


Robert T. Clemen

MAKING HARD DECISIONS

An Introduction to
Decision Analysis

SECOND EDITION



MAKING HARD DECISIONS

An Introduction to Decision Analysis

SECOND EDITION

Robert T. Clemen, Duke University

Here's a straightforward book that teaches the fundamental ideas of decision analysis—without an overly technical explanation of the mathematics used in management science.

This new Second Edition of *Making Hard Decisions*, by Robert T. Clemen, continues to provide readers with a thoroughly enjoyable, easy-to-understand text. The book emphasizes the practical application of decision analysis through interesting case studies and useful exercises that represent the broad range of fields and industries that benefit from systematic decision analysis.

Making Hard Decisions also contains up-to-date coverage of modern computer applications. Special attention to the role of the computer in decision analysis, as well as descriptions of the many software packages available—like DPL, Excel and Data—make this book a handy resource.

NEW TO THE SECOND EDITION

- Value-focused thinking integrated into the text
- Updated coverage of decision analysis software and discussion of dedicated simulation software
- New decision-analytic approach to regression

ABOUT THE AUTHOR

Robert T. Clemen is a Professor at the Fuqua School of Business at Duke University. An active contributor to the field of decision analysis since earning his Ph.D. at Indiana University in 1984, he has become an authority in the use of expert information in risk analysis. Aside from his research interests, Bob has a strong focus on making decision analysis accessible to students at all levels, ranging from creating teaching materials for high school students to teaching in executive education programs.

Making Hard Decisions

**An Introduction to Decision Analysis
2nd Edition**

Robert T. Clemen
Fuqua School of Business
Duke University

Duxbury Press

An Imprint of Brooks/Cole Publishing Company

I (T) P® An International Thomson Publishing Company

Pacific Grove • Albany • Belmont • Bonn • Boston • Cincinnati • Detroit
Johannesburg • London • Madrid • Melbourne • Mexico City
New York • Paris • Singapore • Tokyo • Toronto • Washington

Editor: Curt Hinrichs
Assistant Editor: Julie McDonald
Editorial Assistant: Cynthia Mazow
Marketing Manager: Joanne Terhaar
Advertising Project Manager: Joseph Jodar
Production: Mary Douglas/Rogue Valley Publications
Print Buyer: Karen Hunt

Permissions Editor: Peggy Meehan
Copy Editor: Alice Goehring
Text Design: Wendy LaChance
Cover Design: Ross Carron, Carron Design
Cover Illustration: Rik Olson
Illustrations and Composition: Interactive Composition Corporation
Printer: Quebecor/Fairfield

COPYRIGHT © 1996 by Brooks/Cole Publishing Company A Division of the International Thomson Publishing Inc. I (T) P* The ITP logo is a registered trademark under license. Duxbury Press and the leaf logo are trademarks used under license.

Printed on
acid-free
recycled paper.

Crystal Ball ® screen shots © 1993; reprinted with permission from Decisioneering. DATA ™ screen shots © 1994-1995; reprinted with permission from TreeAge Software, Inc. Microsoft ® Excel screen shots © 1985-1995 by Microsoft Corporation; reprinted with permission from Microsoft corporation. DPL ™ screen shots © 1992-1995 by ADA Decision Systems; reprinted with permission. @RISK and BestFit ® screen shots © 1990-1995 by Palisade Corporation; reprinted with permission. Logical Decisions ® screen shots © 1993-1995 by Logical Decisions; reprinted with permission.

Crystal Ball is a registered trademark of Decisioneering. DATA and Decision Analysis by TreeAge are trademarks of TreeAge Software, Inc. Microsoft is a registered trademark and Windows is a trademark of Microsoft Corporation. DPL and Decision Programming Language are trademarks of ADA Decision Systems. RISK is a trademark of Parker Brothers, Division of Tonka corporation and is used by Palisade under license. BestFit and Palisade are registered trademarks of Palisade Corporation. Logical Decisions is a registered trademark of Logical Decisions. Apple and Macintosh are registered trademarks of Apple Computer, Inc.

The following cases and examples are fictional. Any resemblance to actual people or organizations is purely coincidental: Larkin Oil (p. 31); The Value of Patience (p. 36); Early Bird, Inc. (p. 37); Hiring a Summer Intern (p. 42); The Summer Job (p. 128); GPC'S New Product Design (p. 145); Southern Electronics, Part I (p. 146); Southern Electronics, Part II (p. 147); Strenlar (p. 148); Job Offers (p. 150); Eagle Airlines (p. 155); Dumond International, Part I (p. 182); Strenlar, Part II (p. 183); Job Offers, Part II (p. 184); Decision Analysis Monthly (p. 256); Overbooking (p. 345); Municipal Solid Waste (p. 350); Taco Shells (p. 405); Forecasting Sales (p. 406); Overbooking, Part II (p. 408); Fashions (p. 411); Choosing a Manufacturing Process (p. 428); Organic Farming (p. 430); Overbooking, Part III (p. 433); Dumond International, Part II (p. 456); Interplants, Inc. (p. 497); Strenlar, Part III (p. 501); The Life Insurance Game (p. 524); Nuclear Power Paranoia (p. 525); The Manager's Perspective (p. 526); A Matter of Ethics (p. 572). E. H. Harriman Fights for the Northern Pacific Railroad (p. 462) is a fictionalized account of historical events.

