dialog proposes different approaches within this shared frame.

The alternatives phase provides the following challenges to the team:

- Find creative, fresh alternatives that go beyond variations of “business as usual.”
- Do not be satisfied with a few look-alike alternatives. Find significantly different alternatives that cover the complete range of possibilities.
- Do not squander resources on evaluating undoable, unacceptable alternatives.
- Challenge the common perception of what is acceptable and what is not, what is possible and what is not.
- Look at the problem from a corporate and stockholder perspective. Look to the long term.
- Do not lose mutual respect or the common, shared frame.
- Be enthusiastic about and energetic with all the alternatives.

Warning: During the preparation for the framing and alternatives dialog, there has been an “opening” of perspectives, a creation of new visions, a discovery of new possibilities. This can be difficult for members of the team who felt “the answer is obvious” and “let’s get on with it!” However, many of these people are beginning to get excited by the new perspectives. Prepare these people for the “closing down” which will begin in the analysis dialog.

The principal deliverable for the alternatives dialog is a set of three to five significantly different alternatives, each with:

Theme: A one- to four- word name that characterizes the alternative. This name should be catchy and descriptive—it will be used repeatedly and will come to represent the alternative in everyone’s mind.

Path through the strategy table: A path through the strategy table makes the alternative actionable—a choice in each decision area. The path through the strategy table should be all that is required to instruct the right people in the organization to begin implementation. Of course, implementation will require many tactical decisions, but the direction should be set in the strategy table.

Rationale: The rationale is a brief statement of why the alternative is coherent, compelling, and complete. The purpose of the rationale is to enhance the quality of creative thought and of communication.

- **Coherent:** The choices made in each decision area are consistent with each other and with the strategy theme. Each choice is aimed to achieve the same overall goal.
- **Compelling:** The theme and the choices made in each decision area come together in a way that can motivate people, that can generate a movement toward successful implementation.
- **Complete:** The alternative includes all the elements needed for beginning a successful implementation.

The project team worked with the strategy table and brainstormed a number of strategy themes. It then divided into subgroups, each developing one of these themes in terms of a path through the strategy table and a rationale for that path (Figure 8–7). An alternative called momentum was developed which aimed at minimizing investment (one product) and marketing expenses (only the larger North American market) with a relatively high price. Another alternative was much more aggressive, developing two products, aiming for

Figure 8–7

FoodMachine's Alternatives

Strategy Theme	Product	Distribution	Price
Momentum	A	N. America	\$12/unit
Aggressive	B	Europe	\$15/unit
	A and B	N. America and Europe	\$17/unit

market share by lowering price, and marketing to both North America and Europe.

During the alternatives dialog, the decision team indicated that two other alternatives (not shown in Figure 8-7) should not be pursued further. One of these alternatives involved a joint venture and was rejected as against current company policy. The other rejected alternative involved distribution through the company's Far East distribution channels (an alternative that tested a policy decision). The head of Far East marketing indicated that the product was not well suited to that market.

Analysis Dialog

At the analysis dialog, the decision team is presented with a quantitative evaluation of the alternatives. There is a wealth of material that can be presented at this meeting, and the project team should think carefully about what it wishes to present.

It would be a mistake to think of the analysis dialog as the time when numbers and a computer model are used to find the right answer to a problem. The analysis dialog has a much more ambitious set of goals: understanding, insight, and communication.


- Better understanding is achieved because the model keeps track of complex relationships and uncertainties, something that cannot be done in our heads.
- Organized thought about data, relationships, and results yields insight into the benefits and drawbacks of alternatives.
- Quantitative analysis can be used to communicate to the decision-makers and others the reason why one alternative is preferable—the decision-makers obtain the understanding and insight needed for confidence beyond a gut feel, and they have the financial projections to justify the decision.

The principal challenges in the analysis phase are:

- Never lose sight of the purpose of the model. Start simply and build in complexity by detailing factors that matter to the decision. Maintain simplicity by excluding factors that do not matter to the decision.
- Model to the level of complexity that is required for organizational buy-in, remembering that buy-in is ultimately achieved by insight, not detail.
- Produce insights, not answers. Look to the analysis for insight, not for a machine-made, machine-blessed solution.
- Quantify “intangibles,” “ghost stories,” and issues people avoid talking about.

Warning: During the preparation for the analysis dialog, the “closing down” process begins as alternatives are discarded and as hard reality shows that exciting possibilities are not realistic. This can be difficult for many members of the team. Let the results and insights speak for themselves and gradually convince those unwilling to give up favorite alternatives.

Figure 8–8



The deliverable for the analysis dialog is the same as for any good decision analysis: insight to help the decision-makers choose among alternatives. Analytic results should be presented only insofar as they illustrate and support these insights. Possible analytic results include:

Base Case Evaluation of Alternatives: What does the deterministic model show for the NPV of each alternative? If the project had to stop at this point, it might be reasonable to choose the alternative with the highest NPV.

The project team found that the base case NPV for the momentum alternative was \$21 million while the base case for the aggressive alternative was \$18 million.

Sources of Value: Why is one alternative better than another in the base case? The identification of sources of value provides basic insight into the nature of the decision problem. The “waterfall” chart in Figure 8–8 shows how the various changes in going from Momentum to Aggressive affect the NPV. (Note that changes happen cumulatively from left to right in this chart.)

The team identified that, in the price reduction in Product A (required to enter the European market) was not offset by a sufficient volume increase, and therefore led to a loss in NPV; this loss, however, was almost completely compensated for by profits from adding the European market.

Product B cannibalizes volume from Product A and does not generate enough profit to compensate for this lost volume.

Deterministic Sensitivity to Uncertainty: This sensitivity identifies the uncertainties that have the largest effect on the alternative and which consequently should appear in the decision tree. These are also the uncertainties for which there is often a high value of information.

The “tornado” charts for the Momentum and Aggressive alternatives are shown in Figure 8–9. Price, cost, volume and investment for Product B were estimated as a fraction of those for Product A. As had been expected, the uncertainty in the size of the North American market dominated. However, even with the worst outcome, the NPV was still positive, which was an encouraging sign. Uncertainty in unit cost and on how much volume B would take from A were somewhat less important, as was the uncertainty on the size of the European market. The results on uncertainty in investment were important because people were concerned about the lack of experience in manufacturing with this type of product. (Uncertainty in investment is typically low in the tornado chart.)

Deterministic Sensitivity to Decisions: It is often useful to perform a deterministic sensitivity by varying each of the entries in the strategy table over the range of options in that column. The columns with the largest swing are those decision areas with a high value of control. Although many of the combinations of choices do not make sense, the output of this sensitivity has often suggested modifications to alternatives that has substantially increased the value of the alternatives.

Figure 8–10 shows this sensitivity for the momentum alternative. This plot showed that not much value was created by introducing both products A and B in North America. It did make a lot of sense, even without product B, to add the European market. The price sensitivity included the increase in volume caused

Figure 8–9


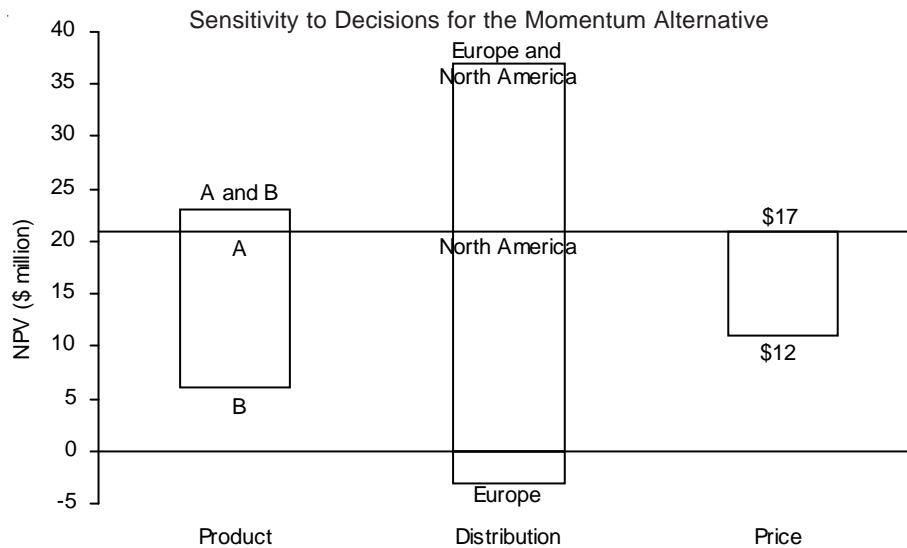


Figure 8–10



by the price reduction; once this was made clear, the decision team understood just how sensitive the venture was to price and the team became very interested in estimating the reaction of the market to pricing.

Probabilistic Analysis: The final step in the analysis phase is to bring everything together in a decision tree/influence diagram analysis. Examples of output from the probabilistic analysis shown in Chapter 6. The list of possible outputs includes: a simple tree with probabilities; probability distribution of the alternatives; sensitivity to probabilities; value of information; value of control; sensitivity to risk tolerance.

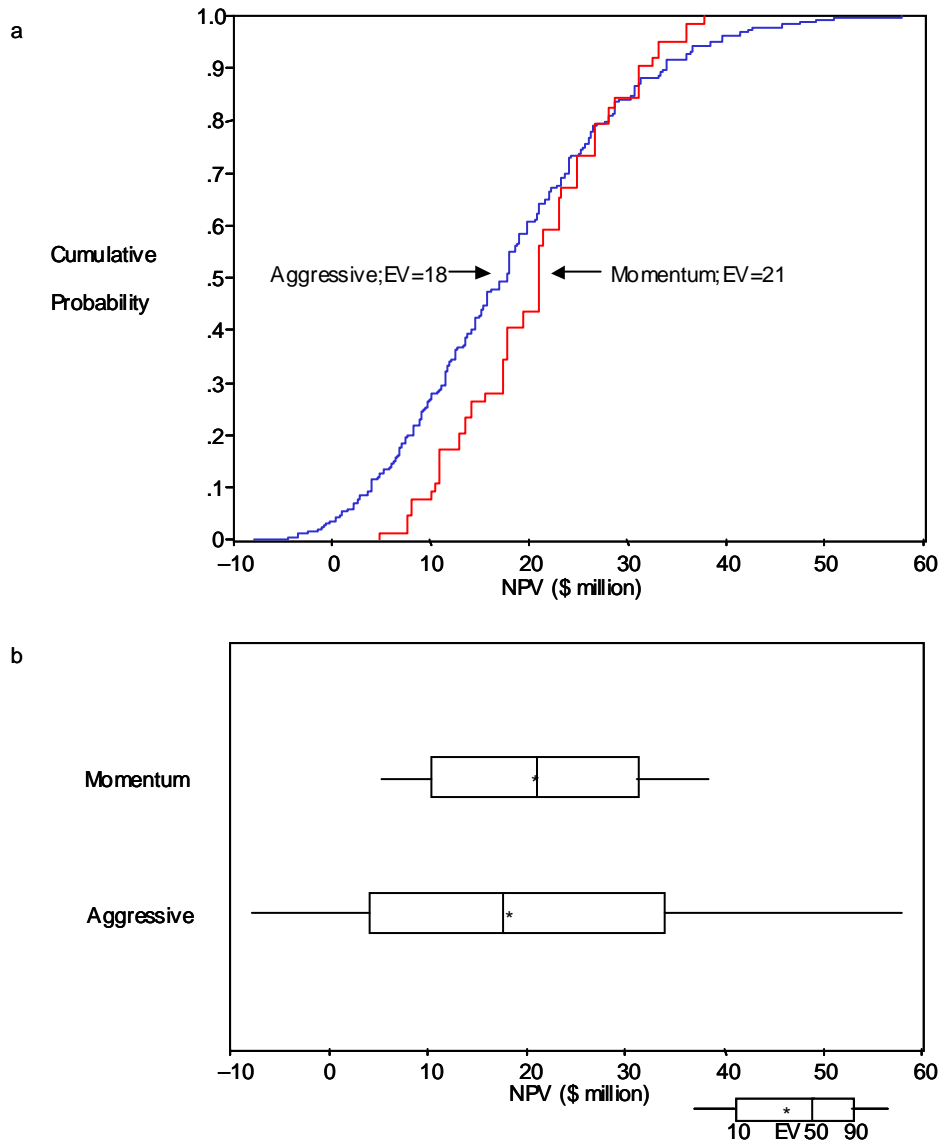
Figure 8–11 shows the probability distribution on NPV for FoodMachines' two alternatives in two forms: the cumulative probability distribution and the simpler form showing the 10 and 90 percentiles around the expected value. The latter was used in the presentation materials. It was clear that the aggressive alternative was not very desirable—it had lower expected value and more downside potential than the momentum alternative.

During the analysis phase dialog, insights often occur that lead to a better alternative. This alternative is created out of the alternatives that have been analyzed, using information from the various sensitivity analyses and from the combined wisdom of the teams. This new alternative is often called a “hybrid” alternative, and it is not unusual for this alternative to be much more valuable than any of the alternatives that were originally considered. For instance, the results shown in Figure 8–10 suggest a hybrid of the momentum alternative in which Product A alone was offered in both North America and Europe.

The decision team tried to absorb this avalanche of information. It was clear that there was considerable uncertainty, but that both alternatives were

Figure 8-11

Probability Distribution for FoodMachines' Alternatives



attractive—neither of the alternatives had a substantial chance of losing money.

The project team indicated that some preliminary analysis showed that the problem with the aggressive alternative was in the low pricing. With some reasonable assumptions on market behavior, higher pricing made the aggressive strategy a clear winner. The decision team requested the project team to analyze a hybrid alternative: the aggressive strategy with somewhat higher prices. In addition, the project team was requested to use a less investment-intensive form of B in the alternative. Finally, several additional sources of information were identified to help clarify some of the market issues.

Finally, the decision team indicated that, if the analysis of the hybrid alternative worked out well, it would be in a position to make a decision at the next meeting.

It is not uncommon for a second analysis dialog to be needed. The first set of alternatives may not be optimal (it is necessary to hold them constant during the analysis phase in order to gather a consistent set of data), and a really complex problem may require the decision-makers to review the analysis more than once.

Decision Dialog

Corporate decisions are not made in a moment, nor are they made once the analysis has been presented. The results need to be discussed, criticized, tested, and assimilated. Often, private discussions need to happen.

The decision dialog has still another purpose. During the project, the team has worked to share a frame, understand a problem, and develop the insights that result in a choice. But the choice will never produce results unless it is related to the larger world of the corporate organization. Both the decision team and the project team must begin to communicate their frame, their insights, and their enthusiasm to the people who will initiate, implement, and live with the chosen alternative. If this does not happen, the chosen alternative may never be implemented or, even worse, may be badly implemented by a recalcitrant, unenthusiastic, or misunderstanding organization.

The decision phase poses the following challenges:

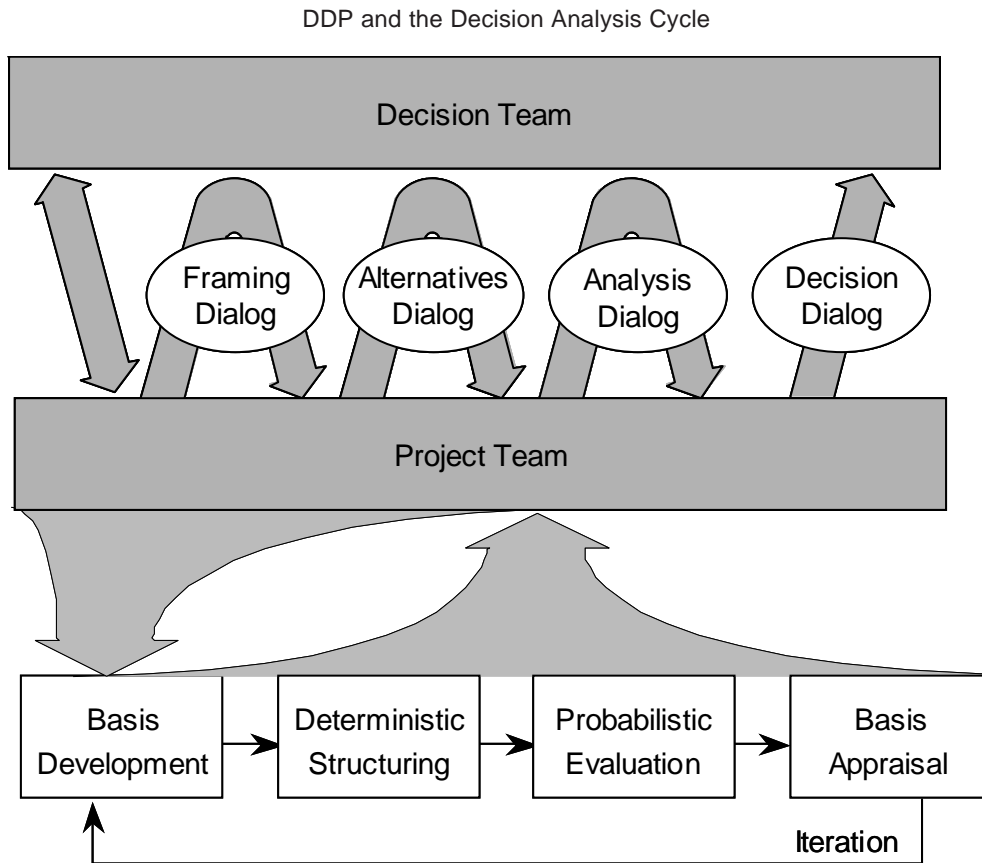
- Choose an alternative.
- Communicate the frame so that all involved see the full picture.
- Communicate the insights so that all involved see the reason for the choice.
- Communicate the robustness of the choice so that participants can see why their favored alternative may not have been chosen.
- Communicate enthusiasm for and commitment to the chosen alternative.
- Obtain organizational ownership of and commitment to the decision.
- Get the implementation started.

Warning: Some team members may feel depressed that what turned out to be a “fun” and exhilarating project is drawing to a conclusion. Rather than let the project end on a flat note, encourage the participants to be disciples of the process and to encourage the creation of more projects.

The deliverable for the decision dialog is simply a well thought out presentation of the chosen “hybrid” alternative—or of several “hybrid” alternatives, if there is still a significant choice to be made. This presentation may contain the elements of an implementation plan and of a plan to communicate the decision to the rest of the company.

FoodMachines’ decision dialog was almost anticlimactic. The hybrid alternative turned out pretty much as had been expected after the analysis dialog. There was some discussion whether the product was worth the marketing cost for European distribution, given the relatively low volume expected in Europe, but the head of European distribution made an impassioned plea for this product. In the end, the hybrid alternative was chosen and a product leader was appointed to begin implementation.

Figure 8–12



DDP and the Decision Analysis Cycle

The tools and deliverables of decision analysis appear throughout the DDP. But how does the DDP relate to the decision analysis cycle introduced in Figure 6–1 of Chapter 6?

The apparent mismatch arises because the DDP involves teams, and these teams are ordinarily composed of people pulled aside from their regular duties for the purposes of this special project. Organizationally, this requires a fixed beginning and end to the effort, with milestones at fixed points during the project. The decision analysis cycle, on the other hand, is structured to “continue the work until the decision-maker is ready to act.”

In practice, there are two correlated projects: the group process with its fixed schedule, and the analytical cyclical process going on in the “back room” (Figure 8–12). The analyst starts early, gets input from the project team and other sources, and makes sure that there are results available for the analysis dialog. Further iterations can (and will) be made for the decision dialog.

Project Staffing and Timing

Who is on the decision team? Most important corporate decisions cross one or more corporate boundaries. For instance, a new product decision will involve resources from marketing, design, engineering, manufacturing, sales, and finance. It would make sense to have representatives from all involved organizations on the decision team.

Who should be on the project team? Normally one would expect roughly the same composition as on the decision team. In addition, the decision facilitator(s) should be on the project team. Often, the decision facilitator serves as the neutral moderator of project team meetings.

How long should the process take? If the full process with four dialog meetings is followed, four calendar months is typical, although the duration varies a great deal in practice.* Among other things, decision-makers’ calendars are typically too crowded to arrange meetings more than once a month. The schedule should allow time for reflection and off-line discussion on the part of the decision team and for analysis and preparation of presentation materials by the project team. Four months may appear an inordinately long time for decision-making, but usually the decision is made at the end of the process and not revisited later. Corporate decisions that are made more quickly have often been reopened and remade, leading to a net expenditure of more time and effort and in the loss of opportunities.

*Major strategic decisions appear to take two to four months in corporations, regardless of the process used. See Kathleen M. Eisenhardt, “Strategy as Strategic Decision Making,” *Sloan Management Review*, Spring 1999, pp. 65-72 and C. J. G. Gersick, “Pacing Strategic Change: The Case of a New Venture,” *Academy of Management Journal*, volume 37, February 1995, pp. 9-45.

Presenting Decision Analysis Results

One of the most important tasks in a decision analysis or Dialog Decision Process is presenting the analysis results simply and convincingly. Decision analysis is supposed to help the decision-making process, and the findings are of little use unless they are well communicated. Unfortunately, in some cases, not enough effort is spent preparing the presentation insights and materials.

The first consideration in preparing a presentation is its purpose. Is it solely intended for the decision-maker(s)? Or is it intended as a communication tool and consensus builder—perhaps serving to make others appreciate the rationale behind the decision and build their enthusiasm for a course of action that they might not originally have championed? In the first case, the presentation should emphasize the conclusions and qualitative insights that came out of the analysis and discuss the next steps. In the second case, emphasis should be placed on how information derived from many different sources within the company comes together in the conclusions.

A second consideration in preparing the presentation is the appropriate level of detail. A good rule of thumb is that the higher in the organization the audience, the less the audience's interest in *how* the analysis was done. Rather, senior managers need to be convinced the work was well done, dealt with their major concerns, and has intuitively reasonable conclusions or recommendations.

In preparing the actual report or presentation, avoid concentrating on the methodology of the probabilistic or deterministic analysis. Generally, the decision-maker is interested only in the overall flow from believable input to reasonable conclusions. Presentations should not contain (unless particularly relevant or necessary) explanations of risk attitude, value of information and control, the techniques of tree evaluation, and the like.

Nor should the facilitator talk too much about the model. In many companies, those attending the presentation will tend to concentrate on what is familiar from most other presentations within the company: deterministic detail. The presenter must be skillful in leading the discussion smoothly but reasonably quickly through the deterministic phase and into the probabilistic phase. Again, the speaker should avoid (unless relevant or necessary) discussions of the discount rate, details of depreciation and tax treatment, undue concentration on the early years of the cash flow, and similar details.

A short presentation might have the following set of slides as a backbone for the presentation for the analysis dialog for a Dialog Decision Process:

- Introduction
- Principal alternatives
- Graphical description of deterministic model (if relevant)
- Base-case input and results for the Momentum alternative—a graph or table of financial and performance results for the first five or ten years is often an efficient way to present model logic and input data
- NPV for the base case for the principal alternatives

- Sources of value—by choice in the strategy table, by product, by region, by whatever yields insight
- Deterministic sensitivity analysis results
- Key probability assessments and information sources, perhaps shown in an abbreviated tree with probabilities
- Probability distributions for the alternatives
- Conclusions
- Value of information or control, probabilistic or risk sensitivity analysis, and the like, only if relevant.

Decision analysis results lend themselves well to graphic rather than tabular presentation—e.g., sensitivity analysis plots, trees, and probability distribution plots. Graphics have been shown to be the most effective way of ensuring both immediate comprehension and subsequent retention of the material presented.

Graphs should usually be presented as smooth, continuous curves. For instance, the staircase cumulative probability curves should be smoothed out, because the variable plotted (such as net present value) is usually a continuous variable. The process to smooth these graphs is the reverse of the discretization process discussed in Chapter 2.

Decision Analysis Capability Building ---

Although decision analysis techniques have been used for several decades to treat difficult corporate problems, the analysis itself has often been done on an *ad hoc* basis, often by outside consultants. To an increasing extent, however, the use of these techniques has become part of the problem-solving apparatus and decision-making process in many companies. While it is difficult to generalize, several traits are characteristic of successful decision analysis implementations or decision-making processes within these companies.

- Decision analysis and associated decision-making processes must be accepted, understood, and required by upper level management within the company. More important, decision analysis must have a strong sponsor at this level to thrive. This does not necessarily mean that upper management must understand the techniques; rather, it means the managers accept as a fact of life that uncertainty can and must be addressed in important decisions.
- Middle management also must be aware of, and sympathetic to, the decision analysis process. After all, these are the managers who commission the analyses and support their execution.
- There must be a talented, experienced technical champion of decision analysis. Since he or she will have to deal with often reluctant “clients” within the company, the individual needs skills in managing people, time, and budgets. A viable and attractive career path for this type of person must be created. More than once, a budding decision analysis effort has failed

when the key technical person was promoted or left the company.