Printed in the United States of America
6 7 8 9 10

For more information, contact Duxbury Press at Brooks/Cole Publishing Company:

Brooks/Cole Publishing Company
511 Forest Lodge Road
Pacific Grove, CA 93950

International Thomson Editores
Campos Eliseos 385, Piso 7
Col. Polanco
11560 Mexico D.F. Mexico

International Thomson Publishing Europe
Berkshire House 168-173
High Holborn
London, WC1V 7 AA, England

International Thomson Publishing GmbH
Konigswinterer Strasse 418 53227 Bonn,
Germany

Thomas Nelson Australia
102 Dodds Street South
Melbourne 3205 Victoria,
Australia

International Thomson Publishing Asia
221 Henderson Road #05-10 Henderson
Building Singapore 0315

Nelson Canada 1120
Birchmount Road
Scarborough, Ontario
Canada M1K 5G4

International Thomson Publishing Japan
Hirakawacho Kyowa Building, 3F 2-2-1
Hirakawacho Chiyoda-ku, Tokyo 102,
Japan

All rights reserved. No part of this work covered by the copyright hereon may be reproduced or used in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems—without the written permission of the publisher.

Library of Congress Cataloging-in-Publication Data

Clemen, Robert T. (Robert Taylor, 1952-
Making hard decisions : an introduction to decision analysis /
Robert T. Clemen. — 2nd ed.
p. cm.
Includes bibliographical references and index.
ISBN 0-534-26034-9 1. Decision-making. I.
Title.

HD30.23.C577 1995
658.4'03-dc20

95-42073
CIP

Brief Contents

1	Introduction to Decision Analysis	1
Section 1	Modeling Decisions	17
2	Elements of Decision Problems	19
3	Structuring Decisions	41
4	Making Choices	101
5	Sensitivity Analysis	155
6	Creativity and Decision Making	187
Section 2	Modeling Uncertainty	217
7	Probability Basics	219
8	Subjective Probability	265
9	Theoretical Probability Models	314
10	Using Data	355
11	Monte Carlo Simulation	410
12	Value of Information	435
Section 3	Modeling Preferences	459
13	Risk Attitudes	461
14	Utility Axioms, Paradoxes, and Implications	503
15	Conflicting Objectives I: Fundamental Objectives and the Additive Utility Function	530
16	Conflicting Objectives II: Multiattribute Utility Models with Interactions	576
17	Conclusion and Further Reading	607

Contents

Preface xix

Chapter 1 Introduction to Decision Analysis 1

Gypsy Moths and the ODA 1

Why Are Decisions Hard? 2

Why Study Decision Analysis? 3

Subjective Judgments and Decision Making 5

The Decision-Analysis Process 5

Requisite Decision Models 8

Where Is Decision Analysis Used? 8

Where Are We Going from Here? 9

Summary 10

Questions and Problems 10

Case Studies: Dr. Joycelyn Elders and the War on Drugs 11

Lloyd Bentsen for Vice President? 12 Du Pont

and Chlorofluorocarbons 13

References 14

Epilogue 15

Section 1 Modeling Decisions 17 Chapter

2 Elements of Decision Problems 19

Values and Objectives	19
Special Objective	20
Decision Context	21
Boeing's Supercomputer	22
Decisions to Make	23
Sequential Decisions	24
Uncertain Events	25
Consequences	27
The Time Value of Money: A Special Kind of Trade-Off	28
Larkin Oil	31
Summary	34
Questions and Problems	34
<i>Case Studies:</i>	
The Value of Patience	36
Early Bird, Inc.	37
References	39
Epilogue	39

Chapter 3 Structuring Decisions 41

Structuring Values	42
Hiring a Summer Intern	42
Fundamental and Means Objectives	44
Getting the Decision Context Right	49
Structuring Decisions: Influence Diagrams	50
Influence Diagrams and the Fundamental Objectives Hierarchy	52
Using Arcs to Represent Relationships	53
Some Basic Influence Diagrams	55
<i>The Basic Risky Decision</i>	55
<i>Imperfect Information</i>	56
<i>Sequential Decisions</i>	59
<i>Intermediate Calculations</i>	61
Constructing an Influence Diagram (Optional)	63
Toxic Chemicals and the EPA	63
<i>Some Common Mistakes</i>	65
<i>Multiple Representations and Requisite Models</i>	67

Structuring Decisions: Decision Trees	67
Decision Trees and the Objectives Hierarchy	69
Some Basic Decision Trees	70
<i>The Basic Risky Decision</i>	70
<i>Imperfect Information</i>	72
<i>Sequential Decisions</i>	72
Decision Trees and Influence Diagrams Compared	74
Decision Details: Defining Elements of the Decision	74 More
Decision Details: Cash Flows and Probabilities	76 Defining
Measurement Scales for Fundamental Objectives	77 Using
Computers for Structuring Decisions	82
<i>Logical Decisions</i>	83
<i>DPL (Decision Programming Language)</i>	85
<i>DATA</i>	88
Summary	89
Exercises	90
Questions and Problems	92 <i>Case</i>
<i>Studies:</i> Cold Fusion	96
Prescribed Fire	97 The
<i>SS Kuniang</i>	97
References	98
Epilogue	100

Chapter 4 Making Choices 101

Texaco Versus Pennzoil	101
Decision Trees and Expected Monetary Value	105
Solving Influence Diagrams: Overview	109
Solving Influence Diagrams: The Details (Optional)	112
Solving Influence Diagrams: An Algorithm (Optional)	117
Risk Profiles	118
<i>Cumulative Risk Profiles</i>	122
Dominance: An Alternative to EMV	123
Making Decisions with Multiple Objectives	127 The Summer Job
Analysis: One Objective at a Time	130 Subjective Ratings for
Constructed Attribute Scales	130 Assessing Trade-Off Weights
Analysis: Expected Values and Risk Profiles for Two Objectives	132
	133

Computer Programs for Decision Analysis 136
 DPL 136 *DATA* 137 *Electronic Spreadsheets* 140
Summary 140
Exercises 141
Questions and Problems 142
Case Studies: GPC's New Product Decision 145
 Southern Electronics, Part I 146
 Southern Electronics, Part II 147
 Strenlar 148 Job Offers 150 SS
 Kuniang, Part II 152
References 152
Epilogue 154

Chapter 5 Sensitivity Analysis 155

Eagle Airlines 155
Sensitivity Analysis: A Modeling Approach 156
Problem Identification and Structure 157
One-Way Sensitivity Analysis 160
Tornado Diagrams 161
Dominance Considerations 162
Two-Way Sensitivity Analysis 164
Sensitivity to Probabilities 165
Two-Way Sensitivity Analysis for Three Alternatives
 (Optional) 169
Investing in the Stock Market 169
Sensitivity Analysis in Action 172
Heart Disease in Infants 173 Sensitivity
Analysis by Computer 174
 Sensitivity Analysis with DPL 174 *Sensitivity Analysis*
 with DATA 176
Sensitivity Analysis: A Built-in Irony 178
Summary 178 Exercises 178 Questions
and Problems 179

<i>Case Studies: The Stars and Stripes</i>	181
Dumond International, Part I	182 Strenlar,
Part II	183 Facilities Investment and
Expansion	184 Job Offers, Part II
184	
References	186
Epilogue	186

Chapter 6 Creativity and Decision Making 187

What Is Creativity?	188 Theories
of Creativity	189 Chains of
Thought	189 Phases of the
Creative Process	190 Blocks to
Creativity	192
<i>Framing and Perceptual Blocks</i>	192
The Monk and the Mountain	192
Making Cigars	193
<i>Value-Based Blocks</i>	195
<i>Cultural and Environmental Blocks</i>	197
Ping-Pong Ball in a Pipe	197
<i>Organizational Issues</i>	199 Value-Focused Thinking for
Creating Alternatives	200
<i>Fundamental Objectives</i>	200
<i>Means Objectives</i>	200
Transportation of Nuclear Waste	201
<i>The Decision Context</i>	202 Other
Creativity Techniques	203
<i>Fluent and Flexible Thinking</i>	203
<i>Idea Checklists</i>	203
<i>Brainstorming</i>	206
<i>Metaphorical Thinking</i>	206
<i>Other Techniques</i>	208 Creating
Decision Opportunities	209 Summary
209 Questions and Problems	210 Case
<i>Studies: Modular Olympics</i>	211
Burning Grass-Seed Fields	212
References	212
Epilogue	214

Section 2 Modeling Uncertainty 217

Chapter 7 Probability Basics 219

- A Little Probability Theory 220
- Venn Diagrams 220
- More Probability Formulas 221
- Uncertain Quantities 226
 - Discrete Probability Distributions* 227
 - Variance and Standard Deviation* 230
 - Measuring Dependence (Optional)* 232
 - Distributions* 236
 - Stochastic Dominance Revisited* 237
 - Stochastic Dominance and Multiple Attributes (Optional)* 238
 - Probabilities Density Functions* 239
 - Expected Value, Variance, and Standard Deviation: The Continuous Case* 240
 - Covariance and Correlation: The Continuous Case (Optional)* 241
- Oil Wildcatting 242
- John Hinckley's Trial 248
- Decision-Analysis Software and Bayes' Theorem 250
- Summary 250
- Exercises 251
- Questions and Problems 255
- Case Studies:*
 - Decision Analysis Monthly 257
 - Screening for Colorectal Cancer 258
 - AIDS 259
- Discrimination and the Death Penalty 262
- References 264
- Epilogue 264

Chapter 8 Subjective Probability 265

- Uncertainty and Public Policy 265
- Probability: A Subjective Interpretation 267
- Accounting for Contingent Losses 268
- Assessing Discrete Probabilities 269
- Assessing Continuous Probabilities 274
- Pitfalls: Heuristics and Biases 281