- A fairly large number of people must have the technical capability to perform the decision analyses and lead the processes. This provides not only stability and continuity in the decision analysis effort, but also contributes to corporate commitment and enthusiasm.
- There must be a mechanism for training new facilitators within the company or for obtaining them from outside the company. A good facilitator requires some form of internship or apprenticeship—not just an academic background in the subject.
- The decision facilitator should be positioned in the company so he or she has access to the decision-maker, is authorized to obtain whatever information is necessary, and is not identified with any particular party or faction within the company.
- Many decision facilitators (and decision analysis groups) have a strong engineering or physical science background or philosophy. Decision analysis uses mathematical tools only insofar as they contribute toward an end. Decision-making processes must also deal with the organizational/people side of the problem. The personality of the facilitator (and group) must be comfortable with this; otherwise, analyses tend to become overly complicated and technical and miss the decision-maker's real needs.
- The first decision analyses within the company should be chosen with care, being neither too simple (“Why spend all this effort on the obvious?”) nor too complex. When problems are too complex, there is the danger of spending an inordinate amount of effort on the analysis and frustrating everyone involved.

Decision analysis may seem to be an expensive and time-consuming process. However, time and experience will show this is not so. When the philosophy, framework, process, and methodology have become established, decisions will be made efficiently and economically.

Once the initial effort has been made, most companies have found the investment in decision analysis capability justified. The variety of options considered, the quality of the knowledge employed, and the clear logic used lead to a decision process of high quality.

Summary

Decision-making in the modern corporation almost always involves cross-organizational teams both to analyze the problem and to make the decision. This adds a dimension of organization/people concerns to effective decision-making.

The Dialog Decision Process meets the challenges of implementing the logic of decision analysis in the team environment. A structured series of dialogs between the project team and the decision team provides the direction, reflection, insight, and communication needed to arrive at a decision.

For long-term, effective utilization of decision analysis in corporate environments, the insights of the analysis must be well-communicated and the facilitators themselves must be appropriately positioned and supported. Presentations are the most important in-house means of communicating the results of a decision analysis. Properly positioning and supporting the facilitators includes personnel selection and training, management understanding and support, and appropriate analysis project selection.

Problems and Discussion Topics

- 8.1 Although the Dialog Decision Process has four distinct meetings, the first two meetings (Framing and Alternatives) are sometimes combined into one short meeting. This often occurs for the evaluation of Research and Development (R&D) projects. Why might this be true? What problems might arise in moving too quickly through these two meetings?
- 8.2 In high-level corporate strategy decisions, the Framing dialog can be the longest and most complex of the dialogs, and sometimes is broken into two meetings. Why might this be true? Why might the Alternatives dialog be especially important in this situation?
- 8.3 Consider some personal decision situation, either past or future, which involves several people. Examples of such decisions are a group choosing a restaurant for dinner, a student choosing which college to attend, a couple deciding whether and when to get married, a class deciding on a class outing.
 - a. Who should be on the project team?
 - b. Who should be on the decision board?
 - c. If there are people who will be affected by the decision, but who are not on either team, what type of communication should be set up?
 - d. How do the answers to a, b, and c lead to decision quality? What potential decision quality problems could occur?
- 8.4 How much effort that should be devoted to the four parts of the decision process: framing, alternatives, analysis, decision? Express your answer in percentage of time/work in each part (four numbers adding to 1) for the following decision situations:
 - a. Choice of a college or graduate school to attend.
 - b. Choice of a restaurant for a special (e.g., birthday or anniversary) dinner.
 - c. A life-changing decision situation such as marriage, decision to have children, choice of career, etc.
 - d. Decision whether to buy a state lottery ticket when the prize

has grown enormous because no body has won for several weeks.

- e. Decision whether to attend a party on Monday night (e.g., Monday night football) or study for an exam Tuesday afternoon.
 - f. Choose between a fixed and variable-rate mortgage.
 - g. Choice between treatment alternatives for some life-threatening medical situation.
- 8.5 In presenting a decision analysis, you often need to clearly and credibly present results to people who may not be familiar with or understand the methodology used to arrive at the results. In what other kinds of business situations is this also the case?
- 8.6 List some of the considerations in deciding what level of detail to include in a decision analysis presentation.
- 8.7 How might you prepare a presentation differently if you were presenting to the operations research staff group as opposed to the vice president of marketing?
- 8.8 What kinds of changes in procedures for making decisions might occur as a company adopts decision analysis? How would the number and function of people involved in decision-making change?
- 8.9 Why might decision analysis have been adopted more rapidly in some industries than in others? Can all industries benefit from decision analysis?
- 8.10 The Lone Star Drilling Company has several prospects in Oklahoma, Texas, and Louisiana. One of these prospects, in the state of Oklahoma, is called Moose Hill. A promising region for natural gas underlies the Moose Hill area at 20,000 feet. Gas discovered at this depth qualifies as “deep gas” and is allowed to sell at a free market price under current regulations. (This is a disguised version of an analysis performed in the late 1970s.) There is little chance of finding oil under Moose Hill.

For gas to be found, there must be a structural trap. Currently available seismic studies indicate 7 chances in 10 there will be a structural trap. Even with a trap, there is a good chance that the water saturation will be too high for a producing gas well. A producing well could yield between 2 and 25 MCF/day the first year (MCF = million cubic feet); yields over 30 MCF/day are unlikely. Over the 10-year life of the well, annual production is expected to decline by 20 to 25 percent per year.

Lone Star is currently drilling a well on the Moosejaw 1 section at Moose Hill. While that well is not expected to reach 20,000 feet for another year, it appears that the cost of drilling a 20,000-foot well will be \$6 million to \$10 million, plus about \$2 million for completing the well if sufficient

gas potential is found. Annual operating costs for similar wells run between \$15,000 and \$25,000.

Property in the area is divided into sections of one square mile. Operators with mineral leases within a given section usually pool together and drill one well per section. However, under Oklahoma's forced-pooling statutes, any mineral leaseholder in a given section can decide to drill and invoke "forced pooling." Holders of the remaining leases in the section must then either join in a drilling operation within 90 days, sharing proportionally in drilling costs and potential gas yield, or offer to sell their rights to the first leaseholder at a price set by the state. The purpose of the statute is to encourage drilling in Oklahoma.

Lone Star holds 99 percent of the mineral rights to Moosejaw 2, a section adjacent to Moosejaw 1. The holder of the other 1 percent of the mineral rights has invoked forced pooling. The state is in the process of setting a "fair" price.

If Lone Star decides to pool and drill on Moosejaw 2, it has the option of negotiating with another exploration company, Delta Resources, for a joint venture—proportional sharing of all future costs and revenues from the property. Delta Resources has expressed an interest in this joint venture opportunity.

Your group is to recommend the best courses of action for the Lone Star Drilling Company. As part of your presentation, include the following.

- What is the minimum amount of compensation Lone Star should accept to sell its current 99 percent share to the owner of the remaining 1 percent, assuming Lone Star must otherwise bear 99 percent of the costs of drilling (no joint venture with Delta Resources)?
- Assuming Lone Star decides to go ahead and drill, what joint venture share should it offer to Delta Resources?
- Assume the state sets \$500,000 as the "fair" price for Lone Star's interests in the lease. Calculate the expected value of perfect information on a few crucial uncertainties.

Make sure the presentation will be acceptable to, and understood by, the president of Lone Star, an old-time driller who never graduated from high school, but who has acquired considerable wealth, experience, and expertise over the years.

- 8.11 Air Wars, Inc., a U.S. manufacturer of fighter planes, is aggressively marketing its popular Galaxy-MX and Scoop-UMi models to several emerging countries of the world. Sales discussions with two such countries, the Democratic Republic of Azultan (which has a reasonably stable government) and Byasfora's new government (which is an uneasy coalition between the Leninist-Marxist wing and the rightist Christian Democrats), are in the final stages in early 1985. Both of

these governments are also concurrently negotiating their air force armament needs with Le Mon Corporation, a European manufacturer. The discussions between the Le Mon Corporation and the governments of Azultan and Byasfora are of serious concern to the management of Air Wars, Inc.

Dr. Ian Winthrop, the CEO of Air Wars, Inc., has called an urgent meeting on the coming Saturday to assess Air Wars' position and to develop a clear strategy to make these sales. Dr. Winthrop, in his memo to senior management, reaffirmed the urgency of the situation and called for their input during the weekend meeting. Dr. Winthrop stressed Air Wars' commitment to growth during the coming years. He also brought senior management up to date on the key items in connection with the potential sale of the planes to Azultan and Byasfora.

Air Wars' Washington representative thought the U.S. government favorably regarded the plane sales to both the Azultan and Byasfora governments. However, the future stability of the new government in Byasfora was in question. A change in Byasfora's government was likely to result in a much more extreme left-wing government supported by neo-communists, creating concern about a reversal of the U.S. government's current support of the sale.

Probability of change of government in Byasfora by 1990: .45

The Washington representative also stressed the importance of the 1988 presidential election in relation to the sales to Byasfora. The most likely contender for president, if elected, is expected to consider the seriousness of the reported human rights violations in Byasfora and oppose the sale.

Probability of administration change in 1988: .80

Probability of opposition by new administration to Byasfora sale: .90

The negotiations with the U.S. government on Air Wars' cost structure for the sale of the Galaxy-MX and Scoop-UMi planes are nearing completion. Air Wars' Finance Department projects the following prices in 1985 dollars (contingent upon three possible U.S. government positions on cost structure).

Probability	Unit Price (\$ million)		
	.3	.6	.1
Galaxy-MX	\$3.9	\$4.6	\$5.2
Scoop-UMi	\$2.65	\$2.8	\$3.15

Last week, Dr. Winthrop met with the Secretary of State and the National Security Advisor. He was briefed on the current U.S. position with regard to the regional strategic balance of power in the Azultan and Byasfora

area. As a result, Dr. Winthrop feels that the current administration is unlikely to approve the sales to both the Azultan and Byasfora governments.

Probability of approving sales to both Azultan and Byasfora during 1987–1988: .01

Recent discussions between Dr. Winthrop and both the Minister of Air Defense of Azultan and the General of Strategic Air Forces of Byasfora resulted in satisfactory agreement on the numbers of planes needed and a shipment schedule for each country:

	Shipment Schedule							
	1989		1990		1991		1992	
	Galaxy	Scoop	Galaxy	Scoop	Galaxy	Scoop	Galaxy	Scoop
Azultan	22	9	21	12	8	4	—	—
Byasfora	—	—	18	13	22	8	11	8

In addition, the Azultan agreement calls for a five-year technology assistance contract at the rate of \$185 million per year beginning in 1989; the Byasfora agreement calls for a six-year \$205 million per year technology assistance contract beginning in 1990. These technology assistance contracts would be terminated if the employees or assets of Air Wars were threatened by any future catastrophic sociopolitical change in these countries.

Probability of Catastrophic Sociopolitical Situation in early 1990s

Azultan	.10
Byasfora	.45

The Finance Department of Air Wars has completed reports on the credit worthiness of Azultan and Byasfora. The credit worthiness was found to be closely tied to the economic condition of these countries. These countries were also dependent on world economic conditions for a portion of their natural resource base revenues. Based on an analysis of world and domestic economic outlook, the probability of their being unable to finance the necessary portion of the sales amount and honoring the technical assistance contracts is as follows:

Probability of Defaulting on Payments After 1990

Azultan	.30
Byasfora	.25

The option of insuring the credit risk is being considered, and Air Wars is making confidential inquiries to determine the fees for such protection.

Dr. Winthrop wants senior management to come up with a clear strategy for Air Wars and explain why it is the best strategy. In addition, Dr. Winthrop is also interested in trade-offs between price and risks. In view of the Le Mon competition, Dr. Winthrop believes that price flexibility to achieve a competitive price is extremely important for closing the sale. Therefore, an analysis comparing Air Wars' risks in relation to potential sales to Azultan and Byasfora is of significant value to Dr. Winthrop.

Your group is to prepare a presentation for Dr. Winthrop that addresses these concerns and clearly lays out the risks and possibilities inherent in this situation. The information given to you above may be redundant, incomplete, inconsistent, or unbelievable. Your group has to make the best of the situation and present a report. Your report should include:

- a. A clear structuring of the uncertainties and their relation to one another and to Air Wars ultimate sales revenues
- b. An analysis of the overall risk and the relative risk imposed by the individual uncertainties
- c. Values of information for the most crucial uncertainties
- d. Recommendations for further study or for possible actions to manage the most serious risks.

9

Decision Quality

Quality in Decision Making

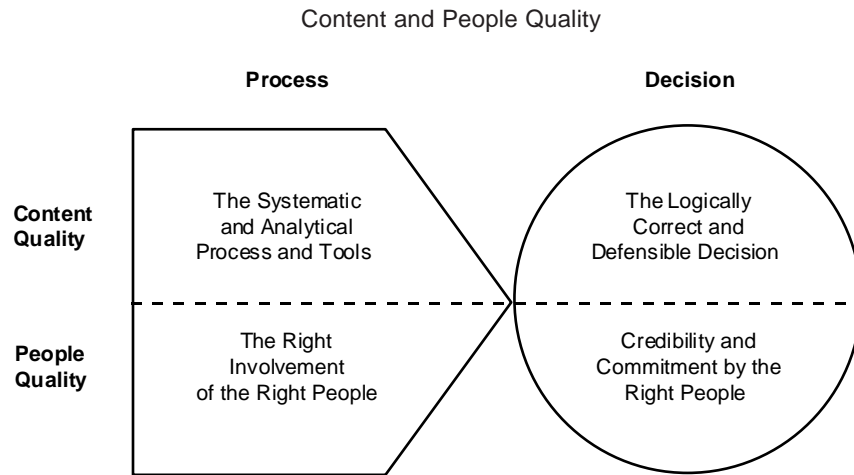
In Chapter 3, we made the fundamental distinction between a decision and an outcome—good outcomes are what we desire, whereas good decisions are what we can to maximize the likelihood of good outcomes. Whether we use the tools and processes outlined in this book or not—whatever the means we use to arrive at our decision—it all comes back to the same question: are we making a good decision?

For the individual decision-maker, it is usually clear when the time comes to make the decision whether the decision is “good,” whether he or she is decision-ready. Tools like the value of information and value of control help in making this judgment, but in the end, it is the personal conviction that further work or delay would not be justified.

On the other hand, the multiple decision-maker environment makes it difficult to determine when the organization is decision-ready. First, there must be agreement on the quality of the alternatives, information, and values, the three elements of the decision basis described in Chapter 6. But now there are additional questions that need to be addressed, such as: “Are we all addressing the same problem?” “Can we engage the organization to act?”

In addition, in the structured team environment of the Dialog Decision Process, there is no provision for extensive iteration. If we are to achieve quality at the end of the process, we need to be able to monitor quality during the course of the process in order to design corrective action when problems arise.

Figure 9-1



To be able to discuss these elements clearly and to judge when the teams are decision-ready, the language of decision quality has been developed.

People Quality and Content Quality

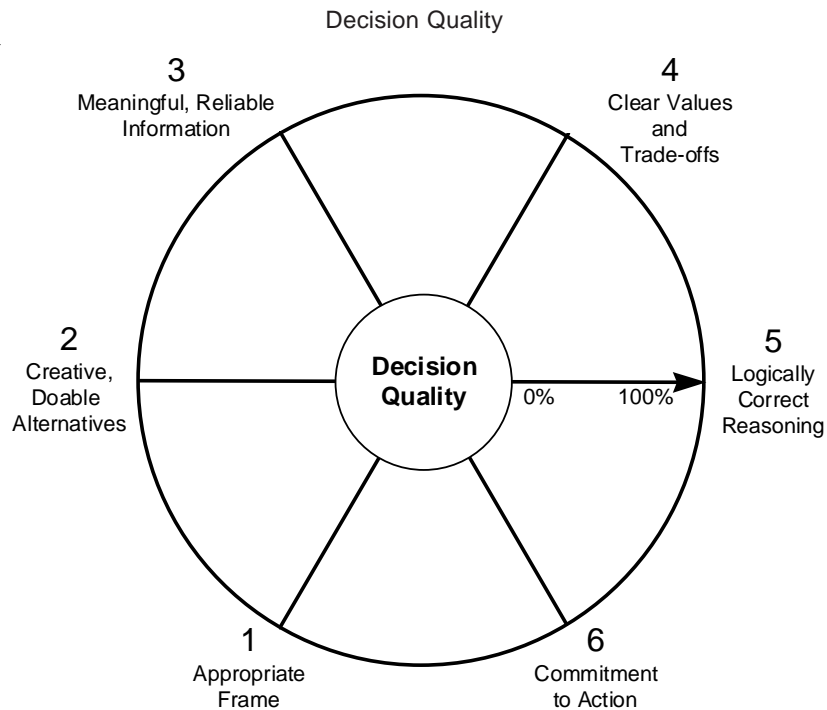
When we speak of quality in a product, we can refer to specific, quantifiable factors such as defect rate, variability, and customer approval. However, when we deal with complex, one-of-a-kind decisions whose outcomes may not be known for years, statistical verification is unrealistic. We can not obtain data regarding the outcomes of our decision until it is too late, and we will never know the outcome of the alternatives we did not choose. And, even if we were to discover the outcomes, we learned at the beginning of Chapter 3 that the quality of a decision involving uncertainty should not be judged by its outcome. How, then, can we assess the quality of such decisions?

In the corporate environment, content quality and people quality should certainly be part of the any definition of decision quality (Figure 9-1).

Content quality: The use of systematic processes and analytical tools leads to a logically correct and defensible decision. Without content quality:

- The choice may be “wrong”—not logically consistent with the company’s alternatives, information, and values.
- The choice can not be clearly communicated to the myriad of people responsible for its implementation.
- Later decisions may not be consistent with the original intent.
- Bad outcomes may be subject to unwarranted “Monday morning quarterbacking.”

Figure 9-2



People quality: The right involvement of the right people leads to credibility and commitment to the decision by the right people. Without the right people:

- The best alternatives and the best information may not be available to the decision-makers.
- The decision may not be implemented or may languish in the state of perpetual reconsideration.
- The decision may be reluctantly and badly implemented.

The Dialog Decision Process was designed to achieve quality in both of these dimensions. Similarly, the elements used in measuring decision quality have both of these dimensions at their heart.

Elements Used in Measuring Decision Quality

Six elements form the basis for measuring the quality of a decision, and each element has its content and people aspects. Decision quality is achieved by achieving quality in all six elements. The "spider diagram" in Figure 9-2 has six spokes radiating from the center, each representing one of the six elements. The outer rim represents 100 percent quality, the point where the cost of improvement exceeds the marginal benefit of the improvement, a concept familiar from the treatment of value of information and value of

control in Chapters 3, 4, and 6. The hub represents zero percent. The goal is to get as close to 100 percent in each of the requirements as possible.

What does the spider diagram provide? It provides a language by which the teams can come to agreement on where to focus efforts at successive points in the Decision Dialog Process. At the end of the process, if there is not consensus that adequate quality has been achieved in all six elements, the teams know that their job is not yet finished.

Appropriate Frame

The Framing Dialog typically does not yield 100% framing quality, but is an important step in getting there. Throughout the DDP, framing should be revisited; experience shows that frame shifts are not infrequent and can furnish much of the value of the DDP.

There are three important elements which contribute to framing quality: clear purpose, defined scope, and conscious perspective.

Clear purpose: When a team begins to work on a decision, its members rarely agree on what has to be done. A clear purpose must be established before the group begins to work. Frequently, this purpose is expressed in the form of a project vision or mission statement and a list of deliverables for each stage of the decision process.

Defined scope: What is inside the frame and what is outside the frame? What are the “givens”, the policy decisions? The scope of the decision process is the set of decisions that lie between the policies and the tactical decisions, as shown in the decision hierarchy.

Conscious perspective: We all use some (usually unconscious) framing rules to deal with the large amount of data we receive. Irrelevant data is filtered out and relevant data is dealt with. To achieve decision quality, team members must enlarge their perspectives to see the full set of relevant data and to appreciate the issues and concerns of everyone involved.

Failure in developing the appropriate frame can lead to:

- “Frame blindness” when team members do not perceive the limitations of their perspectives
- “Plunging in” without a clear vision of where the team needs to go, resulting in project confusion or in project failure
- Too narrow a scope, so that the true opportunity gets missed
- Too wide a scope, so that the project becomes overly complex and could fail because of project goals that can not be achieved or results that can not be communicated
- Unstated assumptions which, if untrue, could undermine the feasibility or desirability of the decision choice
- Solving the wrong problem.

Measuring decision quality is subjective. This measurement often occurs in the context of a team meeting. The following list provides some descriptions and criteria that can help in starting the conversation that leads to assessments of decision quality in the framing dimension.

- 0%—“Plunging in” or “Frame blindness”—no conscious perspective; scope and/or assumptions not stated; decision-makers not identified
- 50%—“Lists of issues, but not fully structured”—issues, perspectives, and concerns identified; decision-makers identified
- 100%—“Conscious, shared perspective”—clear statements of purpose, scope, perspective, and decisions to be addressed; agreement of decision-makers on frame

Experience has shown that, after initial hesitation, team members readily come up with estimates for the level of quality achieved for each spoke in the decision quality spider diagram. And the estimates are surprisingly consistent among team members.

Creative, Doable Alternatives

The Alternatives Dialog should achieve well over 50% quality in this element. It is important to incorporate high quality alternatives into the process from the beginning—the quality of the decision will be limited by the quality of the alternatives considered. After the first round of analysis, 100% quality can be achieved when the alternatives are refined and, perhaps, a “hybrid” alternative developed.

Six factors define high quality decision alternatives: they must be creative, achievable, significantly different, coherent, compelling, and complete.

Creative: Are the alternatives under consideration truly creative or are they minor variations on doing things as usual?

Achievable: Are time and resources being wasted on considering alternatives that are not achievable?

Significantly different: Do the alternatives span the range of possibilities? The interplay of several significantly different alternatives can lead to insight and the development of new, better alternatives.

Coherent: Do the elements of the alternatives make sense as a whole? A good alternative combines many different elements naturally. For instance, if the choice of new product features, market positioning, and pricing are related, a coherent alternative might be a combination of a high priced product with a full feature set positioned for the high-technology market.

Compelling: The theme and the choices made in each decision area come together in a way that can motivate people and generate a movement toward successful implementation.

Complete: The alternatives should include all the elements needed for beginning a successful implementation.

Failure in developing creative, doable alternatives can lead to:

- Considering only one alternative—a decision is a choice between alternatives, and if there are no alternatives, there is no decision and no point to the project
- Missing a great alternative, resulting in value foregone or even in strategic vulnerability if a competitor finds the great alternative

- Considering impracticable alternatives, resulting in wasted time, money and management attention.

The following list provides some descriptions and criteria that can help in starting the conversation that leads to assessments of decision quality in this element.

- 0%—"Business as usual"—only one alternative or several similar alternatives; alternatives that are not compelling or not achievable; good/excellent alternatives neglected
- 50%—"Creative, good alternatives"—alternatives that span the space; feasibility not completely verified; no clear winner, evaluation needed
- 100%—"Better alternatives created from insights developed"—significantly different and creative alternatives; new alternatives combining best features of original alternatives; implementation issues understood for each alternative

Meaningful, Reliable Information

Developing information quality begins during framing, particularly in problems in which a business assessment is required. It continues during alternative generation, and becomes a major concern during analysis.

Four elements define quality in information: knowing what is important, making sure the information is correct and explicit, using appropriate facts, and including the effect of uncertainty in the analysis.

Know what is important: What do you really have to know to make the decision? Frequently, you will find you need to use available data and to gather new data. If you ascertain your information needs early, you can avoid the pitfalls of needlessly analyzing inappropriate data you already have and of neglecting to gather data you need but do not have.

Make sure the information is correct and explicit: What does the information really say? Do not rely on shortcuts or sloppy conceptualization. Instead of using adjectives (e.g., the effect is "small"), state explicit values (e.g., the value is "approximately 2").

Use appropriate facts: What is the basis for the information? Be sure the underlying data support the conclusion. Trace the ancestry of information—a widely accepted and crucial bit of information may be based on little more than a time-hallowed educated guess.