TomW.	281
<i>Representativeness</i>	282
<i>Availability</i>	283
<i>Anchoring and Adjusting</i>	284
<i>Motivational Bias</i>	284
<i>Heuristics and Biases: Implications</i>	284
Decomposition and Probability Assessment	285
Experts and Probability Assessment: Pulling It All Together	291
Climate Change at Yucca Mountain, Nevada	294
Coherence and the Dutch Book (Optional)	296
Summary	298
Exercises	299
Questions and Problems	300
<i>Case Studies:</i> Assessing Cancer Risk—From Mouse to Man	305
Breast Implants	307
The Space Shuttle <i>Challenger</i>	308
References	310
Epilogue	313

Chapter 9 Theoretical Probability Models 314

Theoretical Models Applied	315
The Binomial Distribution	316
The Poisson Distribution	320
The Exponential Distribution	323
The Normal Distribution	325
The Beta Distribution	331
Probability Distributions and Decision-Analysis Software	335
Summary	335
Exercises	335
Questions and Problems	337
<i>Case Studies:</i> Overbooking	345
Earthquake Prediction	346
Municipal Solid Waste	350
References	353
Epilogue	354

Chapter 10 Using Data 355

Using Data to Construct Probability Distributions	355
<i>Histograms</i>	356
<i>Empirical CDFs</i>	357

Halfway Houses 357

Using Data to Fit Theoretical Probability Models 361

Software for Fitting Distributions: BestFit 362

Using Data to Model Relationships 363

The Regression Approach 365

Estimation: The Basics 368 *Estimation: More than One Conditioning Variable* 375 *Regression Analysis and Modeling: Some Do's and Don't's* 378

Regression Analysis: Some Bells and Whistles 383 *Regression Modeling: Decision Analysis Versus Statistical Inference* 386 *An Admonition: Use with Care* 387

Natural Conjugate Distributions (Optional) 387

Uncertainty About Parameters and Bayesian Updating 388 *Binomial Distributions: Natural Conjugate Priors for p* 390 *Normal Distribution: Natural Conjugate Priors for μ* 392 *Predictive Distributions* 394 *Predictive Distributions: The Normal Case* 395 *Predictive Distributions: The Binomial Case* 396

A Bayesian Approach to Regression Analysis (Optional) 397

Summary 397

Exercises 398

Questions and Problems 398

Case Studies: Taco Shells 405

Forecasting Sales 406

Overbooking, Part II 408

References 408

Chapter 11 Monte Carlo Simulation 410

Fashions 411

Using Uniform Random Numbers as Building Blocks 414

General Uniform Distributions 415

Exponential Distributions 416

Discrete Distributions 417

Other Distributions 417

Computer Software for Simulation 417

Crystal Ball 418 *@ RISK* 420

Simulation and Sensitivity Analysis 422

Distributions on Parameters (Optional) 423

Simulation, Decision Trees, and Influence Diagrams	424
Summary	426
Exercises	426
Questions and Problems	426
<i>Case Studies:</i>	
Choosing a Manufacturing Process	428
Organic Farming	430
Overbooking, Part III	433
References	433

Chapter 12 Value of Information 435

Investing in the Stock Market	435
Value of Information: Some Basic Ideas	436
<i>Probability and Perfect Information</i>	436
<i>The Expected Value of Information</i>	438
Expected Value of Perfect Information	439
Expected Value of Imperfect Information	441
Value of Information in Complex Problems	447
Value of Information, Sensitivity Analysis, and Structuring	448
Seeding Hurricanes	448
Value of Information and Nonmonetary Objectives	449
Value of Information and Experts	450
Calculating EVPI: DATA and DPL	451
Summary	451
Exercises	452
Questions and Problems	453
<i>Case Studies:</i>	
Texaco-Pennzoil Revisited	455
Medical Tests	456
Dumond International, Part II	456
References	457

Section 3 Modeling Preferences 459

Chapter 13 Risk Attitudes 461

E. H. Harriman Fights for the Northern Pacific Railroad	462
Risk	463

Risk Attitudes 465
 Investing in the Stock Market, Revisited 467
 Expected Utility, Certainty Equivalents, and Risk Premiums 469
 Keeping Terms Straight 473
 Utility Function Assessment 473
 Assessment Using Certainty Equivalents 474 *Assessment Using Probabilities* 475 *Gambles, Lotteries, and Investments* 477
 Risk Tolerance and the Exponential Utility Function 477
 Risk Tolerance and Sensitivity Analysis: Eagle Airlines, Revisited 480
 Decreasing and Constant Risk Aversion (Optional) 482
 Decreasing Risk Aversion 483 *An Investment Example* 484 *Constant Risk Aversion* 485
 Some Caveats 488
 Summary 488 Exercises 489
 Questions and Problems 490
Case Studies: Interplants, Inc. 497
 Texaco-Pennzoil One More Time 499
 Strenlar, Part III 501
 References 501
 Epilogue 502

Chapter 14 Utility Axioms, Paradoxes, and Implications 503

Preparing for an Influenza Outbreak 503
 Axioms for Expected Utility 504
 Paradoxes 510
 Implications 514
 Implications for Utility Assessment 514 *Managerial and Policy Implications* 516
 A Final Perspective 518
 Summary 519
 Exercises 520
 Questions and Problems 521
Case Studies: The Life Insurance Game 524
 Nuclear Power Paranoia 525
 The Manager's Perspective 526