Include uncertainty in the analysis: One of the important aspects of knowing something is knowing how well you know it. Most important and difficult decisions involve future events which are inherently uncertain. As a result, it is important to identify and deal explicitly with the effects of uncertainty.

Failure in developing meaningful, reliable information can lead to:

- Ignoring uncertainty, resulting in choices that ignore risk and that can be derailed by challenges of "What if...?" and "Have you considered...?"

- Missing interdependencies between factors, especially when uncertainty is involved, resulting in an incorrect statement of the effect of alternatives or in an under- or overstatement of uncertainty
- Focusing on what we know rather than on what is important, resulting in “near sighted” choices that ignore opportunities and threats, or in “safe” choices that are anything but safe
- Overlooking intangibles, resulting in choices that may ignore the critical element in making the decision.

The following list provides some descriptions and criteria that can help in starting the conversation that leads to assessments of decision quality. Calculations such as value of information and control may help in the assessment.

- 0%—“Blissful ignorance”—not knowing how much is known or what is important; ignoring uncertainty and/or “intangibles”
- 50%—“Informed about uncertainty”—knowing information gaps and what is important; uncertainty quantified; interdependencies not explored
- 100%—“Knowledgeable and ready”—information accurate, explicit, and based on appropriate facts; important knowledge gaps filled and limits of knowledge explored; interdependencies understood and taken into account; sources and rationales well documented

Clear Values and Trade-offs

Decisions are made to achieve something the organization places a value on. Frequently, however, several values can compete for attention in the decision process. Developing quality in this area begins in framing when the decision criterion is identified and is further developed through sensitivity studies and through insights developed during analysis.

To achieve quality in this area, clearly identify the company's decision criteria and establish the extent to which management is willing to trade off among these criteria. Two trade-offs arise in almost every decision problem:

Trade-off between long-term and short-term results: How do we value results today compared with results tomorrow? This trade-off is usually expressed in terms of a corporate discount rate such as the weighted average cost of capital (WACC). As discussed in Chapter 5, it is misleading to adjust this time preference to include risk aversion.

Trade-off between risk and return: How much expected return are we willing to give up to avoid risk? As discussed in Chapter 5, risk aversion can be expressed in terms of a risk attitude. However, it would be an unusual project in which it would be appropriate to introduce this technical topic into the dialog with the decision team. It is usually sufficient to display the probability distribution, thus showing the uncertainty (“risk”) in NPV associated with each alternative.

A clear distinction needs to be made between direct values and indirect values. A direct value is one that we seek to maximize in our choices—a decision criterion. An indirect value is one that is useful for management purposes, but is meaningful to the decision only insofar as it relates to the direct value. For instance, market share is often an indirect value, useful in managing a company, but disastrous if used as a decision criterion—one can increase market share by reducing price and profitability.

Failure in developing clear values and trade-offs can lead to:

- Neglecting a key constituency whose values conflict with the decision criterion used in making the decision, resulting in a contentious implementation
- Insufficient clarity in trade-offs, resulting in confusion and misunderstanding among key players
- Neglecting intangibles (e.g., good will on the part of suppliers or customers) which may be essential to successful implementation
- Double-counting the effects of risk, resulting in the choice of “safe” alternatives which may be strategically dangerous.

The following list provides some descriptions and criteria that can help in starting the conversation that leads to assessments of decision quality.

- 0%—“It’s not clear what we want”—preferences not explicit; stakeholders not identified; “intangibles” ignored
- 50%—“Clear value measures”—stakeholders and criteria identified; direct and indirect values distinguished; trade-offs need work
- 100%—“Clear trade-offs”—explicit statement of desired results in terms of decision criteria; explicit trade-offs made between criteria; double counting avoided

Logically Correct Reasoning

Clear logic is needed to convert the mass of alternatives, information, and values into a clear choice. Two aspects are important for quality in this area: reasoning clearly and developing the consequences of the alternatives in terms of the decision criterion.

Reasoning clearly: A good measure of quality in reasoning is whether the results can be explained to an intelligent outsider. Most problems do not require complex logic once the problem is well understood.

Developing the consequences of the alternatives in terms of the decision criterion: If the decision criterion is profit, what is the profit associated with each alternative? If it is ethics, what are the ethical ramifications of each alternative?

Failure in developing logically correct reasoning can lead to:

- Use of incorrect logic, resulting in incorrect results or loss of credibility;
- Models too cumbersome for sensitivity and probability analysis, resulting in numerical analysis without the insights required for

- making good decisions;
- Reliance on deterministic cases, resulting in inadequate estimates of risk and neglect of risk mitigation choices;
- Ignoring the effects of dependencies, resulting in choices that not optimal.

The following list provides some descriptions and criteria that can help in starting the conversation that leads to assessments of decision quality.

- 0%—“Extinction by instinct”—intuitive evaluation by each decision-maker
- 50%—“Visibility is still limited”—key sources of value identified; important uncertainties identified; dependencies not accounted for; logic incomplete and/or at insufficient level of detail
- 100%—“The best choice is clear”—reliable analysis of each alternative; uncertainties, dependencies, and complexities accounted for; analysis is “as simple as possible and no simpler”; clear choice based on frame, alternatives, information, and values

Commitment to Action

The single aspect of quality in this area is the motivation and commitment to action of necessary individuals. A productive way to achieve this is to involve the key implementers in the decision-making process from the beginning, preferably on the decision team or project team. They will then understand the frame, alternatives, information, values, and logic used to arrive at the decision. Their commitment to the process and understanding of the decision will almost always result in enthusiastic implementation.

Failure in developing commitment to action can lead to:

- Poor quality in the other decision quality factors
- Decisions which are never implemented, repeatedly reexamined, or poorly implemented.

The following list provides some descriptions and criteria that can help in starting the conversation that leads to assessments of decision quality.

- 0%—“Unmotivated”—lack of interest of key decision-makers; insurmountable organizational hurdles; insufficient support
- 50%—“Active decision board”—active participation by the right people; commitment to achieve decision quality; commitment to “do it right the first time”; buy-in not yet pervasive in the organization
- 100%—“Take action”—buy-in from project team, decision board, and those affected by the decision; sufficient resources allocated to implement and make the decision stick

Decision Quality and the Smart Organization _____

The goal of decision analysis is to arrive at a quality decision. There are a lot of elements that go into making a quality decision, and a natural question arises: which elements have successful organizations found most important

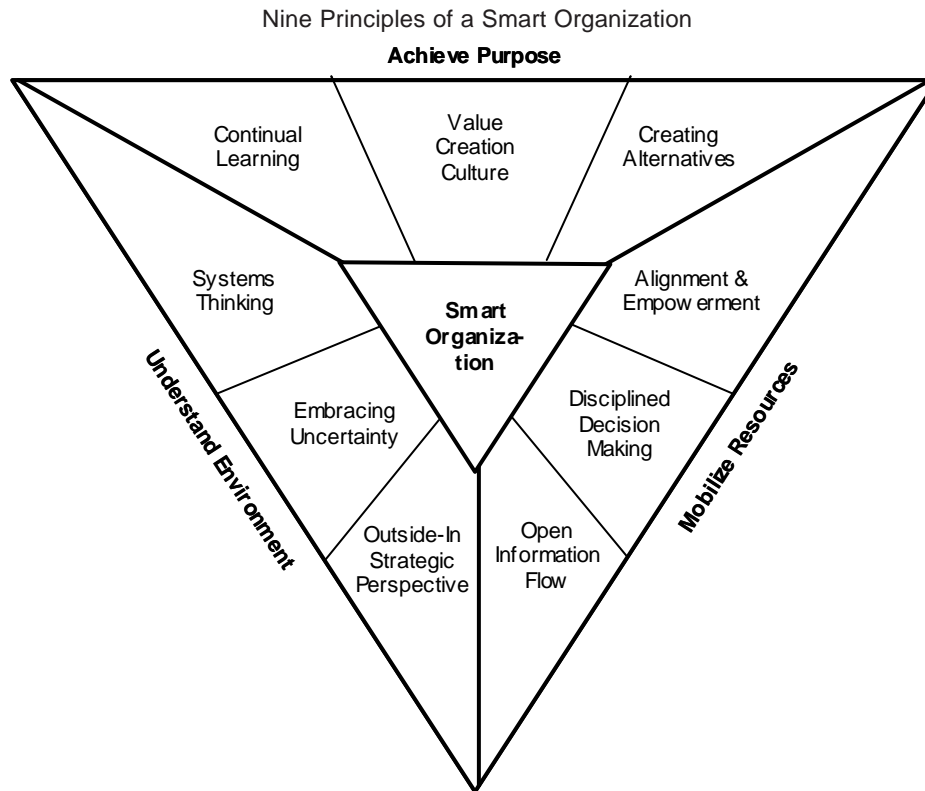
to identify and promote?

You can turn this question around and ask what an organization can do to most effectively change the culture of an organization? Vince Barabba has drawn on his experience at GM to state, “So if you want cultural change to come about—and to stick—stop fiddling with the organizational chart and start changing the decision-making process.”*

It is one thing to know how to use the processes described in this book to make one decision well; it is another thing to have an organization that routinely makes good decisions. There are a surprising number of things, some not obvious, that stand in the way of changing the decision-making process.

Several years ago, two of the authors’ colleagues became interested in what practices organizations used to achieve decision quality consistently over the course of many decisions. † R&D organizations appeared to be a good place to look. These organizations make many important decisions and have

Figure 9-3



*Vincent P. Barabba, *Meeting of the Minds: Creating the Market-Based Enterprise*, Boston: Harvard Business School Press, 1995, p. 219.

†See David Matheson and James Matheson, *The Smart Organization: Creating Value Through Strategic R&D*, Boston: Harvard Business School Press, 1998.

to live with the consequences. Our colleagues conducted a benchmarking study of practices of hundreds of R&D organizations, and correlated these practices with how “good” the organization was at making R&D decisions. This process identified a list of 45 best practices, which is not of direct interest to us here. What *is* of interest is the key to implementing these best practices.

Examination of the enablers (and barriers) for implementation of the best practices led to the identification of the nine principles of the smart organization (Figure 9–3.) These principles are subtle and work at many levels, influencing the way people think and act. The principles appear to be the world view of “smart” companies—companies that are agile, that are capable of delivering a stream of winning products and services, that make the right decisions at the right time.

These principles permeate an organization, defining elements of its culture and setting a context for decision making at many levels. When smart principles are in place, behaviors tend to support decision quality. Other principles are often in operation in an organization, and some of them create barriers to achieving decision quality.

All the nine principles directly support the elements of decision quality. In turn, espousing the goal of decision quality can help create the atmosphere in which the nine principles can guide the organization.

Achieve Purpose: Adhering to the principles of Continual Learning, Value Creation Culture, and Creating Alternatives naturally encourages achieving quality in creating an appropriate frame and creative, doable alternatives for all decisions.

Mobilize Resources: Disciplined Decision Making encourages quality in clarity of values and trade-offs and in logically correct reasoning. Alignment and Empowerment encourages quality in commitment to action. Open Information Flow is important for meaningful, reliable information.

Understand Environment: Outside-In Strategic Perspective and Embracing Uncertainty lead to meaningful, reliable information. Systems Thinking encourages logically correct reasoning.

Each of these principles has implications for the philosophy of the organization, the perspective of the people working within the organization, the culture of the organization, and the organization’s support systems.

Summary

Decision-making involving cross-organizational teams is difficult to manage. Team members bring different concerns, different perceptions, different values to the decision. How can the team manage this diversity and work toward results that are of high quality? How can the team judge when it is decision ready?

Decision Quality provides a practical framework and language by which teams can measure progress, direct future efforts, and judge when the organization is ready to commit to the choice of an alternative.

Decision quality can be achieved for individual decisions by using the processes and tools described in this book. To achieve decision quality in routine and ongoing decisions, the organization must espouse several important principles, which will have profound consequences within the organization.

Problems and Discussion Topics _____

- 9.1 Today's newspaper probably carries several stories relating to decisions announced by a public figure or organization. Choose one of these decisions and describe the tasks that would need to have been done to make this a quality decision. Using your judgment as a member of the public, how would you rate the actual quality of this decision? Note that you do not have to agree with the decision to judge its quality.
- 9.2 R&D decisions are frequently framed as a basic Go/No Go or Continue/Stop decision for a specific project. What do you think is especially important to achieving quality in an R&D decision?
- 9.3 Corporate strategy decisions set the direction of the company for the next few years. Acquisitions, divestitures, and shutting down facilities may be part of the strategy. What do you think is especially important to achieving quality in a corporate strategy decision?
- 9.4 Personal decisions such as choice of college, major, career, or marriage partner are decisions most of us face rarely, but which have great significance in our lives. Pick a personal decision that is (or will be) important to you. What do you think is especially important to achieving quality in this decision?

Part IV

Advanced Topics

10

Probability Theory

Theory Overview

In both business and personal life, we must confront the reality of uncertainty in the world and be able to describe it. As we have emphasized throughout this book, the natural language to describe uncertainty is the language of probability.

The words “probability theory” can induce feelings of apprehension in people who have little experience (or desire to acquire experience) in mathematics. Yet virtually everyone has some familiarity with probability. For example, almost everyone agrees that you win only half the time when calling the outcome of the flip of a fair coin, and most people are even comfortable with more sophisticated statements, such as “There is a 70 percent chance that it will rain today” or “There is only one chance in four that my alma mater will win the football game tomorrow.”

The probability theory used in this book does not involve any abstruse concepts or difficult mathematical formalism. The development is as intuitive and simple as possible. This chapter reviews those elements of probability theory that are important for decision analysis. The text of the book does not depend explicitly on the material in this chapter, and, thus, readers already familiar with probability theory can use this as a refresher.

Definition of Events

What is an event? An event is something about which you can say “It happened” or “It did not happen”—if you have sufficient information. This intuitive definition links the “real world” and the formulation of probability theory. The following examples may help to clarify the definition.

- “The spot price of oil was less than \$20/barrel at some time during 1986.” This is a statement about an event about which you can say “It happened.” Note that an event does not necessarily imply some dramatic change in conditions.

- “Our company’s net income in 2007 will be greater than \$2 million.” This statement describes an event which, in 2008, you will be able to characterize as having happened or not. Today, all you can say is that the event may or may not happen.
- “Company Q has developed a new product that will directly compete with our product.” You may have heard a rumor to this effect. People at Company Q know whether this event has happened or not. However, your information is insufficient to tell you whether the event has happened or not.

The Venn diagram furnishes a convenient way of representing events. In the Venn diagram, each point represents a different possible world. For instance, one point in the square might represent a world in which the price of oil in 2008 is greater than \$300 per barrel, in which the great-great-granddaughter of John F. Kennedy becomes president of the United States, in which Company Q does not come out with its rumored new product, and in which many other events, trivial or important, happen or do not happen. However, only one point in the square represents the world that will be realized as time unfolds.

This rather abstract construct is useful for representing a real problem when the points are arranged into areas in which specific events do or do not happen.

Throughout this chapter, we use the fictitious case of Medequip as a simple example. All Medequip discussions are set in italics.

Medequip manufactures a complete line of medical diagnostic equipment. Medequip’s planning department is studying one of the products in this line. The product in question has been a good performer, but competitive pressures have caused unit revenues to decline steadily over the past five years. There is some concern that, even given Medequip’s experience in the area, the unit costs will not decline enough over the coming years to keep the product profitable. The planning department has chosen 2007 as a good year to represent the end of the long-term trends. It has defined the following three events in terms of unit costs in that year:

- C_1 —Unit Cost less than \$1,000
- C_2 —Unit Cost between \$1,000 and \$1,500
- C_3 —Unit Cost more than \$1,500.


These three events are graphically represented in the Venn diagram (Figure 10–1). The area labeled C_1 , for instance, includes all possible worlds in which event C_1 happens.

Distinctions

Thoughts and conversations about *events* are made possible by our capability to make *distinctions*. A distinction is a thought that separates one large thing (or concept) into several smaller things (or concepts.) Distinctions define the lines in Venn diagrams like the diagram in Figure 10–1.

Figure 10-1

Venn Diagram Divided into Regions with Different Unit Costs



Some familiar distinctions are furnished by the words in italics below:

- Is the unit cost *less than \$1,000/between \$1,000 and \$1,500/greater than \$1,500?*
- Will this R&D project will be *technically successful/unsuccessful?*
- Will this cup of coffee *taste good/bad?*

A high quality distinction needs to be both clear and useful. Both clarity and usefulness are relative terms, and depend on the decision context.

A *clear* distinction is one that passes the clairvoyance test described in Chapter 2. The clairvoyant is a person who can accurately answer any question, even about the future, but who possesses no particular expertise or analytical capability. The clairvoyance test is a mental exercise to determine if the clairvoyant can immediately answer a question or if he needs to know other things first. “Unit cost” is fairly clear, but might need a specification of the product characteristics and the manufacturing conditions before the clairvoyant could answer the question above. “Technical success” would need to be defined before the clairvoyant could answer the question. “Tastes good” is a very personal judgment, and the question could probably never be answered by the clairvoyant.

A *useful* distinction is one that means what we want it to mean, and helps us achieve clarity of action. The mark of an excellent decision facilitator is his or her ability to create distinctions that elegantly cut to the heart of the matter. If we are trying to choose a restaurant for breakfast, the distinction concerning the taste of the coffee could be crucial in achieving clarity in the choice—if only we could create a clear definition! Technical success and unit costs are useful distinctions in arriving at clarity in decisions concerning new product introductions.

Figure 10-2

Representation of the Operations "And," "Or," and "Not"




Figure 10-2: Representation of the operation "And." The shaded region represents the event A and B.




Figure 10-2: Representation of the operation "Or." The shaded region represents the event A or B.

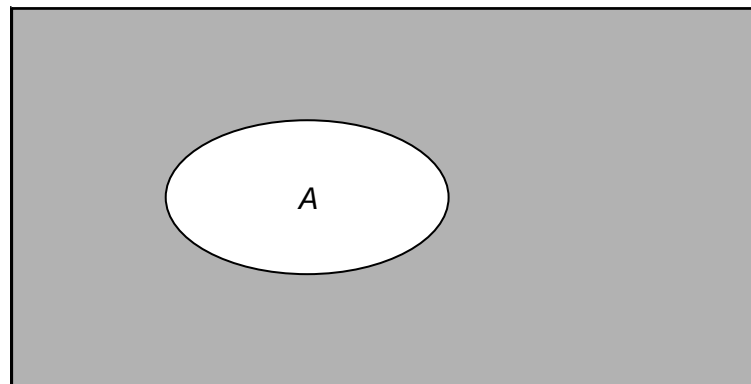


Figure 10-2: Representation of the operation "Not." The shaded region represents the event not A.

Algebra of Events

The Algebra of Events is a powerful formalism developed to deal with events and whether they happen or not. We present only a very small portion of this formalism.

Three important operations are used to combine or modify events. If A and B are any two events, then the effect of these operations can be seen graphically in Figure 10-2.

It is convenient to define two special subsets of the possible worlds that can occur. These are the universal set (I) and the null set (\emptyset), as defined below.

- I : All possible worlds. Graphically, this is the whole area of the diagram.
- \emptyset : No possible worlds. Graphically, this is none of the diagram.


The definition of events and the operations defined above appear in a number of quite different disciplines. Four of these are Algebra of Events, Formal Logic, Boolean Algebra, and Set Theory.

Mutually Exclusive and Collectively Exhaustive Events

Probability analysis (and decision analysis) often moves in the direction of dividing up all the possible worlds into finer and finer subsets. This process allows better discrimination among the quantities of interest and better use of the data available. The set of events used to characterize the decomposition or subdivision should be mutually exclusive and collectively exhaustive.

Figure 10-3

Set of Events (K_1 and K_2) That Is Neither Mutually Exclusive nor Collectively Exhaustive



Mutually Exclusive

Let us imagine that we have a set of m events, X_i , with $i = 1, 2, \dots, m$. This set of events is mutually exclusive if

$$X_i \text{ and } X_j = \emptyset \quad (10-1)$$

Collectively Exhaustive

The second property desired of a list of events is that it include all possibilities. A set of m events, X_i , with $i = 1, 2, \dots, m$ is collectively exhaustive if

$$X_1 \text{ or } X_2 \text{ or } \dots \text{ or } X_m = I \quad (10-2)$$

If the set of events is collectively exhaustive, then any point in the Venn diagram is in at least one of the regions X_i . If the set is also mutually exclusive, then any point will be in one (and only one) of the regions X_i .

An example of a set of events that is *neither* mutually exclusive *nor* collectively exhaustive is the set K_1 and K_2 defined below.

- K_1 : Company Q introduces an inexpensive new product (priced less than \$2,000)
- K_2 : Company Q introduces an expensive new product (priced greater than \$1,500)

The set in Figure 10-3 is not mutually exclusive because there is some overlap of the regions K_1 and K_2 ; a price of \$1,750 falls in both K_1 and K_2 . The set is not collectively exhaustive since there is an area outside the regions K_1 and K_2 ; the event that Company Q does not come out with its new product falls outside the regions K_1 and K_2 .

Medequip was satisfied that its set of events— C_1 , C_2 , and C_3 —was mutually exclusive. However, the planning department found that it is difficult to verify that a list of events is collectively exhaustive. Even in this simple case, there were some problems. For instance, where do the values \$1,000/unit and \$1,500/unit fall? In event C_2 ? What of the possibility that the product is not even produced in 2007? Is this represented in event C_3 ? The planning department refined its definitions as follows:

- C_1 —Unit Cost \leq \$1,000
- C_2 —\$1,000 < Unit Cost \leq \$1,500
- C_3 —\$1500 < Unit Cost

The possibility of not manufacturing the product was judged so remote that it could be effectively included in the high-cost scenario, C_3 .


In cases of real difficulty, it is possible to define an “all other” event.

$$X_{m+1} = \text{not}(X_1 \text{ or } X_2 \text{ or } \dots \text{ or } X_m) \quad (10-3)$$

This will make the set of events X_i ($i = 1, 2, \dots, m, m+1$) collectively exhaustive. However, this event is useful only as a reminder to keep looking for the (as yet) unknown events needed to completely describe all the possible worlds.

Figure 10-4

Venn Diagram Separated into Regions of Different Unit Revenue



Joint Events

As mentioned above, probability analysis (and decision analysis) often moves in the direction of dividing up all possible worlds into finer and finer subsets. This process is frequently accomplished by combining two different sets of events to give a joint set of events. If A and B are any two events, a joint event is defined as

$$A \text{ and } B \tag{10-4}$$

In addition to being concerned about unit costs, Medequip’s planning department was also concerned about unit revenues. For the initial phase of the analysis, they were content to define the following very simple set of mutually exclusive and collectively exhaustive events in terms of unit revenue in 2007 for the product:

- R_1 —Unit Revenue \leq \$1,750
- R_2 —\$1,750 < Unit Revenue

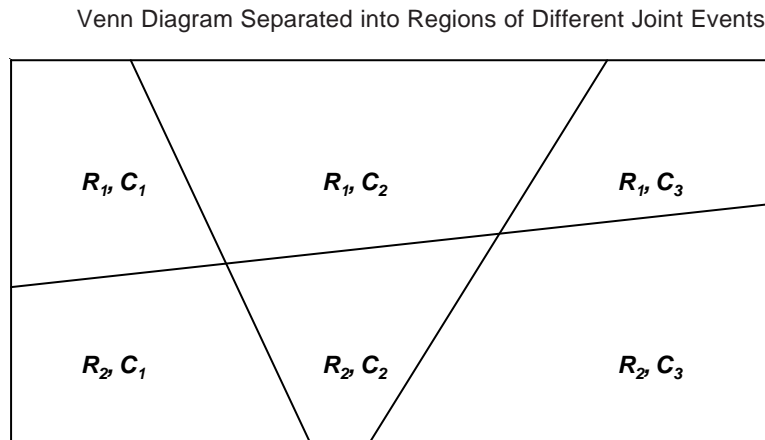
Graphically, this second set could be represented as in Figure 10-4.

These two sets of events combine to give six joint events, M_i , which describe the margin in 2007 (margin equals unit revenue minus unit cost). These six joint events can be represented in the Venn diagram (Figure 10-5).

- | | |
|---------------------------|---------------------------|
| • M_1 — R_1 and C_1 | • M_4 — R_2 and C_1 |
| • M_2 — R_1 and C_2 | • M_5 — R_2 and C_2 |
| • M_3 — R_1 and C_3 | • M_6 — R_2 and C_3 |

Note that the operation “and” is often denoted by a comma, as in Figure 10-5.