References 527
Epilogue 529

Chapter 15 Conflicting Objectives I: Fundamental Objectives and the Additive Utility Function 530

Objectives and Attributes 532
Trading Off Conflicting Objectives: The Basics 534
 Choosing an Automobile: An Example 534 The
Additive Utility Function 536
 Choosing an Automobile: Proportional Scores 538 *Assessing Weights:
 Pricing Out the Objectives* 539 *Indifference Curves* 540
Assessing Individual Utility Functions 542
 Proportional Scores 542 *Ratios* 544 *Standard Utility-Function
 Assessment* 546
Assessing Weights 546
 Pricing Out 547 *Swing Weighting* 547 *Lottery Weights* 550
Keeping Concepts Straight: Certainty Versus Uncertainty 552 An
Example: Library Choices 553 The Eugene Public Library 553
Using Software for Multiple-Objective Decisions 560 Summary
560 Exercises 561 Questions and Problems 562 *Case Studies:*
The Satanic Verses 569
 Dilemmas in Medicine 570
 A Matter of Ethics 572
 FDA and the Testing of Experimental Drugs 573
References 574
Epilogue 575

Chapter 16 Conflicting Objectives II: Multiattribute Utility Models with Interactions 576

Multiattribute Utility Functions: Direct Assessment 577
Independence Conditions 579
 Preferential Independence 579 *Utility Independence* 580
Determining Whether Independence Exists 580

Using Independence 582
Additive Independence 583
Substitutes and Complements 586
Assessing A Two-Attribute Utility Function 586
The Blood Bank 587
Three or More Attributes (Optional) 591
When Independence Fails 592
Multiattribute Utility in Action: BC Hydro 593
Strategic Decisions at BC Hydro 593
Summary 598
Exercises 599
Questions and Problems 600
Case Studies: A Mining-Investment Decision 603
References 605
Epilogue 606

Chapter 17 Conclusion and Further Reading 607

A Decision-Analysis Reading List 608

Appendixes 613

A Binomial Distribution: Individual Probabilities 614
B Binomial Distribution: Cumulative Probabilities 622
C Poisson Distribution: Individual Probabilities 630
D Poisson Distribution: Cumulative Probabilities 635
E Normal Distribution: Cumulative Probabilities 640
F Beta Distribution: Cumulative Probabilities 644

Answers to Selected Exercises 653

Author Index 654 Subject Index

657

Preface

This book provides a one-semester overview of decision analysis for advanced undergraduate and master's degree students. My inspiration to write it has come from many sources, but perhaps most important was a desire to give students access to up-to-date information on modern decision-analysis techniques at a level that could be easily understood by those without a strong mathematical background. At some points in the book, the student should be familiar with basic statistical concepts normally covered in an undergraduate applied statistics course. In particular, some familiarity with probability and probability distributions would be helpful in Chapters 7 through 12. Chapter 10 provides a decision-analysis view of data analysis, including regression, and familiarity with such statistical procedures would be an advantage when covering this topic. Algebra is used liberally throughout the book. Calculus concepts are used in a few instances as an explanatory tool. Be assured, however, that the material can be thoroughly understood, and the problems can be worked, without any knowledge of calculus.

I subscribe to the notion that the objective of decision analysis is to help a decision maker think hard about the specific problem at hand, including the overall structure of the problem as well as his or her preferences and beliefs. Decision analysis provides both an overall paradigm and a set of tools with which a decision maker can construct and analyze a model of a decision situation. Above all else, I want students to understand that the purpose of studying decision-analysis techniques is to be able to represent real-world problems using models that can be analyzed to gain insight and understanding. It is through that insight and understanding — the hoped-for result of the modeling process — that decisions can be improved.

New in the Second Edition

The past five years have seen changes in decision analysis, and I have tried to capture the most important advancements in this second edition of *Making Hard Decisions*. The greatest change is the inclusion of material from Ralph Keeney's *Value-Focused Thinking* (Cambridge, MA: Harvard University Press, 1992). Keeney stresses the primacy of the decision-maker's values in the decision-making process. Values provide the very motivation for decision making in the first place. Although it sounds trite and oversimplified, "When in doubt, think about what your values are!" is very good advice for anyone who faces a difficult decision or who has the vague uneasy feeling that things could be better.

Because values play a vital role in so many aspects of decision making, incorporating value-focused thinking required rethinking many parts of the book. The most radical restructuring is in the first section. Readers will find important new sections in Chapter 3 on structuring objectives, modeling multiple objectives using decision trees and influence diagrams, and developing measurement scales. A simple example demonstrating how objectives can be modeled with an additive value function is included in Chapter 4. An important element of value-focused thinking is that it can provide guidance in the creation of new alternatives, and this material is incorporated in Chapter 6. Finally, the incorporation of value-focused thinking early in the book had implications for the presentation of multiattribute utility in Chapters 15 and 16.