Figure 10-5

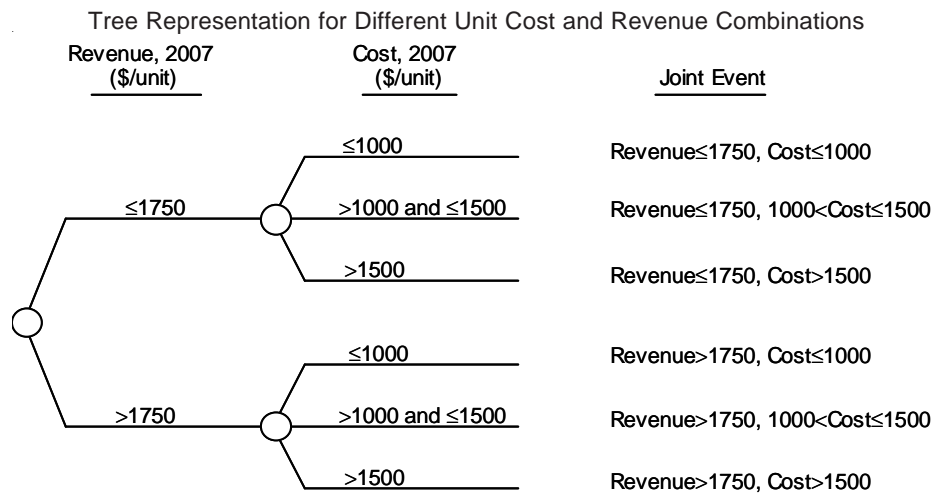


Tree Representation of Events

Because the tabular definitions and the Venn diagram representation used above become cumbersome for all but the simplest problems, we use tree forms as a simple way to represent even complicated problems. Each set of events is represented by a “node” or branching point in the tree. Each branch emanating from a node represents an event. Each path through the tree (moving from left to right) describes one joint event composed of all the events on the path. By convention, the sets of events used at a node are always mutually exclusive and collectively exhaustive.

Figure 10-6 is the tree for the margin for Medequip’s product.

Figure 10-6



Probability and States of Information

There are some events for which we have sufficient information to say “It happened” or “It did not happen.” For instance, most people would agree that the event “Thomas Jefferson became the first president of the United States” did not happen. However, there are other events about which we do not possess enough information to say whether they did or did not happen. We may simply not know what happened in the past. For instance, we may not know whether our competitor had a successful R&D outcome or not. More typically, the event in question is a possible future event. We do not know if our unit cost in 2007 will be less than \$1,000. The answer to that question cannot be given until 2008.

While we may not have sufficient information to give a definitive answer to “Did it or will it happen?,” we can assign some probability (or likelihood) to whether an event did or will happen. This is something we do informally in everyday business and personal situations.

Let us define

$$p(A|S) \quad (10-5)$$

to be the probability we assign that event A did or will happen, given our present state of information, S . While the S is commonly dropped from the notation, if it is dropped there will often be confusion about just what state of information underlies a probability. For instance, someone may judge the probability that event C_1 (unit costs in 2007 less than \$1,000) will occur is 1 in 10:

$$p(C_1|S) = .1 \quad (10-6)$$

However, some days later, after learning that a large deposit of a rare and critical raw material has recently been discovered, the person may revise his probability assessment, given the new state of information, S' , to 1 in 4:

$$p(C_1|S') = .25 \quad (10-7)$$

If we define the event

D : Large deposit of raw material discovered,

then

$$S' = D \text{ and } S \quad (10-8)$$

As mentioned above, the “and” operation is frequently denoted by a comma in probability notation. We can then write

$$p(C_1|S') = p(C_1|D, S) \quad (10-9)$$

In performing decision analyses, it is often necessary to combine or compare probabilities obtained from different people or from the same person at different times. Explicit reference to the underlying state of information is essential to keeping the calculations consistent and meaningful.

Probability Theory

The theoretical underpinnings of probability theory are really quite simple. There are only three axioms necessary, given our understanding of the events above. If A and B are any two events, then we can state the axioms as follows:

$$p(A|S) \geq 0 \quad (10-10)$$

$$p(I|S) = 1 \quad (10-11)$$

If A and $B = \emptyset$, then


$$p(A \text{ or } B|S) = p(A|S) + p(B|S) \quad (10-12)$$

If we take the Venn diagram and rearrange and stretch its surface so the areas of each region are proportional to the probability that the event that defines the region happens, then the axioms have the following graphical interpretation:

1. There are no regions of negative area.
2. The area of the total square is unity. This is the definition of the unit of area.
3. If two regions are nonoverlapping (A and $B = \emptyset$) as in Figure 10-7, then the combined region (A or B) as represented by the shaded region has area equal to the sum of the areas of the two component regions.

Figure 10-7

Venn Diagram with A and B Not Overlapping



Given these three axioms, we can easily deduce properties commonly associated with probabilities. The following four properties are frequently used. The proof of these properties is left to the problems at the end of this chapter.

1. Probabilities are numbers between zero and one.

$$0 \leq p(A|S) \leq 1 \quad (10-13)$$

2. Probabilities sum to one. More precisely, if X_i with $i = 1, 2, \dots, m$ is a set of mutually exclusive and collectively exhaustive events, then

$$\sum_{i=1}^m p(X_i | S) = 1 \quad (10-14)$$

3. For A (any event) and X_i with $i = 1, 2, \dots, m$ a set of mutually exclusive and collectively exhaustive events,

$$p(A | S) = \sum_{i=1}^m p(A, X_i | S) \quad (10-15)$$

This property is often called the Expansion Theorem.

4. For A and B (any two events), then

$$p(A \text{ or } B | S) = p(A | S) + p(B | S) - p(A, B | S) \quad (10-16)$$

The last term on the right in the above expression compensates for the double counting (if the events are not mutually exclusive) of the overlapping (double shaded) region in Figure 10-8.


Joint, Marginal, and Conditional Probabilities

Rarely is the uncertainty of a problem well described in terms of a single set of mutually exclusive and collectively exhaustive events. More frequent is a description in terms of a set of joint events—several sets of mutually exclusive and collectively exhaustive events are used to subdivide the possible worlds into fine enough detail for the work at hand. The probability of the joint events occurring is called the joint probability. If A and B are any two events, the joint probability for the two events occurring is

$$p(A, B | S) \quad (10-17)$$

Figure 10-8

Venn Diagram with A and B Overlapping



Medequip’s planning department used all the information it possessed to estimate the joint probabilities for unit revenue and unit cost. It gathered all the information it could on the principal competitor’s process and on his pricing policy. The department studied historical trends on raw material cost and used the judgment of the managers of the production process to estimate future cost trends. All this information was used to estimate the joint probabilities shown in the table below.

$p(R_i, C_j S)$			
	C_1	C_2	C_3
R_1	.10	.25	.03
R_2	.22	.26	.14

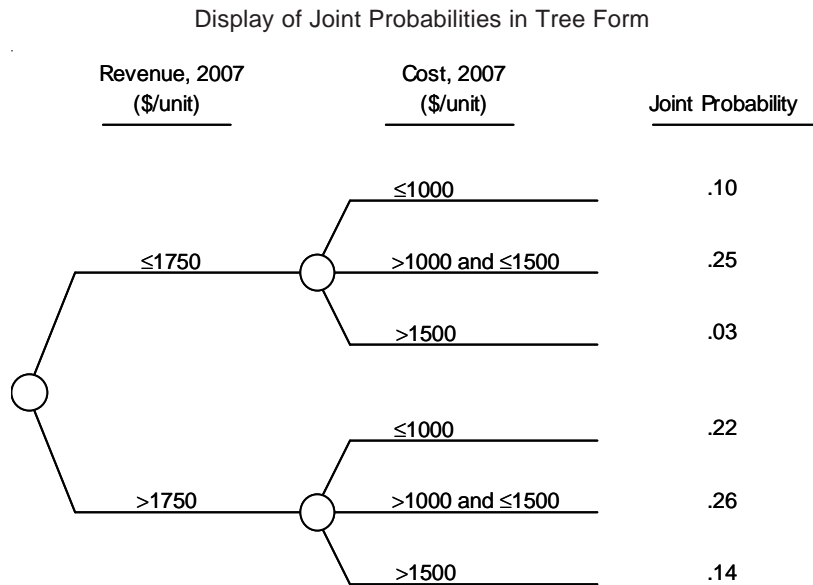
This set of joint probabilities is also displayed in the tree in Figure 10–9.

The marginal probabilities are those placed at the right or bottom edges (margins) of the table of joint probabilities and are obtained by summing probabilities across rows or down columns. This process uses the Expansion Theorem:

$$p(A | S) = \sum_{i=1}^m p(A, X_i | S) \tag{10-18}$$

In the above expression, A is any event and X_i ($i = 1, 2, \dots, m$) is a set of mutually exclusive and collectively exhaustive events.

Figure 10–9



In the Medequip case, the marginal probabilities are easily obtained from the table above. For instance,

$$p(C_1 | S) = p(R_1, C_1 | S) + p(R_2, C_1 | S)$$

$$p(C_1 | S) = .10 + .22 = .32 \tag{10-19}$$

The full set of marginal probabilities is presented in the following table.

	$p(R_i, C_j S)$			$p(R_i S)$
	C_1	C_2	C_3	
R_1	.10	.25	.03	.38
R_2	.22	.26	.14	.62
$p(C_j S)$.32	.51	.17	

Conditional probabilities are defined from the joint and marginal probabilities. If A and B are any two events, then

$$p(A | B, S) \tag{10-20}$$

is the conditional probability—the probability that A occurs, given that B occurs and given the state of information S . This conditional probability is obtained from the joint and marginal probabilities by the following definition:

$$p(A | B, S) = \frac{p(A, B | S)}{p(B | S)} \tag{10-21}$$

For the Medequip case, the conditional probabilities are easily obtained from the table above. For instance, the probability of C_1 occurring, given that R_1 occurs is

$$p(C_1 | R_1, S) = \frac{p(R_1, C_1 | S)}{p(R_1 | S)} = \frac{.10}{.38} \tag{10-22}$$

Tables of the values of the conditional probabilities $p(C_j | R_i, S)$ and $p(R_i | C_j, S)$ are given below.

	$p(C_j R_i, S)$		
	C_1	C_2	C_3
R_1	10/38	25/38	3/38
R_2	22/62	26/62	14/62

	$p(R_i C_j, S)$		
	C_1	C_2	C_3
R_1	10/32	25/51	3/17
R_2	22/32	26/51	14/17

These probabilities show, for instance, that high unit costs are much more likely when unit revenues are high than when unit revenues are low.

In tree form, the probabilities written at the nodes of the tree are, by definition, the probabilities conditional on all the nodes to the left of the node in question. The probabilities of the leftmost node are its marginal probabilities.

For the Medequip case, the tree can be written with the unit revenue node on the left. In the display in Figure 10–10, the values and symbols for the probabilities are both written. The symbols are usually omitted.

Bayes' Rule

Bayes' Rule is a simple rule that relates conditional and marginal probabilities. It is of central importance to decision analysis.

Bayes' Rule solves the following problem: Let X_i ($i = 1, 2, \dots, m$) and Y_j ($j = 1, 2, \dots, n$) be two sets of mutually exclusive and collectively exhaustive events. Given the following marginal and conditional probabilities,

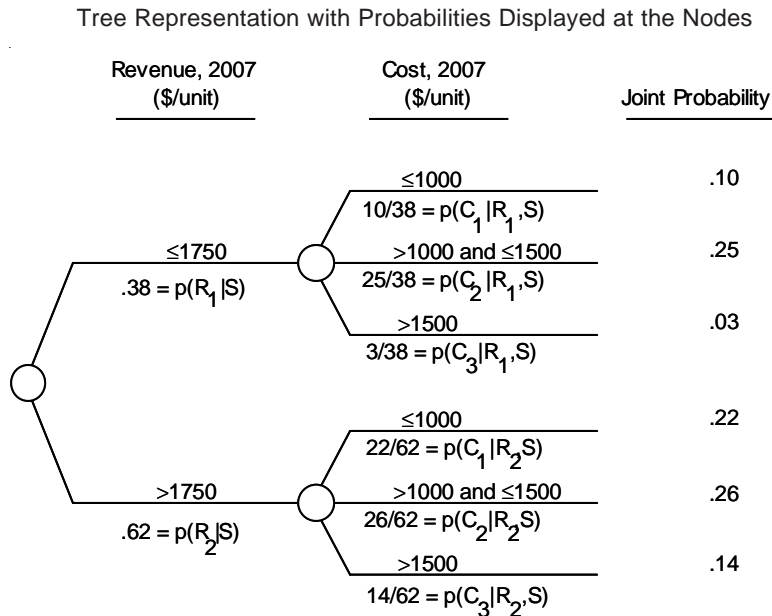
$$p(Y_j | S) \tag{10-23}$$

$$p(X_i | Y_j, S) \tag{10-24}$$

how do we calculate the other marginal and conditional probabilities? This operation is called for when we reverse the order of nodes in a tree. We can write the joint probability in terms of the given probabilities as follows:

$$p(X_i, Y_j | S) = p(X_i | Y_j, S)p(Y_j | S) \tag{10-25}$$

Figure 10–10



We can just as well write the joint probability in the reversed order, as follows:

$$p(X_i, Y_j | S) = p(Y_j | X_i, S)p(X_i | S) \quad (10-26)$$

Equating the right-hand sides of the two equations above enables us to solve for the conditional probability term, as shown below.

$$p(Y_j | X_i, S) = \frac{p(X_i | Y_j, S)p(Y_j | S)}{p(X_i | S)} \quad (10-27)$$

The only unknown in the right-hand side of this equation is $p(X_i | S)$. This can be obtained by using the Expansion Theorem (Equation 10-15) to write this probability in known terms, as illustrated below.

$$p(X_i | S) = \sum_{j=1}^n p(X_i, Y_j | S) \quad (10-28)$$

The joint probability on the right-hand side of the above equation can then be written in terms of the known conditional and marginal probabilities. The probabilities we are seeking are then written in terms of known quantities, as follows:

$$p(X_i | S) = \sum_{j=1}^n p(X_i | Y_j, S)p(Y_j | S) \quad (10-29)$$

$$p(Y_j | X_i, S) = \frac{p(X_i | Y_j, S)p(Y_j | S)}{\sum_{k=1}^n p(X_i | Y_k, S)p(Y_k | S)} \quad (10-30)$$

In the Medequip example, assume we know the marginals $p(R_i | S)$ and the conditionals, $p(C_j | R_i, S)$. This is the information given on the tree at the end of the previous section (Figure 10-10). Then, for example, we could calculate the following probabilities from those given.

$$\begin{aligned} p(C_2 | S) &= [p(C_2 | R_1, S)p(R_1 | S)] + [p(C_2 | R_2, S)p(R_2 | S)] \\ p(C_2 | S) &= [(25/38) \times .38] + [(26/62) \times .62] = .51 \end{aligned} \quad (10-31)$$

$$p(R_2 | C_1, S) = \frac{p(C_1 | R_2, S)p(R_2 | S)}{p(C_1 | R_1, S)p(R_1 | S) + p(C_1 | R_2, S)p(R_2 | S)}$$

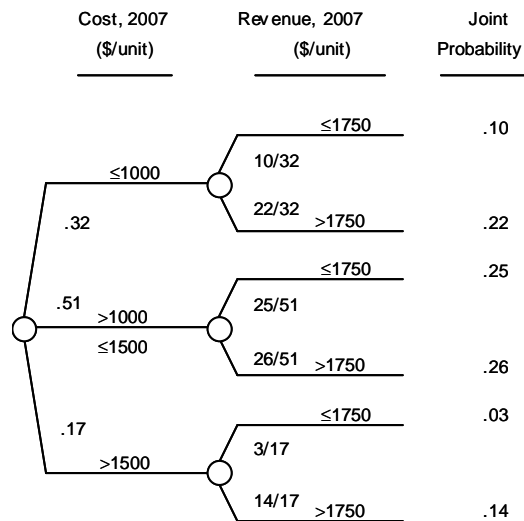
$$p(R_2 | C_1, S) = \frac{(22/62) \times .62}{(10/38) \times .38 + (22/62) \times .62} = 22/32 \quad (10-32)$$

For Medequip, the tree when reversed is as shown in Figure 10-11.

In this tree, we can see an easy, graphical way of performing the calculations to reverse the order of nodes. First, the joint probabilities are taken from the tree in Figure 10-10 and put at the end of the appropriate branches in the tree in Figure 10-11. The probability for each Cost branch is then obtained by summing the joint probabilities for the Revenue branches

Figure 10–11

Reversing the Tree in Figure 10–10



following that Cost branch. The individual Revenue probabilities, then, are simply the joint probability divided by the Cost probability (so that the Cost and Revenue probabilities give the right joint probability when multiplied together).

Probabilistic Independence

Sometimes conditional probabilities turn out to be the same regardless of which conditioning event occurs. For instance, prices could be set by a competitor using a process quite different from our own. In this case, the probabilities for prices and for costs might be independent of each other. If there is no correlation among the probabilities, they are said to be probabilistically independent. If A and B are two events, they are probabilistically independent if

$$p(A|B,S) = p(A|S) \tag{10-33}$$

Multiply or Add Probabilities?

Inevitably, the question arises “When should I multiply probabilities and when should I add them?” For joint (“and”) events, conditional probabilities for the component events are multiplied together. For combined (“or”) events, probabilities are added (if the events are mutually exclusive).

This is illustrated in terms of the Medequip example. A joint (“and”) probability is

$$p(R_2, C_3 | S) = p(R_2 | S)p(C_3 | R_2, S) = .62 \times (14/62) = .14 \quad (10-34)$$

A combined (“or”) probability is (since the set C_i is mutually exclusive)

$$p(C_1 \text{ or } C_2 | S) = p(C_1 | S) + p(C_2 | S) = .32 + .51 = .83 \quad (10-35)$$

Events, Variables, and Values

Many events can be defined by a qualitative description. For instance, “The president of the United States in 2017 will be a Democrat” is a description of a possible event.

In many quantitative situations, however, the event is defined by a variable or parameter taking on a specific value or having a value within a specified range. For instance, in the Medequip example, events have definitions such as “Unit Cost less than or equal to \$1,000.”

In quantitative situations of this sort, it is convenient to define events by the value the variable takes on—rather than by defining the event by a range of values. Values can be discrete (the variable can take on any one of a finite number of different values) or continuous (the variable can take on any value out of a continuum). Some variables have discrete possible values, such as the marginal income tax rate. Most variables, however, have a continuum of possible values, such as unit cost. Unfortunately, continuous values are difficult to work with, and we will always approximate a continuous set of values by a few discrete values. A good process for making this approximation is discussed in Chapter 2.

To make it clear that an event is defined by a value, we will write the value as a lowercase letter rather than the uppercase letters we have been using for general definitions.

Probabilities and the values associated with them are often referred to as a probability distribution for the variable in question.

Medequip’s planning department proceeded to assign a value to represent each of the ranges in the definitions of the events. For unit cost in 2007, they defined the following values:

- c_1 —\$800
- c_2 —\$1,250
- c_3 —\$1,700

For unit revenue in 2007, they defined the following values:

- r_1 —\$1,500
- r_2 —\$2,000

This yielded the following set of values for the margin in 2007:

- $m_1 - \$700 = \$1,500 - \$800$
- $m_2 - \$250 = \$1,500 - \$1,250$
- $m_3 - \$200 = \$1,500 - \$1,700$
- $m_4 - \$1,200 = \$2,000 - \$800$
- $m_5 - \$750 = \$2,000 - \$1,250$
- $m_6 - \$300 = \$2,000 - \$1,700$

Representations of Probabilities for Discrete Values _____

There are a number of ways to represent the probabilities for a set of discrete values. These representations include tabular form, tree form, mass distribution plot, cumulative probability plot, and histogram. We will illustrate each of these representations using the example of the margins from the Medequip case.

Tabular Form

The values and their probabilities are presented in the following simple tabular form:

i	m_i	$p(m_i S)$
1	\$700	.10
2	\$250	.25
3	-\$200	.03
4	\$1,200	.22
5	\$750	.26
6	\$300	.14

Tree Form

These values can be presented in the tree form displayed in Figure 10-12.

Mass Density Plot

Probabilities can be graphed directly against values in what is called, for discrete values, the mass density plot. The plot in Figure 10-13 graphs the probability distribution on margin.

Cumulative Probability Graph


Another graphical representation frequently used is the cumulative probability graph. This graph plots the probability that the value is less than or equal to the value shown on the horizontal axis. Formally, the cumulative probability is defined as

$$P_{\leq}(x | S) = \sum_i p(x_i | S) \quad (10-36)$$

where the sum is over all values of i for which $x_i \leq x$. In practice, the outcomes are placed in a table ordered from lowest to highest value; the cumulative

Figure 10–12

Representation of Probabilities and Values in Tree Form




probability at each value is calculated as the running sum of the probabilities down to each row. For the margin example, we have the following table:

m_i	$p(m_i S)$	$P(m_i S)$
-\$200	.03	.03
\$250	.25	.28
\$300	.14	.42
\$700	.10	.52
\$750	.26	.78
\$1,200	.22	1.00

Figure 10–13

Representation of Probabilities and Values in a Mass Density Plot



The cumulative probability is then plotted (Figure 10–14). Note that the curve is flat except at points where the value on the horizontal axis is equal to one of the values in the table.

Histogram

A final representation is by histogram form. Although this plot is an approximation, it takes the mass density plot and converts it into a form that is more readily interpreted by the eye. The horizontal axis is divided into bins of equal width, and a bar for each bin shows the sum of the probabilities for all the events in that bin. In Figure 10–15, bins of a width of \$400 were chosen, with bin edges falling at values that are integer multiples of \$400.

Mean, Median, Mode, Variance, and Standard Deviation

Another way to represent a set of probabilities is to use a few values to characterize the whole set. A common measure is the mean (or average) of the distribution. In decision analysis, the mean is usually called the expected value. If x_i for $i = 1, 2, \dots, m$ is a set of values that define a set of mutually exclusive and collectively exhaustive events, the mean is calculated using the

Figure 10–14
Representation of Probabilities and Values in a Cumulative Probability Plot





Figure 10–15

Representation of Probabilities and Values in a Histogram



following formula:

$$\text{mean} = \sum_{i=1}^m x_i p(x_i | S) \quad (10-37)$$

The median is the value at which it is just as probable to be below that value as above it. The mode is the most probable value in the set. The mean, mode, and median are measures that identify the center or most important values in the distribution. Of these measures, though, the mean is most sensitive to the shape of the distribution. The median is often the easiest measure to understand (it is the 50/50 probability point). The mode often works well for describing highly asymmetric distributions (such as the cost). Neither the mode nor the median is very important in applications of probability theory to decision analysis, though they may be important in helping people relate to and understand a probability distribution.

The mean of the probability distribution on Medequip's margin in 2007 is calculated as follows:

$$\text{mean} = (700 \times .10) + (250 \times .25) + (-200 \times .03) + (1200 \times .22) + (750 \times .26) + (300 \times .14)$$

$$\text{mean} = \$628 \quad (10-38)$$

The median is somewhere around \$700, since there is a 42 percent chance that the value is less than \$700 and a 48 percent chance that the value is greater than \$700. (In a set of values this small, there will seldom be a value that is exactly the median.) Finally, the mode of the distribution is \$750, since it has the largest probability of any of the values (26 percent).

The variance is a more complicated measure. If the mean measures the average value of the distribution, the variance is a measure of how far off the value might be from this average; it is a measure of how “wide” the distribution is.

$$\text{variance} = \sum_{i=1}^m (x_i - \text{mean})^2 p(x_i | S) \quad (10-39)$$

The standard deviation (written as σ) is the square root of the variance and is a direct measure of the spread or width of the distribution.

The planning department at Medequip calculated the variance in the probability distribution on margin in 2007 as follows. Note that the units of variance here are (\$)².

$$\begin{aligned} \text{variance} = & .10(700 - 628)^2 + .25(250 - 628)^2 + .03(-200 - 628)^2 + .22(1,200 - 628)^2 \\ & + .26(850 - 628)^2 + .14(300 - 628)^2 \\ \text{variance} = & 147,719 \end{aligned} \quad (10-40)$$

The standard deviation is the square root of the variance.

$$\sigma = (147,719)^{1/2} = \$384 \quad (10-41)$$

Moments and Cumulants

A more complete and systematic set of parameters to represent a probability distribution is provided by the set of moments. If x_i for $i = 1, 2, \dots, m$ is a set of values that define a mutually exclusive and collectively exhaustive set of events, then the moments are defined by the following equation:

$$\mu_n = \sum_{i=1}^m x_i^n p(x_i | S) \quad (10-42)$$

The zeroth moment, μ_0 , is 1; this is just the normalization condition that probabilities sum to one. The first moment, μ_1 , is the mean of the distribution. These moments are sometimes called the raw moments and provide a complete description of the distribution.