In addition to the incorporation of value-focused thinking, references to software have been updated throughout the book. Although software changes so rapidly that the material in the book may be dated by the time it appears in print, I have included up-to-date material on a few of the most useful decision-analysis programs available for personal computers: DPL for both influence diagrams and decision trees, DATA for decision trees, Logical Decisions for modeling multiple objectives, BestFit for using data, @RISK and Crystal Ball for Monte Carlo simulation.

The important changes and new features in the second edition of *Making Hard Decisions* are

- Chapters 3 and 4 have been completely rewritten to incorporate value-focused thinking, including discussions of principles, techniques, and examples.
- Chapter 6 on creativity has been refocused to adopt a social-psychology perspective. Also, it incorporates creativity techniques from value-focused thinking.
- New sections on the use of expert judgment are included in Chapter 8.
- New material on regression analysis has been added to Chapter 10.
- Chapter 11 has been rewritten to focus on the use of simulation software (@RISK and Crystal Ball).
- Chapter 15 has been restructured to highlight the additive utility model.
- The BC Hydro example included in Chapter 16 demonstrates the use of multiattribute utility models.
- Updated references to software are included throughout.

Influence-diagram graphics have been changed throughout to conform with DPL software.

Guidelines for Students

The book covers most of the concepts that I consider important for a basic understanding of decision analysis. Although I have tried to write an elementary introduction to decision analysis, this does not mean that the material is itself elementary. In fact, the more I teach decision analysis, the more I realize that the technical level of the math is low, while the level of the analysis is high. Students must be willing to think clearly and analytically about the problems and issues that arise in decision situations. Good decision analysis requires clear thinking; sloppy thinking results in worthless analysis.

Of course, some topics are more demanding than others. The more difficult sections are labeled as "optional." My faith in students and readers compels me to say that anyone who can handle the "nonoptional" material can, with a bit more effort and thought, also handle the optional material. Thus the label is perhaps best thought of as a warning regarding the upcoming topic. On the other hand, if you do decide to skip the optional material, no harm will be done.

In general, I believe that really serious learning happens when problems are tackled on one's own. I have included a wide variety of exercises, questions, problems, and case studies. The exercises are relatively easy drills of the material. The questions and problems often require thinking beyond the material in the text. Some concepts are presented and dealt with only in the problems. Do not shy away from the problems! You can learn a lot by working through them.

Many case studies are included in *Making Hard Decisions*. A few of the many successful applications of decision analysis show up as case studies in the book. In addition, many issues are explored in the case studies in the context of current events. For example, the AIDS case at the end of Chapter 7 demonstrates how probability techniques can be used to interpret the results of medical tests. In addition to the real-world cases, the book contains many hypothetical cases and examples, as well as fictional historical accounts, which I have tried to make as realistic as possible.

Some cases and problems are realistic in the sense that not every bit of information is given. In these cases, appropriate assumptions are required. On one hand, this may cause some frustration. On the other hand, incomplete information is typical in the real world. Being able to work with problems that are "messy" in this way is an important skill.

Finally, many of the cases and problems involve controversial issues. For example, the material on AIDS (Chapter 7) or medical ethics (Chapter 15) may evoke strong emotional responses from some readers. In writing a book like this, there are two choices. We can avoid the hard social problems that might offend some readers. Or we can face these problems that need careful thought and discussion. I have taken the second approach because I believe these issues require our attention. Moreover,

even though decision analysis does not provide the answers to these problems, it does provide a useful framework for thinking about the difficult decisions that we as a society must make.

Computers and Decision Analysis

Computers play a large role in decision analysis. As a result, sections in several chapters discuss current software available for the construction and analysis of decision models on personal computers.

One of the most intriguing software developments in decision analysis is the publication of DPL. This remarkable product integrates both influence diagrams and decision trees. Created initially for consulting use by Applied Decision Analysis, Inc., DPL has also enjoyed considerable success as a vehicle for teaching decision analysis. To facilitate the use of DPL in courses, a student edition of DPL is available at a discount when bundled with *Making Hard Decisions*.

In addition to dedicated decision-analysis software, much of the discussion in the text revolves around the use of electronic spreadsheets, such as Lotus 1-2-3 or Microsoft Excel. Spreadsheets provide a very flexible modeling environment for analyzing decision situations, and many students and decision makers may have access to spreadsheets rather than special-purpose software. One of the most useful skills that an analytically-minded decision maker could have these days is expertise in the use of a spreadsheet.

Do not take this business about computers wrong, though. Decision-analysis concepts and techniques can indeed be learned, understood, and used quite effectively without ever touching a computer. In some cases, though, such as simulation or sensitivity analysis, working with a computer more readily demonstrates the insights that can be obtained from the analysis.

A Word to Instructors

Many instructors will want to supplement *Making Hard Decisions* with their own material. In fact, a few topics that I cover in my own courses are not included here. But, in the process of writing the book and obtaining comments from colleagues, it has become apparent that decision-making courses take on many different forms. Some instructors prefer to emphasize behavioral aspects, while others prefer analytical tools. Other dimensions have to do with competition, negotiation, and group decision making. *Making Hard Decisions* does not aim to cover everything for everyone. Instead, I have tried to cover the central concepts and tools of modern decision analysis with adequate references (and occasionally cases or problems) so that instructors can introduce their own special material where appropriate. For example, in Chapters 8 and 14 we discuss judgmental aspects of probability assessment and decision making, and an instructor can introduce more behavioral material at these points. Likewise, Chapter 15 delves into the additive utility function for decision making. Some instructors may wish to present goal programming or the analytic hierarchy process here.