There is a second set of moments that can also be used to describe the distribution—the central moments. For the same values, x_i , the central moments are defined by the following equation:

$$v_n = \sum_{i=1}^m (x_i - \mu_1)^n p(x_i | S) \quad (10-43)$$

The zeroth central moment, v_0 , is just 1; again, this is just the normalization condition for probabilities. The first central moment, v_1 , is identically zero. The second central moment, v_2 , is the *variance*. The third central moment, v_3 , is the *skewness*. The fourth central moment, v_4 , is the *kurtosis*. The central moments, together with the mean, provide a complete description of the distribution.

Occasionally, it is necessary to transform one representation into the other. In practice, the first five moments are all that are likely to be used in representing a distribution. The mean, μ_1 , is common to both representations. The equations transforming raw moments into central moments are:

$$\begin{aligned}v_2 &= \mu_2 - \mu_1^2 \\v_3 &= \mu_3 - 3\mu_2\mu_1 + 2\mu_1^3 \\v_4 &= \mu_4 - 4\mu_3\mu_1 + 6\mu_2\mu_1^2 - 3\mu_1^4 \\v_5 &= \mu_5 - 5\mu_4\mu_1 + 10\mu_3\mu_1^2 - 10\mu_2\mu_1^3 + 4\mu_1^5\end{aligned}\quad (10-44)$$

The equations transforming central moments into raw moments are:

$$\begin{aligned}\mu_2 &= v_2 + \mu_1^2 \\ \mu_3 &= v_3 + 3v_2\mu_1 + \mu_1^3 \\ \mu_4 &= v_4 + 4v_3\mu_1 + 6v_2\mu_1^2 + \mu_1^4 \\ \mu_5 &= v_5 + 5v_4\mu_1 + 10v_3\mu_1^2 + 10v_2\mu_1^3 + \mu_1^5\end{aligned}\quad (10-45)$$

A third set of parameters used to represent a distribution is provided by the set of cumulants, c_i . The special property of cumulants is that if we wish to add two probabilistically independent distributions, we can simply add the cumulants of each distribution to obtain the cumulants of the sum.*

The first three cumulants are the mean, variance, and skewness, and for many purposes these are sufficient. It would be very unusual to need more than the first five cumulants:

$$\begin{aligned}c_1 &= \mu_1 \\ c_2 &= \mu_2 - \mu_1^2 \\ c_3 &= \mu_3 - 3\mu_2\mu_1 + 2\mu_1^3 \\ c_4 &= \mu_4 - 4\mu_3\mu_1 + 12\mu_2\mu_1^2 - 3\mu_2^2 - 6\mu_1^4 \\ c_5 &= \mu_5 - 5\mu_4\mu_1 - 10\mu_3\mu_2 + 20\mu_3\mu_1^2 - 60\mu_2\mu_1^3 + 30\mu_2^2\mu_1 + 24\mu_1^5\end{aligned}\quad (10-46)$$

$$\begin{aligned}c_2 &= v_2 \\ c_3 &= v_3 \\ c_4 &= v_4 - 3v_2^2 \\ c_5 &= v_5 - 10v_3v_2\end{aligned}\quad (10-47)$$

*The theory of moments and cumulants can be found in Alan Stuart and J. Keith Ord, *Kendall's Advanced Theory of Statistics, Volume 1 Distribution Theory*, Sixth Edition, New York-Toronto:Halsted Press, 1994.

In analyzing portfolios, it is often necessary to transform from cumulants to moments. The equations are:

$$\begin{aligned}\mu_1 &= c_1 \\ \mu_2 &= c_2 + c_1^2 \\ \mu_3 &= c_3 + 3c_2c_1 + c_1^3 \\ \mu_4 &= c_4 + 4c_3c_1 + 3c_2^2 + 6c_2c_1^2 + c_1^4 \\ \mu_5 &= c_5 + 5c_4c_1 + 10c_3c_2 + 10c_3c_1^2 + 15c_2^2c_1 + 10c_2c_1^3 + c_1^5\end{aligned}\quad (10-48)$$

$$\begin{aligned}v_2 &= c_2 \\ v_3 &= c_3 \\ v_4 &= c_4 + 3c_2^2 \\ v_5 &= c_5 + 10c_3c_2\end{aligned}\quad (10-49)$$

In analyzing portfolios, it is also often necessary to represent graphically the probability distribution characterized by these moments. See Problem 10-24.

Representations of Probabilities for Continuous Variables _____

Throughout this book, we use discrete approximations instead of continuous values. In the interest of completeness, however, we discuss a few properties and definitions for probability distributions on continuous variables. Continuous probability distributions are usually defined by graphs or by a functional form rather than using an infinite tabular form. For instance, the normal or Gaussian probability distribution is defined by the following equation:

$$p(x | \mu_1, v_2, normal) = \frac{1}{\sqrt{2\pi v_2}} e^{-(x-\mu_1)^2 / 2v_2} \quad (10-50)$$

The state of knowledge is completely specified in this case by the value of the mean, μ_1 , the value of the variance, v_2 , and the knowledge that it is a normal distribution. The value of the function $p(x | S)$ is the probability that the value falls in the range between x and $x + dx$. In the graph of the normal probability distribution (Figure 10-16), σ is the standard deviation, the square root of the variance. In Microsoft Excel, the normal distribution is given by `NORMDIST(x, μ_1, v_2, k)`, where if k is TRUE the cumulative probability at x is given, and if FALSE the probability at x is given.

The shaded area equals the probability that x lies within one standard deviation, σ , of the mean, μ_1 . For the normal or Gaussian distribution in Figure 10-16, this probability is 68 percent.

The cumulative probability graph is smooth (instead of the staircase found with discrete distributions) and is defined by the following integral:

$$P_{\leq}(x | S) = \int_{-\infty}^x p(x' | S) dx' \quad (10-51)$$

Figure 10-16





Figure 10-17

Cumulative Probability Graph for the Normal or Gaussian Probability Distribution



For the normal or Gaussian distribution, the cumulative probability distribution is as shown in Figure 10–17.

Note that for a normal or Gaussian distribution, $84 - 16 = 68$ percent of the probability lies within one standard deviation of the mean and $98 - 2 = 96$ percent of the probability lies within two standard deviations of the mean. Also, the 10/90 width is 2.56σ for this distribution.

The moments of a probability distribution for a continuous variable are calculated by the equations below.

$$\mu_n = \int_{-\infty}^{\infty} x^n p(x|S) dx \quad (10-52)$$

$$\nu_n = \int_{-\infty}^{\infty} (x - \mu_1)^n p(x|S) dx \quad (10-53)$$

As you can imagine from the equations above, the use of continuous distributions almost always leads to problems of overwhelming analytical complexity for real situations. For this reason, discrete approximations are almost always used in actual decision analysis applications.

Problems and Discussion Topics

- 10.1 State whether each of the following statements is or is not an event and if not, why not.
- The temperature will be greater than 53°F in Bombay tomorrow.
 - General Motors stock sold for more than \$55 per share on the New York Stock Exchange.
 - The weather is nice today.
 - The final version of the first printing of this book contains 1,034,246 characters.
 - Elm trees are taller than oak trees.
- 10.2 Draw the Venn diagram for the following events: 1,000 - Units Sold < 5,000; Profit Margin per Unit < \$2; Profit Margin per Unit \leq \$2. In which region(s) of the Venn diagram do the following events occur? (Profit = Units Sold \times Profit Margin.)
- Margin = \$2 and Profit = \$4,000
 - Margin = \$1.50 and Profit = \$2,000
 - Profit < \$2,000 and Sales = 500 units
 - Profit = \$15,000
- 10.3 The New England Patriots and the Cincinnati Bengals both have one game left in the season. They are not playing each other, and each game will go into overtime if necessary to produce a winner.
- Draw the Venn diagram for this situation.

If the Patriots win and the Bengals lose, then the Patriots go to the play-offs. If the Patriots lose and the Bengals win, then the Bengals go to the play-offs. Otherwise, you are not sure who goes to the play-offs.

- b. Show the region on the Venn diagram where you know the Patriots go to the play-offs for certain.

The Miami Dolphins also have one game left in the season and are playing neither the Patriots nor the Bengals. If Miami wins and both the Patriots and the Bengals lose, then the Patriots go to the play-offs. If Miami loses and both the Patriots and Bengals win, then the Patriots also go to the play-offs. Otherwise, the Miami game is not relevant.

- c. Redraw the Venn diagram and show the regions where the Bengals go to the play-offs for sure.

- 10.4 For each of the following events or list of events, complete the list to make it mutually exclusive and collectively exhaustive.

- a. The number of passenger automobiles assembled in the United States in 1985 was at least 20 million and less than 22 million.
- b. The average length of all great white sharks reported caught off the coast of Australia is less than 15 feet.
- c. The variable unit cost for producing the product is \$1.75.
- d. The market demand for adipic acid is at least 100 pounds per year and less than 150 million pounds per year. The market demand for adipic acid is greater than 400 million pounds per year.
- e. A competitive product is introduced before our product is introduced.
- f. Our market share is twice that of our nearest competitor.

- 10.5 A soldier is taking a test in which he is allowed three shots at a target (unmanned) airplane. The probability of his first shot hitting the plane is .4, that of his second shot is .5, and that of his third shot is .7. The probability of the plane's crashing after one shot is .2; after two shots, the probability of crashing is .6; the plane will crash for sure if hit three times. The test is over when the soldier has fired all three shots or when the plane crashes.

- a. Define a set of collectively exhaustive events.
- b. Define a set of mutually exclusive events.
- c. Define a set of mutually exclusive and collectively exhaustive events.
- d. What is the probability of the soldiers shooting down the plane?

- e. What is the mean number of shots required to shoot down the plane?

10.6 Define the joint events for the following two sets of events.

- m_1 —Market Share < 5 percent
 m_2 —5 percent - Market Share < 10 percent
 m_3 —10 percent - Market Share
 d_1 —Development Cost - \$2 million
 d_2 —\$2 million < Development Cost - \$5 million
 d_3 —\$5 million < Development Cost

10.7 For the data in the preceding problem, assume that the events have been approximated by discrete values as follows.

- m_1 —Market Share = 3 percent
 m_2 —Market Share = 7 percent
 m_3 —Market Share = 13 percent
 d_1 —Development Cost = \$1.5 million
 d_2 —Development Cost = \$3 million
 d_3 —Development Cost = \$7 million

- a. Define the joint set of events.
 b. If Market Size = \$100 million and Revenue = (Market Size × Market Share) - Development Cost, calculate the Revenue for each joint event.
 c. Calculate the Revenue for each joint event given a Market Size of \$60 million.

10.8 On the air route between Chicago and Los Angeles, there is either a head wind or tail wind. Depending on which way the wind is blowing and how fast, flights from Chicago to Los Angeles may be early, on time, or late. We define the following events:

- w_1 —Head Wind
 w_2 —Tail Wind
 a_1 —Arrive Early
 a_2 —Arrive on Time
 a_3 —Arrive Late

The joint probabilities are as follows:

- w_1 and a_1 — .06
 w_1 and a_2 — .12
 w_1 and a_3 — .22

w_2 and a_1 — .39

w_2 and a_2 — .18

w_2 and a_3 — .03

- a. What is the marginal probability of a head wind?
- b. What are the conditional probabilities for arriving early, on time, or late given a tail wind?
- c. Given you arrive on time, what is the probability you had a tail wind? If you arrive early? If you arrive late?

- 10.9 The Surprise Dog man at Fenway Park sells all his hotdogs for the same price, but he does not tell you in advance what you are getting. You could receive a regular dog or foot-long dog, either of which could be a cheese or chili dog. We define the following events:

l_1 —You get a foot-long dog

l_2 —You get a regular dog

c_1 —You get a cheese dog

c_2 —You get a chili dog

The marginal probability of getting a foot-long dog is .25. The probability of getting a foot-long chili dog is .225, and the probability of getting a regular cheese dog is .45.

- a. What is the marginal probability of getting a cheese dog?
- b. What is the probability of getting a regular chili dog?

- 10.10A weather forecaster said that San Francisco and Los Angeles have probabilities of .7 and .4, respectively, of having rain during Christmas day. She also said that the probability of their both having rain is .28.

- a. Find the probability of rain in San Francisco on Christmas day given rain in Los Angeles on Christmas day.
- b. Find the probability of rain in Los Angeles on Christmas day given rain in San Francisco on Christmas day.
- c. Find the probability of rain in San Francisco or Los Angeles (or both) on Christmas day.

- 10.11 Your resident expert on Soviet deployments, Katyusha Caddell, has just given you his opinion on recent Soviet missile developments. The Soviets are building silos that may be of type 1 or type 2 (it is too early to tell), and Katyusha is unsure about which of two possible missile types the Soviets will be deploying in them. He describes the following events:

s_1 —Silo of type 1 built

s_2 —Silo of type 2 built

m_1 —Type 1 missile deployed

m_2 —Type 2 missile deployed

Katyusha puts the probability of the silos being type 2 at .6 and figures that type 2 silos mean a .7 probability of type 1 missiles, while type 1 silos mean a .8 probability of type 2 missiles. He further puts the marginal probability of type 2 missile deployment at .6. Do the marginal probabilities agree?

- 10.12 You and a friend are pondering buying tortilla chips and salsa at a baseball game. Your friend tells you he has made a systematic study of the different varieties of salsa and says the possible types are salsa tomatillo, salsa fresca, and traditional salsa. (The workers at the snack bar do not know what kind it is.) Furthermore, the salsa could be hot (spicy) or not hot. Your friend makes the following predictions:

The chance of hot salsa tomatillo is .08.


The chance of hot salsa fresca is .15.

The chance of not-hot traditional salsa is .18.

The chance of getting salsa fresca is .3 and of getting traditional salsa is .6.

- a. What is the probability of getting not-hot salsa tomatillo?
 - b. What is the conditional probability that the salsa is hot, given that it is salsa tomatillo?
 - c. What is the marginal probability that the salsa is hot?
 - d. What is the conditional probability of getting traditional salsa, given that it is not hot?
- 10.13 Frequently, people use tests to infer knowledge about something. A current (controversial) example is the use of a blood test to see if a person has the AIDS virus or not. The test results reflect current knowledge of the virus' characteristics, and test accuracy may be a matter of concern. How should the information represented by the blood test result be used to update knowledge of the test subject's condition? Bayes' Rule gives the answer to this question.

Suppose a number of people have taken an XYZ virus test with the result shown below. (The numbers are purely illustrative and are not intended to reflect current understanding of the AIDS blood test.)



If a person has taken this test and the result turns out to be positive, what is the probability that he or she does not have the XYZ virus?

(Testing for the AIDS virus involves serious issues of rights to privacy and due process. This problem addresses only the information gained by using a test where the outcome of the test is not a perfect indicator of the underlying condition.)

- 10.14 Your professor tells you that only 50 percent of the students in her class will do the homework and pass the class; 25 percent will not do the homework and will still pass the class; 8.3 percent will do the homework and study too much (missing sleep) and still pass. The professor figures 30 percent will not do the homework, 60 percent will do the homework, and 10 percent will work too much.

According to the professor, are doing the homework and passing the class probabilistically dependent or independent?

- 10.15 You suspect that your corns hurt when your mother is about to call you. However, you think that the chance of getting a call from your mother and your corns not hurting is about .5. Your corns hurt about 10 percent of the time.

What is the marginal probability of your mother calling if her calling and your corns hurting are probabilistically independent?

- 10.16 Use the information from problem 10.7 to perform the following calculations:
- Formulate the joint events and calculate the probabilities and revenues for them. Assume probabilities .25, .50, and .25 for m_1 , m_2 , and m_3 and for d_1 , d_2 , and d_3 , respectively.
 - Plot the cumulative probability distribution for revenue.

- c. Plot the histogram for revenue. (Choose the bin size to give a good representation of the data.)
- d. Calculate the mean, variance, and standard deviation of the distribution.
- 10.17 The annual revenues from a new gasoline additive depend on annual U.S. gas consumption and on the average price of gasoline over the next year. It is estimated that 1 bottle of additive will be sold for every 1,000 gallons of gasoline consumed. The price will be set at twice the price for a gallon of gas. Discretized estimates of U.S. gas consumption next year put a .3 chance on consumption being 1 billion gallons, a .6 chance on consumption being 1.5 billion gallons, and a .1 chance on consumption being 2 billion gallons. Similarly, average gas prices have a .25 chance of being \$0.50, a .5 chance of being \$1.00, and a .25 chance of being \$1.25.
- a. Formulate a probability tree for revenue.
- b. Calculate the probabilities and revenues for the joint events.
- c. Plot the cumulative probability distribution for revenue.
- d. Plot the histogram for revenue.
- e. Calculate the mean, variance, and standard deviation of the distribution.
- 10.18 You are offered an opportunity to engage in a series of three coin flips (.5 probability of winning or losing). For the first flip, you would bet \$1 and either double your money or lose it. For the second flip (if you had won the first flip), you would reinvest your \$2 and either double your money or lose it; if you had lost the first flip, you would bet another \$1 and double or lose it. The process is repeated for the third coin flip, with having either the money you won on the second flip or a new \$1 investment if you lost the second flip.
- a. Draw the probability tree for the three coin flips.
- b. Calculate your winnings or losses for each joint event and the associated probabilities.
- c. Plot the cumulative probability distribution for your proceeds from the flips (wins or losses).
- d. Plot the histogram for your proceeds from the flips.
- e. Calculate the mean, variance, and standard deviation for your proceeds from the flips.
- 10.19 Explain graphically why the following relationships are true for the events I and \emptyset :

$$I \text{ and } \emptyset = \emptyset$$

$$I \text{ or } \emptyset = I$$

Use these relationships and the probability axioms to prove the following probability:

$$p(\emptyset | S) = 0$$

- 10.20 Let A and B be any two events and let $A' = \text{not } A$. Explain graphically why the following relations are true:

$$\begin{aligned} A \text{ or } B &= A \text{ or } (A' \text{ and } B) \\ B &= (A \text{ and } B) \text{ or } (A' \text{ and } B) \\ A \text{ and } (A' \text{ and } B) &= \emptyset \end{aligned}$$

Use these relationships, the results of problem 10.19, and the probability axioms to prove the following relationship among probabilities:

$$p(A \text{ or } B | S) = p(A | S) + p(B | S) - p(A \text{ and } B | S)$$

- 10.21 Let A and B be any two events and let $B' = \text{not } B$. Explain graphically why the following relationships are true:

$$\begin{aligned} A &= (A \text{ and } B) \text{ or } (A \text{ and } B') \\ \emptyset &= (A \text{ and } B) \text{ and } (A \text{ and } B') \end{aligned}$$

Use these relationships and the probability axioms to prove the following simple application of the Expansion Theorem:

$$p(A | S) = p(A, B | S) + p(A, B' | S)$$

- 10.22 Let A and B be any two events. Assume that A is probabilistically independent of B :

$$p(A | B, S) = p(A | S)$$

Prove that B is probabilistically independent of A .

$$p(B | A, S) = p(B | S)$$

- 10.23 Conditional probability is defined as:

$$p(A | B, S) = \frac{p(A, B | S)}{p(B | S)}$$

Show that the definition of conditional probability satisfies the three axioms of probability introduced in Equations 10-10, 10-11, and 10-12.

- 10.24 Suppose that the moment technique described at the end of Chapter 7 has been used to evaluate a business portfolio. The mean of the distribution is \$230 million, the variance is 18,212 in units of (\$ million)², and the skewness is 5,000,000 in units of (\$ million)³.

- a. Large numbers like this are difficult to interpret. More convenient are the standard deviation, σ , and the skewness coefficient, skewness/ σ^3 . A skewness coefficient outside the range -2 to 2 means the distribution is quite skewed. What is the standard deviation and skewness coefficient for this portfolio?

- b. Decision-makers relate better to graphics than to numbers like variance and standard deviation. Use the normal distribution (Equation 10–50) to plot a cumulative probability distribution that has the given mean and variance.
- c. Creating a plot that matches mean, variance, and skewness is more difficult. Use the shifted lognormal distribution

$$\frac{1}{\sqrt{2\pi s(x-x_0)}} e^{-\frac{1}{2}\left(\frac{\ln(x-x_0)-m}{s}\right)^2}$$

to plot a cumulative probability distribution that matches the given mean, variance, and skewness. (The more familiar, non-shifted lognormal distribution has $x_0 = 0$.)

- d. Compare the results of b and c. How different are the distributions? How would the plot differ if skewness were 50,000,000? 500,000?

Hint: Use Microsoft Excel or another spreadsheet program to create the graph. Plot the distribution from (mean $- 2\sigma$) to (mean $+ 2\sigma$).

For the lognormal distribution with $x_0 = 0$, m and s are the mean and standard deviation of $\ln(x)$. The moments of x are given by

$$\mu_1 = e^{m + \frac{1}{2}s^2}$$

$$\mu_k = \mu_1^k e^{\frac{1}{2}k(k-1)s^2} = e^{km + \frac{1}{2}k^2s^2}$$

The following calculations can be used to obtain x_0 , m , and s from the moments μ_1 (mean), v_2 (variance), and v_3 (skewness). All that is required to obtain these results is some tedious algebra and solving a cubic equation to determine d . First calculate

$$w = 2v_2^3 / v_3^2$$

$$r = \sqrt{1 + 2w}$$

$$d = \frac{v_2^2}{v_3} + v_2 \frac{(1+w+r)^{1/3} + (1+w-r)^{1/3}}{(2v_3)^{1/3}}$$

$$A = v_2 + d^2$$

$$B = v_3 + 3dv_2 + d^3$$

These values can then be used to calculate the parameters we need:

$$m = \frac{3}{2} \ln(A) - \frac{2}{3} \ln(B)$$

$$s = \sqrt{\frac{2}{3} \ln(B) - \ln(A)}$$

$$x_0 = \mu_1 - d$$

Note that if the skewness is negative, you need to use the reflected shifted lognormal. To do this, use the absolute value of the mean μ_1 and skewness v_3 in the above calculations, and then substitute $(x_0 - x)$ for $(x - x_0)$ in the equation for the shifted lognormal above.

11

Influence Diagram Theory

Theory Overview

The decision facilitator learns early that decision trees grow with amazing rapidity. Add a three-branch node, and the tree becomes three times larger. Add a few more nodes, and you need a wall-size piece of paper to represent the tree. And yet you are discussing only a relatively small number of uncertainties.

Influence diagrams are a means of representing the same decision problem much more compactly. For each chance and decision variable in the tree, there is a single object in the graph. Arrows connecting these objects represent probabilistic and informational relationships.

Influence diagrams are popular and important in the practice of decision analysis for three reasons.

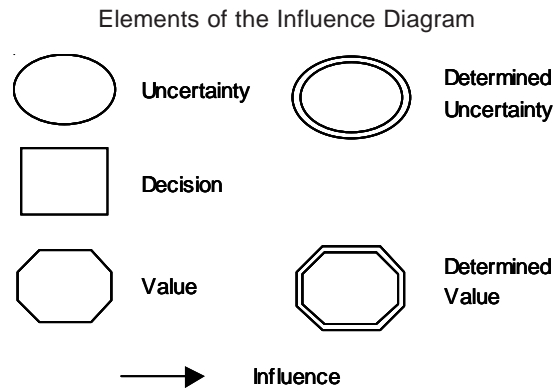
- The influence diagram is compact enough that even the most complex problems can be developed and discussed on a single large sheet of paper. This is a tremendous advantage when structuring the problem, organizing analytical tasks, monitoring problem analysis, and presenting an overview of the problem.
- Influence diagrams appear to be the easiest way to introduce and work with probabilistic dependence, an unfamiliar and difficult concept for many decision-makers.
- An influence diagram is a theoretical construct and an evaluation device with all the power of a decision tree.

In this book, we use the influence diagram for the first and second reasons.

Elements of Influence Diagrams

Six basic elements are used in an influence diagram (Figure 11-1).

Figure 11-1



To define each of these elements, we will use the fictitious Medequip example introduced in Chapter 10.

Medequip's planning department continued its study of the product's declining revenues. The planners were considering changing the long-term pricing strategy for the product. Two alternatives suggested themselves. Medequip could choose a Premium Pricing strategy: promote the special characteristics of the product, target specific segments of the market, and set a price 20 percent above the principal competitor's price. Or it could choose a Market Pricing strategy: set price at the competitor's price and compete in all segments of the market.