In restructuring the book to incorporate value-focused thinking, it became clear that my colleagues were of several minds as to the best way to organize and present this topic. Some preferred to have all of the multiple-objective material put in the same place (Chapters 15 and 16), whereas others preferred to integrate the material throughout the text. Ultimately (after using value-focused thinking to help me decide), I chose the latter, especially stressing the role of values in structuring decision models. In particular, students must read about structuring values at the beginning of Chapter 3 before going on to structuring influence diagrams or decision trees. The reason is simply that it makes sense to understand what one wants before trying to structure the decision. From my own teaching experience, introducing the notion of objectives early works well for students. In Gordon Hazen's words, "Once I ask the students about their objectives in a particular problem, I can't get them to shut up!"

At the same time, I attempted to organize the remaining material in Chapters 3 and 4 so that the instructor who wants to avoid the more technical aspects of value-focused thinking may do so. For example, material on creating attribute scales is left until the end of Chapter 3, and a multiobjective example is in its own section in Chapter 4. A reader can readily skip these sections and return to them later if the instructor wishes to cover all of the multiobjective material at once.

Keeping up with Changes

The world changes quickly, and decision analysis is changing with it. The good news is that the Internet, and especially the World Wide Web (WWW), can help us keep abreast of new developments.

For students and instructors, I encourage you to visit the WWW site of the Decision Analysis Section of the Institute For Operations Research and Management Science (INFORMS). The current address is <http://www.fuqua.duke.edu/faculty/daweb/>. This organization provides a focus for decision analysts worldwide and many others with interests in all aspects of decision making. And on the Society's web page you will find links to many related sites.

While you are keeping up with changes, I hope that you will help me do the same. I encourage you to provide feedback on the second edition of *Making Hard Decisions* by sending email to Clemen@mail.duke.edu. You may also send regular mail to Robert Clemen, Fuqua School of Business, Duke University, Durham, NC 27708. Please let me know of (hopefully the few) mistakes or typos that you may find in the book, innovative ways to teach decision analysis, new case studies, or interesting applications of decision analysis.

Acknowledgments

It is a pleasure to acknowledge the help I have had with the preparation of this text. First mention goes to my students, who craved the notes from which the first edition grew. For resource support, thanks to the Lundquist College of Business at the University of Oregon, Decision Research of Eugene, Oregon, Applied Decision Analysis, Inc., of Menlo Park, California, and the Fuqua School of Business of Duke University.

Oddly enough, I owe a special debt of gratitude to the authors of competing decision-analysis texts. Their treatments of various topics typically inspired me to try to find a way to bring the insights I gained from their work to my students. Ron Howard's enthusiasm for influence diagrams to some extent influenced my decision to incorporate them into the first edition. Adam Borison's efforts likewise were instrumental in getting DPL incorporated into the second edition. Also, I have been influenced by Detlof von Winterfeldt and Ward Edwards's advanced text *Decision Analysis and Behavioral Research*, Derek Bunn's *Applied Decision Analysis*, and, of course, Ralph Keeney's *Value-Focused Thinking*.

A number of individuals have provided comments on portions of the book at various stages. The contributions of the following reviewers were appreciated: Gordon Hazen, Northwestern University; Donald Keefer, Arizona State University; Don N. Kleinmutz, University of Illinois; and Irving H. LaValle, Tulane University. Thanks to Deborah Amaral, Sam Bodily, Cathy Barnes, George Benson, Dave Braden, Bill Burns, Peter Farquhar, Ken Gaver, Andy Golub, Max Henrion, Ralph Keeney, Robin Keller, Craig Kirkwood, George MacKenzie, Allan Murphy, Bob Nau, Roger Pfaffenberger, Steve Powell, H. V. Ravinder, Terry Reilly, Gerald Rose, Sam Roy, Rakesh Sarin, Ross Shachter, Jim Smith, L. James Valverde, M.I.T., Bob Winkler, and Wayne Winston. Special thanks to Deborah Amaral for guidance in writing the Municipal Solid Waste case in Chapter 9; to Dave Braden for outstanding feedback as he and his students used manuscript versions of the first edition; to Susan Brodt for guidance and suggestions for rewriting the creativity material in Chapter 6; and to Kevin McCardle for allowing me to use numerous problems from his statistics course. Thanks also to all of the editors who have worked closely with me over the years: Patty Adams, Marcia Cole, Mary Douglas, Curt Hinrichs, and Michael Payne.

Finally, thanks once again to my wife, Sheri, and to my children for their forbearance while I worked many long hours on both the first and second editions. I have not promised anything, but I am sure they are hoping that the second edition will be the last!

Robert T. Clemen

Introduction to Decision Analysis

Have you ever had a difficult decision to make? If so, did you wish for a straightforward way to keep all of the different issues clear? Did you end up making the decision based on your intuition or on a "hunch" that seemed correct? At one time or another, all of us have wished that a hard decision was easy to make. The sad fact is that hard decisions are just that — hard. As individuals we run into such difficult decisions frequently. Business executives and governmental policy makers struggle with hard problems all the time. For example, consider the following problem faced by the Oregon Department of Agriculture (ODA) in 1985.