Critical to the competitive position of Medequip's product was the availability of an essential raw material. It was suggested that a raw material survey could be initiated that could help predict future availability and cost of the raw material. Information like this could help Medequip enter into advantageous long-term supply contracts.

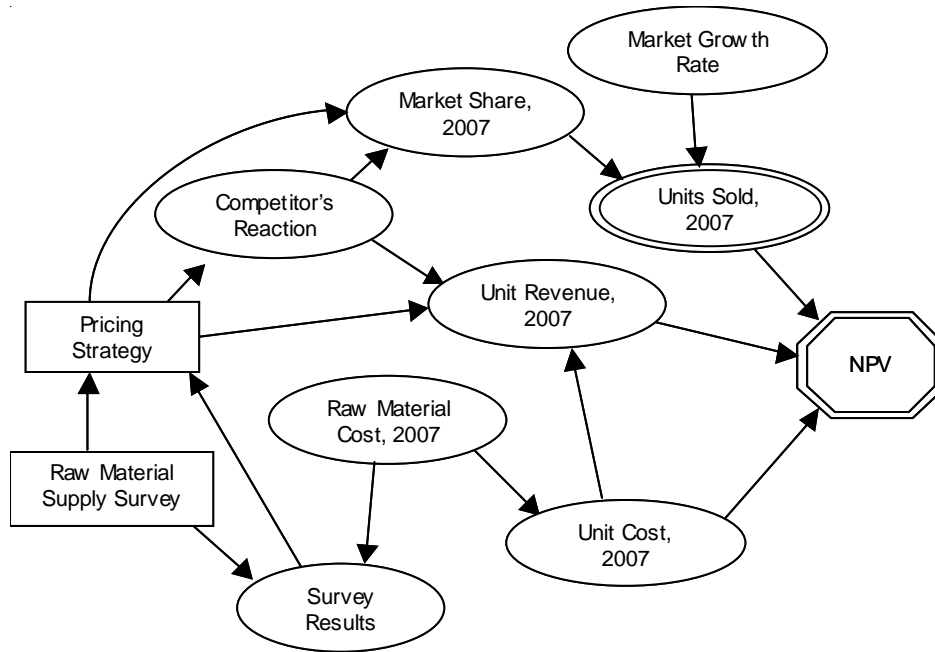
Medequip brought together a number of people from management and planning and created an influence diagram (Figure 11-2) that represented their perception of the problem.

Uncertainty

Uncertainties are represented by ovals. Within the oval is a label that indicates the mutually exclusive and collectively exhaustive set of events among which the node distinguishes. The most desirable label is a variable name, where the value of the variable defines the events. An oval represents a set of possible events and the probabilities assigned to these events.

Figure 11-2

Influence Diagram for Medequip's Decision Problem



The planning department was uncertain about how rapidly the market for the product would grow. The oval labeled Market Growth Rate in the influence diagram represented the uncertainty on the average annual market growth rate during the period until 2007. After the influence diagram had been completed, several of those most knowledgeable about the market assessed the probability distribution represented by the oval. The information contained “inside” the oval is shown in Figure 11-3. This tree representation for the data at a node is called a distribution tree and represents a probability distribution on a single set of events.

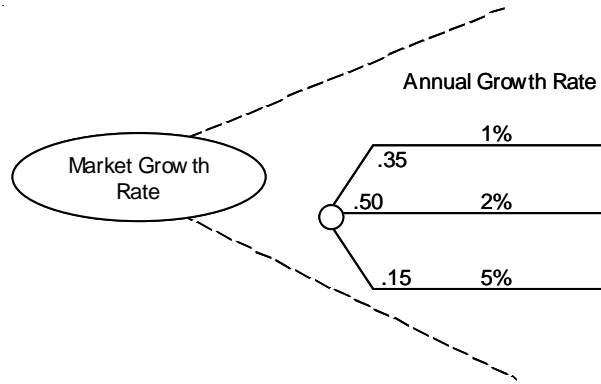
Decision

Decisions are represented by rectangles. Written in the rectangle is a label for the set of significantly different alternatives being considered.

Representatives from manufacturing were concerned about the cost of the critical raw material used in making the product. It was suggested that a raw material survey would help predict availability and cost of the raw material, and this in turn would have an impact on product unit cost. The rectangle labeled Raw Material Supply Survey represents the decision on whether to perform this survey (Figure 11-4).

Figure 11-3

Distribution Tree for the Data Contained in the Market Growth Rate Node



Influence

Influences are indicated by arrows and represent a flow of information and, with decision nodes, a time sequence. As discussed in Chapter 10, all probability assignments are based on the state of knowledge of a particular person at a particular time. Decisions are also made based on a state of knowledge. It is essential that all the probability assignments and decision nodes in an influence diagram share a common state of information, S . The arrow indicates that, in addition to this common state of information, S , there is information concerning the node at the base of the arrow available at the node at the head of the arrow. This concept is made more explicit in the four cases presented below.

Figure 11-4

Alternatives Represented by the Raw Material Supply Survey Node


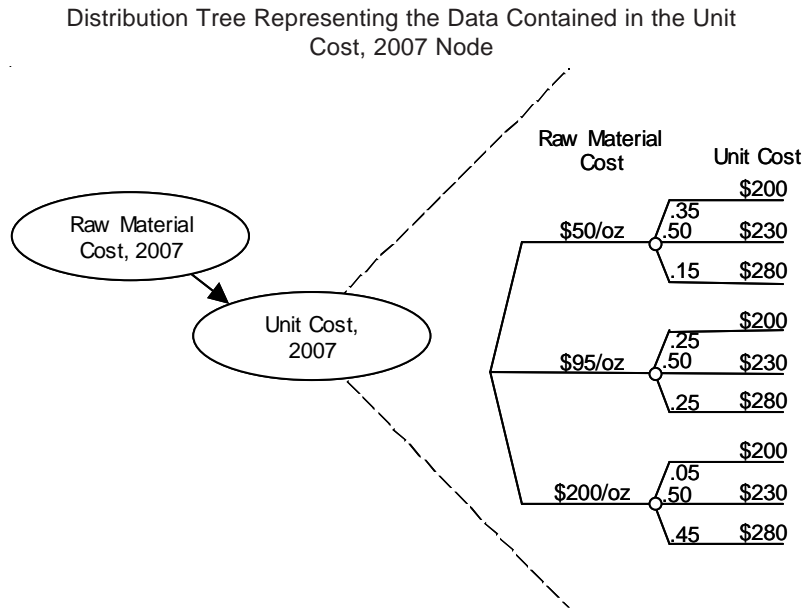


Figure 11-5



Arrow from Uncertainty Node to Uncertainty Node

If the arrow is between two uncertainty nodes, the probability distribution for the node at the head of the arrow is probabilistically dependent (conditional) on the node at the base of the arrow. If there is an arrow from node A to node B, the diagram is read as “A influences B.”* The information flow in this case does not necessarily imply a causal link or a time sequence.

Recently, the term “relevance” has begun to replace the term “influence” (which has a causal connotation) in this context. If there is an arrow from node A to node B, the diagram is read as “A is relevant to B.” Influence diagrams that contain only uncertainties are referred to as “relevance diagrams.”

Representatives of manufacturing felt much more comfortable assigning probabilities to Unit Cost, 2007 conditional on the cost of raw material around that time. Accordingly, they drew an arrow between the Raw Material Cost, 2007 node and the Unit Cost, 2007 node. If A_i and C_i are a set of mutually exclusive and collectively exhaustive events describing Raw Material Cost and Unit Cost, respectively, and if S is the state of knowledge common to all the nodes in the influence diagram, then the probability distribution represented by the node Unit Cost, 2007 is

$$p(C_i | A_i, S)$$

At a later meeting, the probability assignments (Figure 11-5) were assessed by the representatives of manufacturing.

* If there are arrows from several nodes to node B, it is this set of nodes (rather than each individual node) that influences node B.

The tree representation for the conditional probability distribution above is a more complex form of distribution tree than the one shown in Figure 11-4. The difference between a distribution tree and a decision tree is that the nodes on the left side of a distribution tree are there only to specify the state of knowledge used in assessing the probabilities for the node on the right. They can be arranged in whatever order helps the assessor assign probabilities. The order of nodes in a distribution tree is not necessarily related to the order of nodes in the decision tree.

Arrow from Decision Node to Uncertainty Node

If the arrow is from a decision node to an uncertainty node, the probability distribution for the node at the head of the arrow is probabilistically dependent (conditional) on the alternative chosen at the node at the base of the arrow. This implies that the decision is made before the uncertainty is resolved and that there is some sort of causal link between the decision and the resolution of the uncertainty.

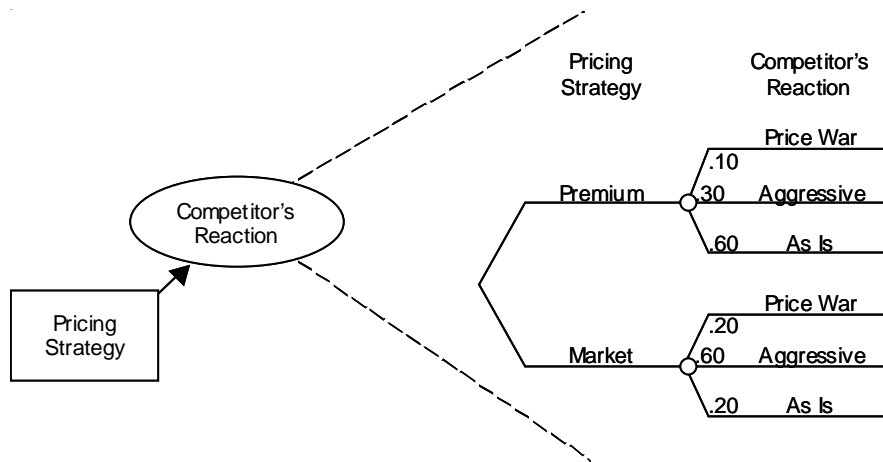
The representatives of the marketing department were uncertain just how their principal competitor would react to either of the pricing strategy alternatives, but they were sure that the probabilities they assigned would depend on the pricing strategy chosen. For this reason, an arrow was drawn from the Pricing Strategy node to the Competitor's Reaction node. The assessed probability distribution is shown in Figure 11-6.

Arrow from Decision Node to Decision Node

An arrow between two decision nodes means that the decision-maker remembers which alternative was chosen at the node at the base of the arrow when he or she comes to make the decision at the node at the head of the arrow. There is a strong chronological assertion here: the decision represented

Figure 11-6

Distribution Tree for Data Contained in Competitor's Reaction Node

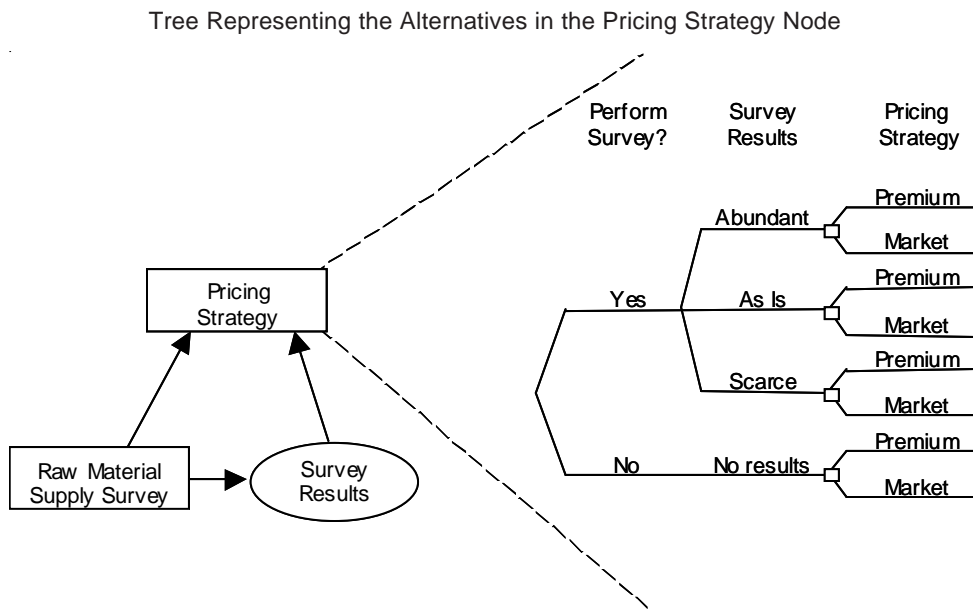


by the node at the base of the arrow will be made before the decision represented by the node at the head of the arrow.

Arrow from Uncertainty Node to Decision Node

If the arrow is from an uncertainty node to a decision node, the uncertainty is resolved before the decision is made and the decision-maker learns what happened before making the decision. Note that, in this case, there is a strong assertion about chronology: the uncertainty is resolved and the information received by the decision-maker before the decision is made.

Figure 11-7



The Medequip team decided that a decision on whether to perform the raw material supply survey would be made before the pricing strategy decision. Furthermore, it was decided that if a raw material supply survey were decided on, the pricing strategy decision would not be made before the results of the survey were available. The decision-maker would know the survey results when he or she made the decision. An arrow was drawn from the Raw Material Supply Survey node to the Pricing Strategy node and from the Survey Results node to the Pricing Strategy node (Figure 11-7).

The last example shows one way to represent an asymmetry in an influence diagram. The set of events represented by the node may differ, depending on the event occurring at a node that influences it. In this case, the events represented by the Survey Results node include a single event (no results), which occurs if the survey is not performed.

Determined Uncertainty

As an influence diagram grows and influences (arrows) are added, the uncertainty at a node frequently disappears: once the outcomes are known at all the influencing uncertainty and decision nodes, there is no more uncertainty about the actual event at the influenced node. The node can be left in the diagram, denoted by a double oval. This type of node is usually called a “deterministic node.” (The node can also be removed from the diagram, provided you assure the proper information flow by connecting all incoming arrows to all outgoing arrows.)

A deterministic node usually represents either a formula or calculation for which the influencing nodes supply input. Often this calculation is complicated enough that it is made by a computer spreadsheet or other computer program. Occasionally, the deterministic node represents a table of values, one entry for each combination of the events of the nodes that influence it.

One of the first uncertainties the Medequip team identified was the uncertainty represented by the node Units Sold, 2007. As the influence diagram grew, two nodes were created (Market Growth Rate and Market Share, 2007), both of which influenced the node Units Sold, 2007. At this point, there is no uncertainty left at the node Units Sold, 2007: once you know what happened at the influencing nodes, all that remains is a simple calculation (Figure 11–8).

The node Units Sold, 2007 could be removed from the diagram if the arrows coming into it are rerouted to the node that it influences—that is, if arrows are drawn from the Market Growth Rate node to the NPV node and from the Market Share, 2007 node to the NPV node—thus assuring that information flows are maintained.

Value and Determined Value

Implicit in an influence diagram is a rule by which decisions are made. As discussed earlier in the text, decisions are made by choosing the alternative that maximizes the certain equivalent of the value measure. In simpler terms, we choose the alternative that gives us the most of what we want.

Figure 11–8

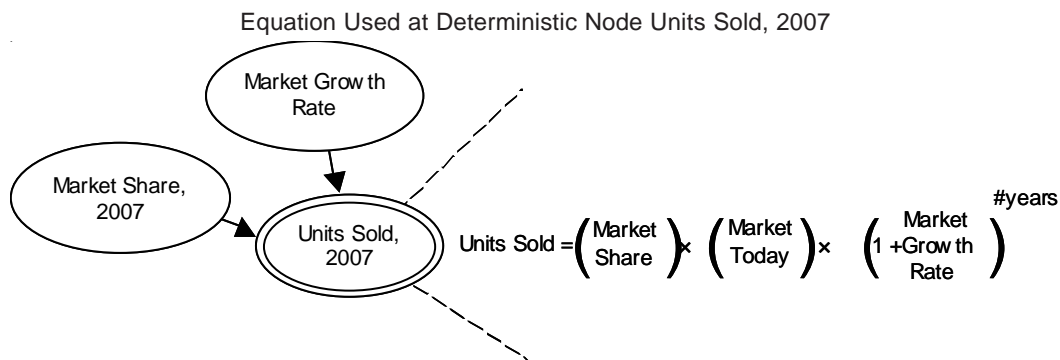
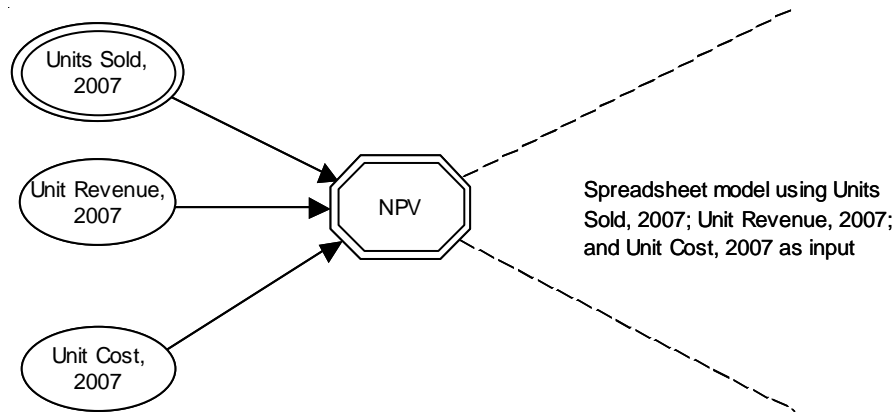


Figure 11–9

Model Used at Deterministic Node NPV



If there are decision nodes in an influence diagram, there must be one (and only one) node in the influence diagram that represents the value measure used in making the decisions. This value node is always an uncertainty node or a deterministic (determined uncertainty) node. To identify its role, an octagon (double octagon if it is deterministic) is used rather than an oval.

Very early, the Medequip team had chosen NPV of cash flow as the value measure used in making decisions. Given the outcomes of the influencing nodes (Units Sold, 2007; Unit Revenue, 2007; and Unit Cost, 2007), it was felt that most of the uncertainty in NPV would be resolved. For this reason, the node representing NPV was represented by a double octagon (Figure 11–9).

Rules for Constructing Influence Diagrams

The following four rules must be obeyed to create meaningful influence diagrams for the type of decision problems dealt with in this book:

1. *No Loops*—If you follow the arrows from node to node, there must be no path that leads you back to where you started. This is most readily understood in terms of the information flow indicated by the arrows.
2. *Single Decision-Maker*—There should be just one value measure (octagon), and all decisions should be made to maximize the same function (expected value or certain equivalent) of this value.

3. *No Forgetting Previous Decisions*—There should be an arrow between all pairs of decisions in the diagram. This will establish the order in which the decisions are made and also indicate that the decision-maker remembers all his or her previous choices.
4. *No Forgetting Previously Known Information*—If there is an arrow from an uncertainty node to a decision node, there should be an arrow from that uncertainty node to all subsequent decision nodes.

Procedures for Manipulating Influence Diagrams _____

There are two useful procedures for manipulating influence diagrams without changing the information in the diagram. These procedures are necessary when solving an influence diagram directly, but may be useful in restructuring a diagram to facilitate drawing a decision tree.

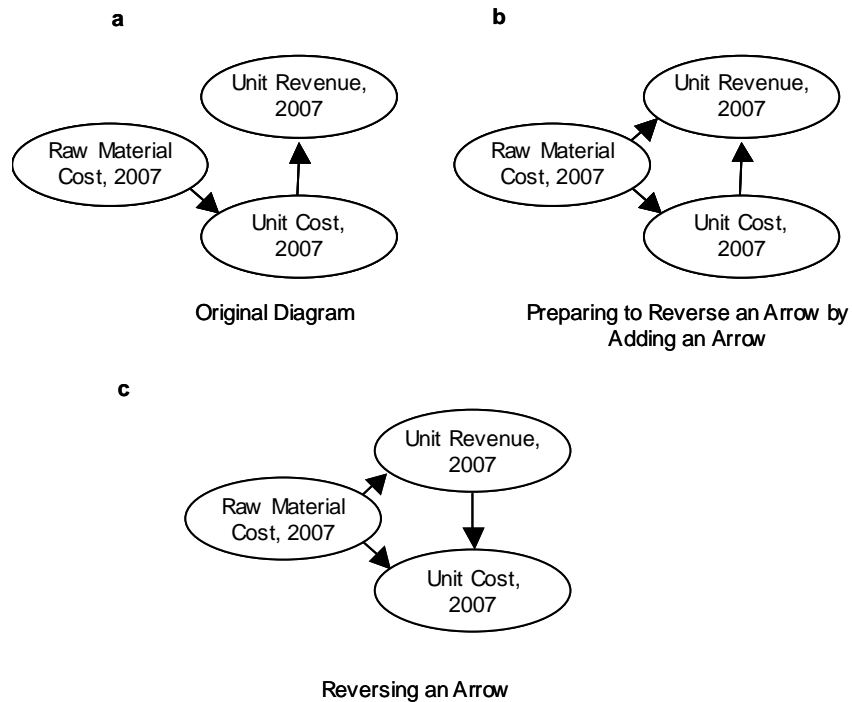
1. *Adding an Arrow*—An arrow can be added between two uncertainty nodes in an influence diagram, provided it does not create a loop. The lack of an arrow between two nodes is a statement that information from one node is not needed to assess the probabilities at the other. Adding an arrow can be thought of as making this information available, even though it does not affect the probabilities assessed.
2. *Reversing an Arrow*—An arrow between two uncertainty nodes can be reversed if the state of knowledge available at both nodes is the same. Therefore, to reverse an arrow between two nodes, all the arrows into both nodes must be the same (except, of course, the arrow between the two). Once the state of information is the same, the operation of reversing an arrow in an influence diagram is the same operation as applying Bayes' Rule to a probability distribution and "flipping" a probability tree.

The Medequip team had recognized that the easiest and most reliable means of assessment was first to have procurement personnel provide probabilities for raw material cost, then to have manufacturing estimate unit cost (conditional on raw material cost), and finally to have sales and marketing provide probabilities for unit revenue conditional on unit cost (Figure 11–10a).

There was a desire to see whether assessing the probabilities in a different order made a difference. First, sales and marketing would estimate unit revenues, and then manufacturing would estimate unit cost conditional on unit revenues. (There was some feeling that perhaps cost rose when revenue was high and was squeezed down when revenue was low.) This assessment entails reversing the arrow between Unit Cost, 2007 and Unit Revenue, 2007. Before this can be done, an arrow from the Raw Material Cost node to the Unit Revenue node must be added (Figure 11–10b) so that both the Unit Revenue and Unit Cost nodes share the same state of knowledge. (For this exercise, the arrows from Pricing Strategy and Competitor's Reaction to Unit Revenue, 2007 were dropped from the diagram.)

Figure 11–10

Manipulating Influence Diagrams



Now the arrow between the Unit Revenue and Unit Cost nodes can be reversed (Figure 11–10c). At this point, the team realized that the job of the sales and marketing representatives was harder than they had thought. Unit Revenue probabilities had to be assessed conditional on Raw Material Cost!

Turning an Influence Diagram into a Decision Tree

The procedure for turning an influence diagram into a decision tree is quite simple.

1. Arrange the decision nodes such that all arrows with a decision node at their base or head point to the right-hand side of the page. Arrows pointing to or emanating from decision nodes imply a chronology that must be followed in decision trees. By convention, the chronology of decisions in a decision tree flows from left to right, and therefore these arrows must point from left to right.

2. Arrange the uncertainty nodes so that no uncertainty node is to the left of a decision node unless there is an arrow from that node to the decision node. In a decision tree, the outcome of a node to the left of a decision node is known to the decision-maker when he or she makes the decision; in an influence diagram, this means there is an arrow from the uncertainty node to the decision node.
3. Arrange, insofar as possible, the uncertainty nodes so that all arrows point to the right. This will cause conditional probabilities to be displayed simply on the tree.

The rules for manipulating influence diagrams can be used to reverse arrows between nodes and make all arrows point to the right. However, at least one decision tree program (Supertree) can accept input in which probabilities at a node depend on nodes that follow it in the tree—that is, for which the arrow to the node points left—and so this last step of arrow reversal is not necessary.

4. Make the deterministic value node a tree endpoint node. This will usually involve calculations to find the value associated with each combination of events at the nodes that influence the value node.
5. Number the nodes, give each node a node name, and input the structure and data into Supertree.

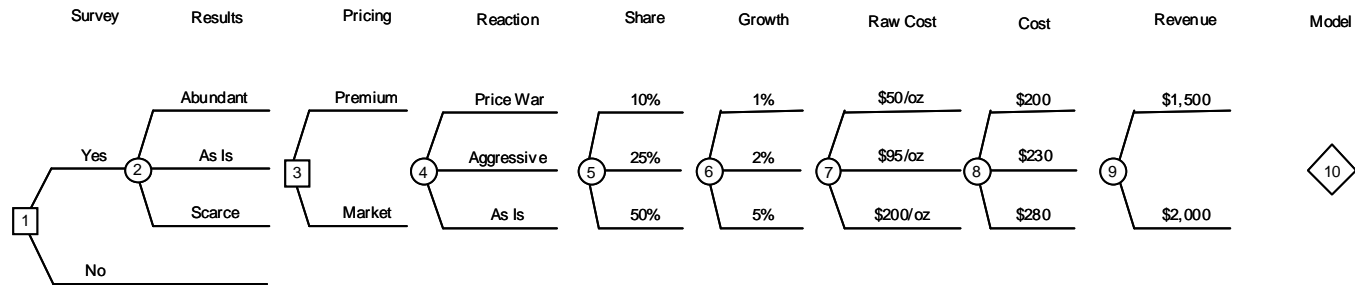
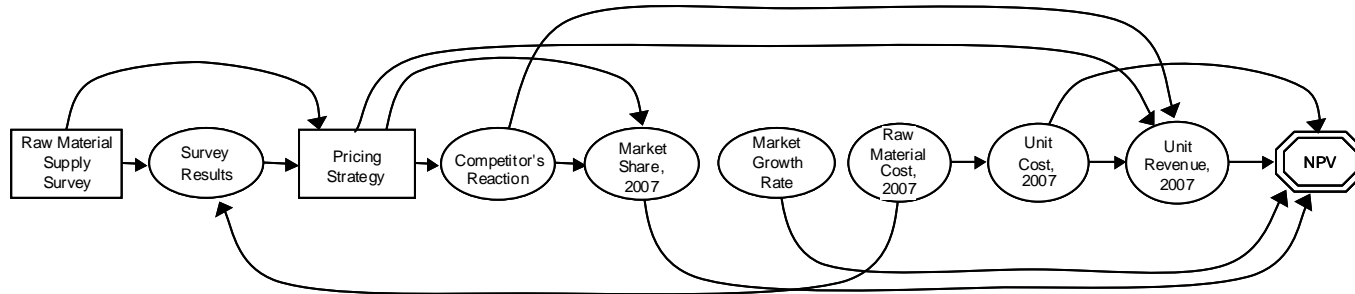
The Medequip facilitators decided to put the influence diagram into tree form. First, they eliminated the Units Sold, 2007 determined uncertainty node, rerouting the two arrows coming into it to the node it influenced, the NPV node. Following the rules given above, they rearranged the diagram and formed the tree shown in Figure 11–11. Since the tree would be entered in Supertree, they did not need to reverse the arrow from the Raw Material Cost, 2007 node to the Survey Results node because Supertree does this automatically.

The tree is ready to be evaluated. All the information necessary for inputting the nodes is shown except for the probabilities. The probabilistic dependence is shown by the arrows in the influence diagram. For instance, the probabilities of node 9 depend on nodes 3 and 4, and the probabilities for node 2 depend on node 7. The endpoint, node 10, depends on nodes 5, 6, 8, and 9 and will probably be calculated by a spreadsheet. Note the difference in orientation between “influences” and “depends on”: if node 3 influences node 9 in the influence diagram, then in the language often used for trees, the probabilities at node 9 depend on node 3.

Particular care should be taken to distinguish the use of the word “successor” between influence diagrams and decision trees. In influence diagrams, it is quite natural to say that a node at the head of an arrow is the successor to the node at the base of the arrow. Thus, in Figure 11–11, the node Survey Results is the successor (direct successor, to be precise) of the node Raw Material Cost, 2007. In decision trees, successor refers only to the order in the decision tree and does not imply any information flow or dependence between the nodes. Thus, in the tree in Figure 11–11, Market Growth Rate is the successor to Market Share, 2007.

Figure 11-11

Turning the Medequip Influence Diagram into a Decision Tree



Summary


This example illustrates the power of the influence diagram to represent a problem compactly and to lead people naturally through the complexities of probabilistic dependence. There are 3,888 paths through this tree; at 66 lines per page, the printout of the full tree would be 59 pages long. The dependent probabilities weave a complex pattern through the tree. On the other hand, for the influence diagram we need only one page to draw the diagram and several other pages on which to list the events and probabilities represented by the nodes.

In practice, influence diagrams are currently used mostly in the initial phases of a decision analysis. The process of model construction, data gathering, and sensitivity analysis normally leads to insight and simplification of the structure of the problem and to the natural construction of a simple tree. However, given the rapid evolution of software tools and consulting practice, we can expect to see the unification of these two representations of the problem in a single system.

Problems and Discussion Topics


- 11.1 How does an influence diagram contribute to making good decisions? (Refer to the elements of a good decision.) What elements of a good decision does an influence diagram not help with and why?
- 11.2 Describe at least one way that using influence diagrams helps you draw better decision trees and one way that being familiar with decision trees helps you draw better influence diagrams.
- 11.3 How do you know when an influence diagram has become complicated enough? Relate your answer to the problem of assessing probabilities and to the clairvoyance test.
- 11.4 Think of a significant decision you have made. Draw the influence diagram for that decision. Were there significant uncertainties? How did you identify and deal with them at the time? Do you have any new insights into the decision? (Relate this last answer to the good decision/good outcome distinction.)
- 11.5 Draw the influence diagram for the date and time when a specific close relative walks through your front door. Make sure the uncertainty passes the clairvoyance test and try to summon all your information and experience on the factors influencing the uncertainty. Has the exercise changed your understanding of the uncertainty at all? Could you now draw a decision tree and do a meaningful probability assessment? Draw the tree and explain what (if anything) would prevent you from assessing probabilities and calculating the expected date and time that the relative walks through your door.

- 11.6 Draw an influence diagram for the probability of a major war within the next ten years. Make sure your uncertainties pass the clairvoyance test. How is this problem different from the previous one? Is there anything preventing you from drawing a tree and calculating a probability distribution for this problem?
- 11.7 Draw the influence diagram for the number of times you eat pizza within the next month. Again, make sure the uncertainty passes the clairvoyance test. Are there any difficulties in completing this problem and, if so, what are they?
- 11.8 In the influence diagram used to construct the tree in Figure 11–11, there is an arrow pointing to the left.
 - a. Reverse this arrow to make the diagram a “decision tree network,” one in which the nodes can be arranged so that all the arrows point to the right.
 - b. Why is it necessary to reverse this diagram to create a tree? (Hint: How would you display the probabilities at node 2 in the tree?)
- 11.9 Adding an arrow between an uncertainty and a decision node is related to the value of information calculation described in chapters 2 and 4.

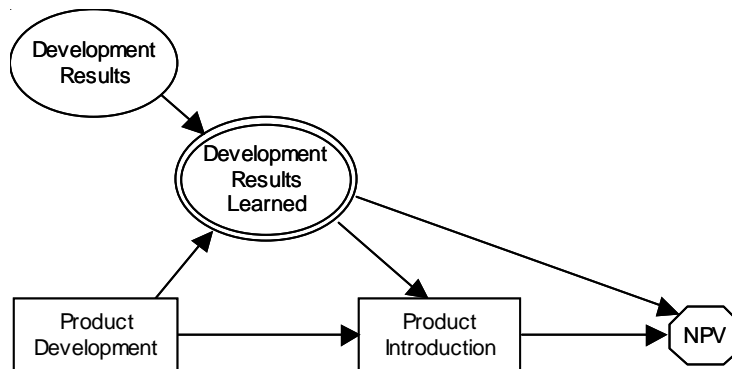


- a. Draw the trees represented by influence diagrams A and B above.
 - b. How are influence diagrams A and B related to the value of perfect information on Market Growth Rate?
- 11.10 In the Howard canonical form of an influence diagram, there are no arrows from a decision node to any uncertainty node aside from the value node. For the purposes of this definition, groups of uncertainty nodes can be amalgamated into a larger uncertainty node (the value node) provided no loops are created.

A company has several different routes it could pursue in developing a new product. The influence diagram representing its problem is shown below.




a. Is the influence diagram above in Howard canonical form? The one below?



- b. Suppose some preliminary work could predict the results of the product development effort before all the necessary development work was done. Which influence diagram could be used to calculate the value of information about development results?
- c. Which influence diagram is in Howard canonical form? Draw a tree showing the logic contained in the deterministic node that would make the second diagram equivalent to the one that preceded it.
- d. Can the second influence diagram be manipulated to a form from which the value of information about development results can be calculated?

11.11 It is possible to use deterministic nodes to represent asymmetries in the problem in a straightforward way.



- Redraw the influence diagram in Figure 11–11 using the deterministic node defined above. What arrows should be drawn from these nodes to the remainder of the diagram?
 - What is the logic contained in the deterministic node Survey Results Learned?
 - Is the tree drawn from the new diagram different from the tree in Figure 11–11?
 - Make this diagram into a decision tree network. (See problem 11.8 for the definition of a decision tree network.)
 - Is the influence diagram in Figure 11–11 in Howard canonical form? (See problem 11.10.)
 - Is the influence diagram drawn as part of this problem in Howard canonical form?
- 11.12 In Chapter 10, a joint probability distribution was given for the two sets of mutually exclusive and collectively exhaustive events, R_i and C_j .

	$P(R_i, C_j)$		
	C_1	C_2	C_3
R_1	.10	.25	.03
R_2	.22	.26	.14

- (R, C)
- R
- C
- $R \rightarrow C$
- $C \rightarrow R$

Draw the distribution trees for the nodes in the five influence diagrams

above.

11.13 In Figure 11-2, there is no arrow going to the Market Growth Rate node. The lack of arrows is of great significance to the facilitator, since the absence of arrows makes modeling and probability assessment relatively simple.

- a. There is no arrow between the nodes Pricing Strategy and Market Growth Rate. What does this indicate about the nature of the market and Medequip's place in the market? Under what circumstances should an arrow be drawn between these two nodes?
- b. There is no arrow between the nodes Competitor's Reaction and Market Growth Rate. What does this indicate about the nature of the market and the companies that supply products in this marketplace? Under what circumstances should an arrow be drawn between these two nodes?
- c. Under what circumstances should there be an arrow from both the Pricing Strategy and Competitor's Reaction nodes to the Market Growth Rate node?
- d. We are not allowed to draw an arrow from Pricing Strategy to Competitor's Reaction, from Competitor's Reaction to Market Growth Rate, and from Market Growth Rate to Pricing Strategy (perhaps to represent a pricing adjustment to changes in market dynamics). Why is this not allowed? How might you represent an adjustment of pricing strategy to market dynamics?

12

Encoding a Probability Distribution

Although most facilitators are familiar with data-gathering methods and problems, decision analysis uses two types of data that require quite different methods to gather and pose quite different sorts of problems. One type of data is the utility curve. (The methods and conditions for encoding utility curves are discussed in Chapter 5.) The other type of data is the probability distribution, which is used throughout decision analysis. Encoding a probability distribution can be one of the most difficult tasks to perform well in decision analysis. This chapter provides a review of some of the problems encountered in encoding a probability distribution and an outline of the encoding process.

Levels of Detail in Encoding Probability Distributions _____

Ideally the most important uncertainties in an analysis will be described by probability distributions obtained through the full encoding process: many points for the cumulative probability distribution assessed from some expert. For really important uncertainties, probability distributions may be assessed from several experts.

However, practice over the past decade has shown the importance of obtaining quality in obtaining the 10/50/90 values for uncertain variables. The 10/50/90 values describe the uncertainty in terms of three numbers: a low value (10% chance the true value will turn out to be less than this value), a base value (50% chance the true value will turn out to be less than this value), and a high value (90% chance the true value will turn out to be less than this value, and 10% chance it will turn out to be higher). These values give three points on the cumulative probability curve (see Chapter 2) and are used in the early stages of analyzing the problem (see probabilistic sensitivity analysis in Chapter 6).

There is a temptation to obtain these numbers quickly and carelessly and assume that more care can be spent on encoding the truly important uncertainties later. However, in practice the second encoding session rarely happens. It is well worth a minor investment of time to make sure that, however briefly, all the steps of the encoding session are touched on in assessments of 10/50/90 ranges.

Problems in Encoding Probability Distributions _____

In Chapter 2, we stated that since probability is used to represent a person's state of knowledge about something, there are no correct or incorrect probabilities. However, two major questions arise in encoding and using probability distributions.

The first question concerns choosing a person to furnish the probability distribution. Does that person have an adequate state of relevant knowledge? Are sufficient and correct experience and data available; does the person have the training and intelligence to assimilate this experience and data? These questions are familiar to everyone who wonders whom to believe or rely on.

Decision analysis does not offer a new solution to this problem; it simply ensures that the decision-maker considers the person chosen to be knowledgeable or expert in the relevant area. Usually, the facilitator and decision-maker can agree easily on who the appropriate experts are. If there are several possible experts, the facilitator can encode probability distributions from each of them and compare the effect the distributions have on the problem. Different probability distributions may or may not lead to different decisions.

The second question is whether the encoded probability distributions really represent the state of knowledge of the persons involved. A surprising number of biases can sneak in (often at an unconscious level) and make carelessly encoded probability distributions inadequate representations of the subject's state of knowledge. Decision analysis handles this problem by devising methods to deal with the biases that inevitably arise.

Motivational Biases _____

Motivational biases arise when the encoded probabilities do not reflect the expert's conscious beliefs. The cause of this behavior is usually desire for a reward or fear of punishment, which can be economic, psychological, or physical. An example of this type of bias can occur when a salesman is asked for an estimate of future sales. Sales personnel are often conditioned to respond artificially high (possibly to obtain the rights to sell the product) or artificially low (perhaps to appear good when they exceed their estimates). Another example of motivational bias occurs when people give estimates based on "real, objective data" (which may have limited relevance), rather than take the responsibility to express their own judgment and intuition.

The long-term cure for motivational biases is to align the reward structure to encourage truthful responses. In the short term, however, the facilitator

can show the expert the importance to the company of his responses and emphasize that he is recording a state of knowledge—not asking for a prediction or commitment. If the expert has a personal stake in having a particular course of action chosen, the facilitator can divide the quantity being assessed into several subfactors, which brings the assessment down to a level of detail where the expert does not know how to give answers that serve his self-interest. As a last resort, if the above cures do not appear to be working, we can disqualify the expert.

Three forms of behavior can be loosely classified as motivational biases: the manager's bias, the expert's bias, and the facilitator's bias.

The manager's bias is typified by the statement: "If that's what the boss wants, I'll get it done!" In this case, the encoded uncertainty may be very small, while the manager has shifted all the uncertainty to some other unsuspected factor, such as the cost of getting it done.

The expert's bias is typified by the attitude that experts are expected to be certain of things rather than uncertain. Given this attitude, an expert will give a narrow distribution. The obvious errors associated with this sort of reasoning should be discussed with the expert.

The facilitator's bias results from not wanting to appear in disagreement with official company information. It often occurs in an organization with an official in-house forecast. This forecast, which may merely be the result of the latest computer run, is intended primarily to impose consistency on company analyses. Except at top-management levels, it is rare to find anyone who is willing to disagree publicly with this forecast. The encoder must show the expert the origins, limitations, and purpose of the forecast and encourage candid responses. (The in-house forecast can also create a powerful anchor, as discussed below.)

Cognitive Biases

How do people assess the probability of an uncertain event? It turns out there are a limited number of heuristic methods used in this process. (A heuristic involves or serves as an aid to learning, discovery, or problem-solving by experimental, and especially, trial-and-error methods.) In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors.* Improperly using the cognitive heuristics causes biases and leads to incorrect conclusions.

An example of a heuristic is the use of visual clarity to judge the distance of objects: hazy objects are thought to be distant. An example of an incorrect use of this heuristic occurs in the clear air of deserts where distant objects appear to be quite near. Although biases tend to occur at an unconscious

* There is much literature on the subject. A seminal paper in this area is the excellent and readable article: A. Tversky and D. Kahneman, "Judgment Under Uncertainty: Heuristics and Biases," *Science* 185 (September 27, 1974): 1124-1131. This article is reprinted in *Readings on the Principles and Applications of Decision Analysis*, ed. R.A. Howard and J.E. Matheson, 2 vols. (Menlo Park, California: Strategic Decisions Group, 1984) 2: 903-910.

level, they are correctable. For example, in a desert, one learns to say “Objects are farther away than they seem!”

Availability Bias

One type of common bias in probability encoding is called the availability bias. In this case, the heuristic is to think of occurrences of an event (or of similar events); the easier it is to think of occurrences, the more likely the event is judged to be. For instance, if eight out of ten small businesses you can think of failed within their first year, you may judge it quite likely that a new small business will fail. This heuristic is quite useful and generally works well. However, a problem occurs because of quirks in the way we store and retrieve information in our minds. For instance, information that is ordered, redundant, dramatic, recent, imaginable, certified, or consistent with our world view is much more likely to be stored and recalled than information that does not have these qualities. As a result, probabilities tend to be unduly weighted by the more available information.

We can counteract this bias in several ways. First, we can check to find out if there are reasons some information is more readily available than other information. If there are, we can discuss and research the less available information. Another cure (which is a powerful cure for other biases as well) is to postulate extreme outcomes and request the expert to give some scenarios that would lead to these outcomes. This technique forces the expert to search his or her mind for little-used information. A final cure (also valuable in confronting other biases) is to have the expert pretend to consider the situation retrospectively from the future. A change of perspective often opens the mind up to new possibilities.

Representativeness Bias

A second type of bias is called the representativeness bias. The heuristic in force here is to look at some evidence and then assign a high probability to the event that makes the evidence most likely. What can go wrong? Studies show that when faced with specific evidence, people often tend to discard (or undervalue) previous experience and general information. For instance, a sharp decrease in a company’s net income might lead someone to assign a relatively high probability that the company is in financial trouble. However, mature reflection might reveal that the particular industry has net incomes that tend to vary wildly from year to year and that the latest change has little significance.

Another example of the representativeness bias is the common error of stereotyping: “All people who do X are Y. Therefore, anyone who is Y is very likely to do X.” Very few people stop to consider the possibility that there may be many people who are Y and only a few cases of people who do X.

Besides being very common, this bias can be quite subtle and complex, both in its manifestations and in its causes. (For more discussion of this bias, see the article by Tversky and Kahneman.) The cure for this bias is to separate prior or general information from new evidence and to update prior informa-

tion explicitly using probability theory. (See the treatment of Nature's tree in Chapter 4.)

Adjustment and Anchoring Bias

Adjustment and anchoring give rise to a third common bias. The heuristic is to start from some initial estimate and then adjust the estimate until it seems reasonable. This heuristic is the basis of many engineering or other quantitative estimates. Unfortunately, the adjustment is rarely adequate, and the final estimate tends to be anchored near the initial estimate—no matter how arbitrary the initial estimate was. We see this bias a great deal when encoding probability distributions. For example, people sometimes start by making a best estimate and then attempt to estimate how uncertain the best estimate is, thus arriving at the width of the distribution. What happens is that they get anchored by the best estimate and almost always wind up with too narrow a probability distribution. Other examples of this bias are anchoring on the corporate plan and corporate forecasts—it is difficult to think of futures that are too different. The cure for this bias is to start by discussing extreme outcomes and asking the expert for explanatory scenarios. If the facilitator must start with an initial best estimate, he or she should try to do so in a way that does not commit the expert.

A special case of the anchoring and adjustment bias arises because of conjunctive distortions. We tend to overestimate the probability of success for conjunctive events (where all the individual events must happen for success), because we anchor on the probability of each individual event happening. This bias is particularly important in R&D, where a number of hurdles must be cleared before success is achieved. The obvious cure for this bias is to obtain the probabilities for each event happening and then to multiply them to obtain the probability of success. This type of distortion also occurs when we underestimate the probability of success for disjunctive events (any individual event happening gives success). This bias occurs in research and development when we tend to underestimate the probability of technical success in following parallel but independent research paths.

Implicit Conditioning Bias

Finally, there is the implicit conditioning bias. One form of the heuristic is to tell a story; the more coherent and plausible the story, the more likely we judge the outcome we are trying to explain. The bias is that we can forget to examine the likelihood of the enabling events needed to start the story. Another form of the heuristic occurs when we make certain assumptions to make the assessment process easier. The bias arises when we subsequently forget these assumptions exist. The cures for this type of bias are, first, to postulate extreme outcomes and request explanatory scenarios and, second, to examine probabilities of any enabling events for the scenarios. It also helps to have a checklist of some common assumptions (no fire, strikes, lawsuits, war, competitive breakthrough, etc.). You can also identify hidden assump-

tions by asking the expert what he or she would like to insure against. One way to verify that the results are now accurate is to ask the expert if he or she would invest personal funds using the probability distribution just encoded.

Probability Encoding Process

After the preceding discussion, it may seem virtually impossible to get a reliable probability distribution. However, the process described below incorporates cures for the common biases. The process has a strong effect—experts usually adjust their answers considerably when it is used and usually in the “right” direction of broader probability distributions. Moreover, we can obtain a high degree of consistency in encoded distributions: points tend to lie along smooth curves, and the degree of repeatability in responses is high. Finally, experts have considerable confidence in the results generated through the process compared with those obtained through other methods. As a result, they typically become strong supporters of the analysis.

The process related here (originally described in the 1972 joint ORSA-TIMS-AIEE national meeting in Atlantic City*† is typically conducted in five stages. While the process may seem long and very detailed, many of the steps take only a few seconds, and the process is natural enough that the details are not obvious.

It is usually best to conduct the encoding in private, so that group pressures do not bias the results. After the initial encoding, however, a group review frequently proves useful in resolving differences and sharing information.

Stage 1: Motivating

The motivating stage establishes the necessary rapport with the subject and explores whether a serious potential for motivational biases exists. After explaining his delegated authority to conduct the interview, the encoder explains the decision problem, describes the decision model that has been constructed, and shows how all the uncertain factors will be accounted for in the analysis. This ensures that the subject can clearly focus on the task at hand.

Next, the encoder explains the importance of accurately assessing uncertainty on the quantity and emphasizes that the intent is to measure the expert’s knowledge and best judgment—not to predict the value. The importance of emphasizing this distinction depends on how much the encoder

* Carl S. Spetzler and Carl-Axel S. Staël von Holstein, “Probability Encoding in Decision Analysis,” *Management Science*, Vol. 22, No. 3 (November 1975): 340-358. A slightly different version of this article is reprinted in *Readings on the Principles and Applications of Decision Analysis*, ed. R.A. Howard and J.E. Matheson, 2 vols. (Menlo Park, California: Strategic Decisions Group, 1984) 2: 601-625.

† Much of the material in this section is drawn from work done with M.W. Merkhofer, “Quantifying Judgmental Uncertainty: Methodology, Experiences, and Insight,” *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-17 (Sept.-Oct. 1987): 741-752.

detects the possibility of the manager or expert having the biases mentioned above.

Finally, the encoder discusses the expert's personal involvement with the decision and with the variable being encoded to identify any asymmetries in the payoff structure that might encourage the expert to bias his or her estimates high or low. Note-taking by the encoder is useful because it encourages a more balanced presentation of issues.

The conversation covering the motivating stage for encoding product cost for a new product might resemble the following:

As you may know, the VP of marketing has asked us to look into the possibility of developing and marketing the following new product. One of the important considerations in making this decision is product cost, and your boss has asked us to consult you in this area. More specifically, we are looking into the following alternatives for the new product, and here is how product cost fits in to our analysis. [A simplified version of the strategy table and influence diagram can help at this point.] What we are looking for is an estimate of the uncertainty in what product cost will be over the life of the product; at this point, nobody is asking for predictions or budget commitments. By the way, are you involved in this decision? And are you the right person to supply these estimates?

Stage 2: Structuring

The structuring stage has two purposes: to structure the variable clearly and to explore how the subject thinks about it. The first step is to precisely define the variable whose uncertainty is to be assessed, using, for example, the clairvoyance test described in Chapter 2. The next step is to explore the possibility of decomposing the variable into more elemental quantities and then to assess those elemental quantities individually. The third step is to elicit and list all assumptions the subject is making in thinking about the variable. To identify hidden assumptions, it is useful to ask the subject what he or she would like to insure against. Finally, the encoder selects an appropriate measuring scale—most importantly one that uses the units most familiar to the expert. (One unit to avoid is a growth rate. Encoding the uncertainty on compound growth rates is difficult, because few people adequately appreciate the effects of compounding over long periods of time. It is better to encode the uncertainty on the actual value of the quantity at the end of the growth period and then calculate the implied growth rate.)

The conversation covering the structuring stage might resemble the following:

The variable we hoped to assess is product variable cost, excluding depreciation, fixed costs, and allocated costs. Is this the way you think of product cost, or are you more familiar with reports of full product cost? How do we factor in things like the possibility of supply interruption or labor difficulties? Do you think in terms of dollars/unit or in terms of raw material/unit; if the latter, from whom should we get raw material costs? What is the best way to think of the evolution of this cost over time: Initial cost declining to some steady state cost? What of inflation?

Stage 3: Conditioning

In the conditioning stage, the expert draws out the expert's relevant knowledge about the uncertain variable. Often, discussion will indicate that the expert is basing his or her judgment on both specific information and general information. Given the problems of the representativeness bias, it may be appropriate to encode probabilities from prior information and to update this prior information with the specific information. (See Chapter 4.) Fortunately, it is usually adequate to educate the expert about the representativeness bias.

The next step is to counteract the anchoring and availability biases, something the encoder can do by eliciting extreme values of the variable from the expert and then asking for scenarios that would explain these outcomes. (At this point, hidden assumptions are often uncovered.) The encoder should also explore for the presence of anchors, such as corporate plans or forecasts.

This step is very important. Working with the subject to explore how extreme scenarios may occur (and perhaps even writing down the reasons why) broadens and unlocks (unbiases) people's thinking in an almost magical way.

Another useful method is to explain or demonstrate what we call the 2/50 rule. In many seminars, we have asked attendees to assign probability distributions to the answers to questions drawn from an almanac. Usually, the specific question is to give a range such that the correct answer has only a 20 percent chance of falling outside that range. In other words, if people correctly perceived their uncertainty about the answer, 20 percent of the time the answers would be expected to fall outside the 10 percent and 90 percent fractiles they defined. Over the 25 years we have been giving this demonstration, we have found that over 50 percent of the answers fall outside this range. People tend to think they know things better than they actually do, and probability distributions tend to be far too narrow.*

The conversation covering the conditioning stage might resemble the following:

First, let us deal with the steady state cost, the cost that we get to X years after production starts. Before we get down to serious numbers, we would like to explore the range of possibilities. What is the highest steady state product cost you could imagine? What would need to happen for something like that to occur? Imagine we met in an airport ten years from now and that product cost was 20% higher than even the highest we had estimated; explain to me how this outcome might have happened. How about the lowest steady state cost you could imagine?

* Similar experiences are reported in E.C. Capen, "The Difficulty of Assessing Uncertainty," *Journal of Petroleum Technology* (August 1976): 843-850. This article is reprinted in *Readings on the Principles and Applications of Decision Analysis*, ed. R.A. Howard and J.E. Matheson, 2 vols. (Menlo Park, California: Strategic Decisions Group, 1984) 2: 591-600.

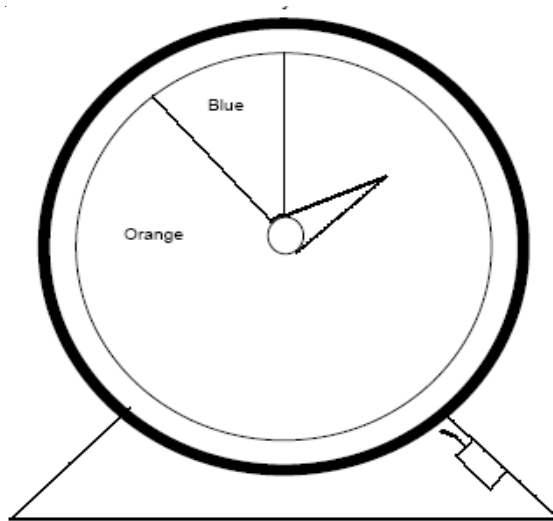
Stage 4: Encoding

Having defined the variable, structured it, and established and clarified the information useful for assessing its uncertainty in the first three stages, the encoder quantifies that uncertainty in this stage.

Of the various encoding methods available, we have found a reference method using a probability wheel (Figure 12-1) to be the most effective.

Figure 12-1


Probability Wheel



The probability wheel is divided into a blue and an orange sector, the relative sizes of which can be adjusted. To use the wheel, the encoder selects a value for the variable the expert thinks is not too extreme—but not the most likely or central value. The encoder then asks the expert, “Would you rather bet that the variable will be less than this value or that when I spin this wheel the pointer lands in the blue region?” The relative sizes of the blue and orange regions are then adjusted and the question repeated until a setting is found for which the subject is just indifferent. In other words, the encoder finds the point where the expert believes the probability of the two events (variable less than stated value or pointer landing in blue) are identical. The quick check of reversing the question—“Would you rather bet that the variable would be greater than this value or that the pointer lands in the orange?”—frequently makes the expert rethink the question and adjust the answer. A scale on the back of the wheel gives the probability of the event. This value is then plotted as one point on a cumulative distribution. Repeating this process for different values of the variable leads to a collection of points that may be connected by a smooth curve (Figure 12-2).

In using the probability wheel, the encoder should follow two important rules. First, the encoder must carefully avoid leading the expert to a value the facilitator thinks makes sense or is consistent. A wiser approach is to strive to confound the expert's possible attempts to mislead or impose false consistency in responses by, for example, varying the form of the questions and by skipping back and forth from high to low values so that the expert really has to think about each individual question.

Figure 12-2



Second, the encoder should plot and number the encoded points out of the expert's view and then look for inconsistencies and odd discontinuities. The encoder should note any changes in the expert's thinking that might be indicated by shifts in the points plotted. Often, he or she will see the curve along which the points lie shift because the subject has suddenly thought of some new piece of information. When this happens, the encoder should discuss the new thought and be prepared to discard all the earlier points if the perspective has been improved.

The conversation covering the encoding stage might resemble the following:

Now it is time to get some specific estimates. What value would you like to put down for a high estimate for steady state product cost. By "high" we mean a value such that there is one chance in ten that the value will be higher, and nine chances in ten that it will be lower. [At this point, it is useful to pull out a probability wheel or similar device to demonstrate what 10% means.] What value would you like to put down for a low estimate. "Low" means one chance in ten the value will be lower. [For a full encoding of a probability distribution, many more points would be assessed on the cumulative distribution, probably using the probability wheel.] Finally, how about a base case value? For "base

case" we like a value for which it is equally likely for the true value to turn out to be higher or lower.

Stage 5: Verification

The last stage of the encoding process is to test the judgments obtained in the encoding stage by explaining the encoded distribution to the subject. It is often useful to convert the cumulative distribution into histogram form and to discuss bimodal shapes or sharp extremes with the expert. To verify the distribution, the encoder can ask whether the expert would willingly invest his or her own money based on the encoded results. The encoder can also form equally likely intervals based on the distribution and see if the expert would have a difficult time choosing which interval to bet on.

The conversation covering the verification stage might resemble the following:

Here are the results we are going to use in the analysis. Does everything look right to you? If you were put in charge of implementing this decision, is there anything else in the product cost area you would be worried about?

Experiences and Insights from Practice _____

One obvious question is how long the encoding process normally takes. If the encoding session goes smoothly, the process may be completed in as little as half an hour, especially if only 10/50/90 ranges are being encoded. A complete encoding process for a critical uncertainty may require one to three hours. Incidentally, if resources permit, having a second facilitator present at the interview to record the data and to observe will expedite the interview and help in the subsequent review of results.

One point that new facilitators are often skeptical about is whether the probability wheel really works. Furthermore, they worry that the executive or expert will not take seriously questions based on the probability wheel and will not tolerate the exercise for the length of time involved. Our experience is that assuring experts that their superiors have approved and requested their participation in the encoding usually leads to a positive response. In virtually all cases, the expert readily accepts the wheel and adapts quickly to its use.

Probability encoding has gained considerable use by practicing decision facilitators. The process yields results, and experts and facilitators seem to like it.* As with quantitative analysis in general, the real value of probability encoding is determined not so much by whether it does or does not produce

* Encoded probabilities have been tracked in the business area. The reported results show that carefully encoded probabilities correspond quite well to the frequencies with which the outcomes actually occur. See the following articles: H.U. Balthasar, R.A.A. Boschi, and M.M. Menke, "Calling the Shots in R&D," *Harvard Business Review* (May–June 1978): 151–160; Irwin Kabus, "You Can Bank on Uncertainty," *Harvard Business Review* (May–June 1976): 95–105; and W.E. Sander, "The Validity of Subjective Forecasts by R&D Project Managers," *IEEE Transactions on Engineering Management* (February 1969): 35–43.

the correct bottom line as by the insights and improved clarity it provides decision-makers. And regardless of whether quantitative methods like probability encoding are perfectly accurate, decision-makers must continue to make decisions based on incomplete knowledge. Methods like probability encoding have proven too valuable to be dismissed lightly.

Summary

Obtaining high-quality probability distributions for crucial uncertainties is a vitally important step in every decision analysis. Although common motivational and cognitive biases tend to reduce the quality of a probability assessment, careful attention to the five steps of the assessment process can greatly improve the quality of the assessment.

Problems and Discussion Topics

- 12.1 Discuss the differences in the processes by which motivational and cognitive biases arise. What implications do these differences have for the methods to overcome them?
- 12.2 List a bias that commonly arises in an area other than probability assessment. How is the bias recognized and overcome? If it is not overcome, is it because it is not possible to do so or not important enough to do so? Or is it because it is not even recognized?
- 12.3 You are about to assess the probability distribution on the average growth rate over the next year for the entire energy industry. The expert is a market facilitator who closely follows the stocks of the large oil companies. What biases might you expect to encounter?
- 12.4 One technique for overcoming several kinds of biases is called the “Rip Van Winkle Technique.” To apply it, you would discuss with the subject the highest and lowest possible outcomes of an uncertain variable. You would then say that it is a number of years after the actual outcome of the variable was discovered. The two of you run into each other again. You inform him that the variable turned out to be 10 percent higher than his highest possible estimate years before. You ask him to explain how it turned out higher than either of you had thought possible.
Why does this technique work?
- 12.5 Find a friend to serve as a subject in the following subjective probability experiment. Alternatively, try the experiment on yourself.
 - a. Tell the subject that a fair coin [$p(\text{head} | S) = .5$; $p(\text{tail} | S) = .5$] will be flipped six times. Assess the subject’s cumulative probability distribution on the number of heads that occur in the six flips. Discourage your subject from trying to make any mathematical calculations of the odds. If you ask the ques-

tions in the right way, it will be very difficult for him or her to make any such calculations.

- b. Now tell your friend that you have three coins (two of which are unfair) with different probability distributions.

1. $p(\text{head} | S) = .25$, $p(\text{tail} | S) = .75$
2. $p(\text{head} | S) = .50$, $p(\text{tail} | S) = .50$
3. $p(\text{head} | S) = .75$, $p(\text{tail} | S) = .25$

Then tell the subject that one of those coins will be randomly selected and flipped six times. Assess his or her subjective cumulative distribution on the total number of heads that result.

- c. Calculate the actual distributions for a and b under the given assumptions. Compare these distributions with the assessments from your subject. Also note any difference between the subject's distributions in a and b. What might explain the differences, if any, between the various distributions?

- 12.6 Break into groups of two or three and encode probability distributions. Role playing by the “expert” and the “facilitator” can help make the exercise more realistic, especially if assumed motivational biases are written down beforehand (but not revealed to the facilitator). The quantity encoded should be a continuous variable for which the uncertainty will be resolved some time after the encoding session. Be sure to spend time describing and structuring the variable and exploring the possibility of biases. You may find it useful to structure a simple influence diagram with the subject before assessing the probability.

Some possible topics for assessment are (make sure the definitions pass the clairvoyance test):

- a. The price of a stock two weeks from now
- b. The difference in temperature between Stockholm and Rio de Janeiro on a particular day
- c. The number of people attending a large undergraduate class on a given day.

- 12.7 Slippery Company produces, among other things, special types of lubricants for specific mechanical applications. There is one type of lubricant it does not produce. This lubricant is currently produced by several large companies from a feedstock of ethylene. Since ethylene prices are rising along with petroleum prices, the cost to produce this lubricant is rising. (This case is a disguised version of an analysis done in the late 1970s.) Slippery knows that the lubricant can be made from the oil of the “oily bean” at a cost that appears competitive today with the ethylene-based process. Since oily bean oil prices are not rising, Slippery is considering constructing a facility to produce the lubricant

from oily bean oil. However, two factors worry Slippery. First, there is a rumor that several other companies are considering the same move, which would saturate the market with cheap lubricant. Second, although oily bean oil prices are fairly constant, droughts make the price jump temporarily every couple of years.

- a. Structure the problem and determine your information needs. Be sure to draw an influence diagram.
- b. For one of the necessary items of information, designate an expert, motivate the expert, and assess a probability distribution on the item.

Bibliography

The literature on decision analysis is extensive, and there are several quite different areas of investigation discussed in this literature. The following books should help the interested reader delve more deeply into these areas.

Allen, Michael S., *Business Portfolio Management: Valuation, Risk Assessment, and EVA Strategies*, New York: John Wiley & Sons, 2000.

Amran, Martha, and Nalin Kulatilaka, *Real Options: Managing Strategic Investment in an Uncertain World*, Boston: Harvard Business School Press, 1999.

Barabba, Vincent P., *Meeting of the Minds: Creating the Market-Based Enterprise*, Boston: Harvard Business School Press, 1995.

Behn, Robert D., and James W. Vaupel, *Quick Analysis for Busy Decision Makers*, New York: Basic Books, Inc., 1982.

Bodily, S., *Modern Decision Making: A Guide to Modeling with Decision Support Systems*, New York: McGraw-Hill Book Co., 1984.

Brown, Rex V., A. S. Kahn, and C. R. Peterson, *Decision Analysis: An Overview*, New York: Holt, Rinehart, and Winston, 1974.

Bunn, D. W., *Applied Decision Analysis*, New York: McGraw-Hill, 1984.

Clemen, Robert T., *Making Hard Decisions: An Introduction to Decision Analysis*, Duxbury Press, 1996.

Cotovello, Vincent T., and Miley W. Merkhofer, *Risk Assessment Methods: Approaches for Assessing Health and Environmental Risks*, Plenum, 1993.

Golub, Andrew L., *Decision Analysis: An Integrated Approach*, New York: John Wiley & Sons, 1996.

Hammond, John, Ralph L. Keeney, and Howard Raiffa, *Smart Choices: A Practical Guide to Making Better Decisions*, Boston: Harvard Business School Press, 1998.

Holloway, Charles A., *Decision Making Under Uncertainty: Models and Choices*, Englewood Cliffs, NJ: Prentice-Hall Inc., 1979.

Howard, Ronald A., and James E. Matheson, eds., *Readings on the Principles and Applications of Decision Analysis*, 2 volumes, Menlo Park, California: Strategic Decisions Group, 1984.

Keeney, Ralph L., and Howard Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, New York, NY: Cambridge University Press, 1993.

Keeney, Ralph L., *Value Focused Thinking: A Path to Creative Decisionmaking*, Boston: Harvard University Press, 1996.

Luenberger, David G., *Investment Science*, Oxford University Press, 1997.

Matheson, David, and James Matheson, *The Smart Organization: Creating Value Through Strategic R&D*, Boston: Harvard Business School Press, 1998.

Merkhofer, M. W., *Decision Science and Social Risk Management*, Boston: Reidel, 1987.

Newendorp, Paul D., *Decision Analysis for Petroleum Exploration*, Planning Press, 1998.

Raiffa, Howard, *Decision Analysis*, McGraw Hill College Division, 1997.

Raiffa, Howard, *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*, New York: Random House, Inc., 1986.

Samson, Danny, *Managerial Decision Analysis*, Homewood, IL: Richard D. Irwin, Inc., 1988.

Schlaifer, Robert, *Analysis of Decisions under Uncertainty*, Melbourne, FL: Robert E. Krieger Publishing Co., Inc., 1978.

Skinner, David C., *Introduction to Decision Analysis*, Probabilistic Press, 1999.

Spurr, William A., and Charles P. Bonini, *Statistical Analysis for Business Decisions*, Homewood, IL: Richard D. Irwin, Inc., 1973.

Thompson, Mark S., *Decision Analysis for Program Evaluation*, Cambridge, MA: Ballinger Publishing Co., 1982.

Tummala, V. M., and Richard C. Henshaw, eds., *Concepts and Applications of Modern Decision Models*, Michigan State University Press, 1976.

von Winterfeldt, Detlov, and Ward Edwards, *Decision Analysis and Behavioral Research*, New York, NY: Cambridge University Press, 1988.

Index

A

Adjustment and anchoring bias 325
Algebra of events 271
Analytic complexity 225
Analyzing the tree 26, 55
Availability bias 324
Aversion to risk taking 101

B

Basis appraisal 164
Basis development 143
Bayes' rule 75, 85, 280
Bidding strategies 192
Business portfolios 208
Buying price 133

C

Certain equivalent 30, 160
 Approximate formula 117, 135
Chance node 21
Choice rule 104
Clairvoyance test 19
Cognitive biases 323
Collectively exhaustive events 271
Communicating results 177
Conditional distributions 164
Conditional probabilities 277
Constant dollars 156
Corporate applications 187
 Bidding strategies 192
 Investment decisions 194
 Litigation 190
 New product introduction 187

Options 196
R&D 202
R&D portfolio 205
Real options 200
 Strategies and portfolios 208
Cumulative probability plot 26

D

Decision analysis 1
 Decision Quality 6
 Dialog Decision Process (DDP) 5
 Four levels of 2
 Good decisions and good outcomes 2
 Personal decision-making 8
 Scientific decision-making 8
 Societal decision-making 8
 Stanford University 8
Decision analysis capability building 245
Decision analysis cycle 141
 Basis appraisal 164
 Basis development 143
 Deterministic structuring 151
 Sensitivity analysis 157
 Probabilistic evaluation 161
Decision analysis process 42
Decision criterion 50, 233
Decision facilitator 3
Decision hierarchy 144
Decision node 44
Decision projects 224
Decision tree
 Constructing from influence diagram 45
Decision-dependent outcomes 93
Decision-dependent probabilities 92

Dependent outcomes 79
 Dependent probabilities 76
 Dialog Decision Process 142, 227
 Alternatives dialog 234
 Analysis dialog 236
 Decision dialog 241
 Framing dialog 229
 Discrete approximation 34
 Discretizing 34
 Distinctions 268
 Distribution trees 20

E

Encoded utility function 111
 Encoding probabilities 32
 Encoding probability distributions
 Adjustment and anchoring bias 325
 Availability bias 324
 Cognitive biases 323
 Encoding process 326
 Motivational biases 322
 Representativeness bias 324
 Equivalence rule 104
 Events 267, 283
 Algebra of events 271
 Distinctions 268
 Tree representation 274
 Venn diagram 268
 Evocative node 152
 Expected value 30
 Inadequacy of 101
 Exponential utility function 111

F

Facilitator 3

G

Gaussian distribution 117, 290
 Good decisions 41

H

Histogram plot 28

I

Imperfect control 92
 Inadequacy of expected values 101
 Inflation 156
 Influence diagrams 15, 147, 232, 303
 Decision representation 305
 Determined uncertainty 310
 Elements of 303
 Influence representation 306
 Procedures for manipulating 312
 Relevance 16
 Rules for constructing 45, 311
 Theory overview 303
 Three elements of 15
 Turning into a decision tree 313
 Uncertainty representation 303
 Value and determined value 310
 Information flow 45
 Investment decisions 194

J

Joint events 273
 Joint probabilities 277

L

Litigation decision analysis 190

M

Marginal probabilities 277
 Modeling profit 155
 Motivational biases 322
 Mutually exclusive events 271

N

Nature's tree 80–87
 Net present value (NPV) 53
 New product introduction 187
 No fun in gambling rule 105
 Normal distribution 290

O

Options 196
 Order rule 103

Organizational complexity 225

Outcomes

Decision-dependent 93

Dependent 79

P

Perfect control 62

Perfect information 58

Policy matrix 175

Presenting a decision analysis 244

Probabilistic dependence 75

Probabilistic evaluation 161

Probabilistic independence 75, 282

Probabilistic sensitivity analysis 56

Probabilities 17

Continuous variables 290

Cumulants 288

Cumulative probability 284

Decision-dependent probabilities 92

Definition of 18

Histogram 286

Mass density 284

Mean, median, mode, variance, standard deviation 286

Moments 288

Tabular form 284

Tree form 284

Probability distributions

Adjustment and anchoring bias 325

Availability bias 324

Cognitive biases 323

Encoding process 326

Motivational biases 322

Representativeness bias 324

Probability rule 105

Probability theory 267

Three axioms 276

Probability tree 23

Probability wheel 329

Problem modeling 151

Project vision statement 230

Pruning the tree 161

R

Range data 35

Real Options 200

Relevance 16, 76, 78

Reordering the tree

Interpretive problems 92–94

Representations of probabilities 284

Representativeness bias 324

Research and Development 202

Portfolio 205

Risk attitude 101, 160

Risk penalty 101

Risk premium 101

Risk-averse attitude 101

Risk-free discount rates 124

Roll-forward evaluation 30

Rollback evaluation 30

S

Schematic tree 24

Selling price 133

Sensitivity analysis 238

Deterministic 157

Probabilistic 157

Risk attitude 157

To probabilities 170

To risk attitude 170

Shutting down the business 156

Sources of value 237

Standard deviation 117, 160, 286

States of information 18, 275

Stochastic sensitivity analysis 56

Strategy table 145, 232

Rationale 235

Structured decision process 226

Substitution rule 104

Sunk costs 155

T

Taylor series expansion 117

Terminal value 156

U

- Uncertainty 13
- Utility curve 108
- Utility function 103
 - Definition of 106
 - Deriving existence of 120
 - Encoding of 119
 - Use of 108

V

- Value
 - Future money 53
 - Imperfect control 92
 - Imperfect information 87
 - Nonmonetary 52
 - Perfect control 62
 - Perfect information 58
- Value node 44
- Value of imperfect information 87
- Value of Information
 - With encoded utility function 111
- Value of Perfect Control 169
- Value of Perfect Information 166
- Value trade-offs 170
- Values 283
- Variables 283
- Variance 117, 160, 286
- Venn diagram 268

Authors

Peter McNamee

Peter McNamee obtained a Ph.D. in theoretical physics from Stanford University. After ten years of university teaching and physics research, he decided to leave the realm of academia for the practical world of decision analysis. He feels, however, that both vocations have similar purposes: to find the order in seeming chaos and to communicate the insights obtained in the process.

First with SRI International (formerly Stanford Research Institute) and then with Strategic Decisions Group, Peter has worked for companies in many areas, including bidding, energy, consumer goods, oil, telecommunications, chemicals, and insurance.

Peter has lectured on decision analysis both in the United States and abroad. One of his ongoing interests is education, which led to his co-authoring of this book. He feels computer software is essential to making decision analysis more practical. For this reason, he has played a principal role in developing several generations of Supertree software.

Peter is now one of the founders of SmartOrg, a company dedicated to developing systems to help organizations develop superior management decision skills through education and training, and through the development and implementation of advanced decision systems.

John Celona

John Celona has pursued his long-time interest in making difficult decisions and implementing the choice through a number of disciplines and vocations; first with a degree in industrial engineering and engineering management, then as a management consultant for Strategic Decisions Group, then as a lawyer specializing in litigation risk analysis, then as an independent consultant.

Then, things got really interesting as he tried to sit on the other side of the table by being interim CEO of two successive startup companies. This, of course, requires developing and implementing a business strategy for oneself. The first one never got off the ground. The second one is up and running. Meanwhile, John went back to consulting while considering what next.

SMARTORG, INC.

SMARTORG develops and markets advanced decision support systems based on the processes described in *Decision Analysis for the Professional*. Our products include Decision Advisor®, an intelligent decision system for building decision models; Portfolio Navigator™, a web-based decision support management system; and Supertree® decision analysis software.

SmartOrg was founded in 2000 by former principals and managers of Strategic Decisions Group (SDG). Our goal is to help management make quicker and better decisions that create exceptional value through technology. SmartOrg software and systems reflect the strategic decision-making processes and tools developed and proven at SDG and Stanford University's Department of Management Sciences and Engineering over several decades.

For information about our services and other publications please visit www.smartorg.com.

SMARTORGSM

Decisions Shape the Future

www.smartorg.com

ISBN0-9710569-0-0