GYPSY MOTHS AND THE ODA

In the winter of 1985, the ODA grappled with the problem of gypsy moth infestation in Lane County in western Oregon. Forest industry representatives argued strongly for an aggressive eradication campaign using potent chemical insecticides. The ODA instead proposed a plan that involved spraying most of the affected area with BT (*Bacillus thuringiensis*), a bacterial insecticide known to be (1) target-specific (that is, it does little damage to organisms other than moths), (2) ecologically safe, and (3) reasonably effective. As well as using BT, the ODA proposed spraying three smaller areas near the city of Eugene with the chemical spray Orthene. Although Orthene was registered as an acceptable insecticide for home garden use, there was some doubt as to its ultimate ecological effects as well as its danger to humans. Forestry officials argued that the chemical insecticide was more potent than BT and

was necessary to ensure eradication in the most heavily infested areas. Environmentalists argued that the potential danger from the chemical spray was too great to warrant its use. Some individuals argued that spraying would not help because the infestation already was so advanced that no program would be successful. Others argued that an aggressive spray program could solve the problem once and for all, but only if done immediately. Clearly, in making its final decision the ODA would have to deal with many issues.

The ODA has an extremely complex problem on its hands. Before deciding exactly what course of action to take, the agency needs to consider many issues, including the values of different constituent groups and the uncertainties involving the effectiveness and risks of the pesticides under consideration. The ODA must consider these issues carefully and in a balanced way — but how? There is no escaping the problem: This hard decision requires hard thinking.

Decision analysis provides structure and guidance for thinking systematically about hard decisions. With decision analysis, a decision maker can take action with confidence gained through a clear understanding of the problem. Along with a conceptual framework for thinking about hard problems, decision analysis provides analytical tools that can make the required hard thinking easier.

Why Are Decisions Hard?

What makes decisions hard? Certainly different problems may involve different and often special difficulties. For example, the ODA's problem requires it to think about the interests of various groups as well as to consider only limited information on the possible effects of the sprays. Although every decision may have its own special problems, there are four basic sources of difficulty. A decision-analysis approach can help a decision maker with all four.

First, a decision can be hard simply because of its complexity. In the case of the gypsy moths, the ODA must consider many different individual issues: the uncertainty surrounding the different sprays, the values held by different community groups, the different possible courses of action, the economic impact of any pest-control program, and so on. Simply keeping all of the issues in mind at one time is nearly impossible. Decision analysis provides effective methods for organizing a complex problem into a structure that can be analyzed. In particular, elements of a decision's structure include the possible courses of action, the possible outcomes that could result, the likelihood of those outcomes, and eventual consequences (e.g., costs and benefits) to be derived from the different outcomes. Structuring tools that we will consider include decision trees and influence diagrams as well as procedures for analyzing these structures to find solutions and for answering "what if" questions.

Second, a decision can be difficult because of the inherent uncertainty in the situation. In the gypsy moth case, the major uncertainties are the effectiveness of the different sprays in reducing the moth population and their potential for detrimental ecological and health effects. In some decisions the main issue is uncertainty. For example,

imagine a firm trying to decide whether to introduce a new product. The size of the market, the market price, eventual competition, and manufacturing and distribution costs all may be uncertain to some extent, and all have some impact on the firm's eventual payoff. Yet the decision must be made without knowing for sure what these uncertain values will be. A decision-analysis approach can help in identifying important sources of uncertainty and representing that uncertainty in a systematic and useful way.

Third, a decision maker may be interested in working toward multiple objectives, but progress in one direction may impede progress in others. In such a case, a decision maker must trade off benefits in one area against costs in another. In the gypsy moth example, important trade-offs must be made: Are the potential economic benefits to be gained from spraying Orthene worth the potential ecological damage and health risk? In investment decisions a trade-off that we usually must make is between expected return and riskiness. Decision analysis again provides both a framework and specific tools for dealing with multiple objectives.

Fourth, and finally, a problem may be difficult if different perspectives lead to different conclusions. Or, even from a single perspective, slight changes in certain inputs may lead to different choices. This source of difficulty is particularly pertinent when more than one person is involved in making the decision. Different individuals may look at the problem from different perspectives, or they may disagree on the uncertainty or value of the various outcomes. The use of the decision-analysis framework and tools can help sort through and resolve these differences whether the decision maker is an individual or a group of stakeholders with diverse opinions.

Why Study Decision Analysis?

The obvious reason for studying decision analysis is that carefully applying its techniques can lead to better decisions. But what is a good decision? A simple answer might be that it is the one that gives the best outcome. This answer, however, confuses the idea of a lucky outcome with a good decision. Suppose that you are interested in investing an inheritance. After carefully considering all the options available and consulting with investment specialists and financial planners, you decide to invest in stocks. If you purchased a portfolio of stocks in 1982, the investment most likely turned out to be a good one, because stock values increased dramatically during the 1980s. On the other hand, if your stock purchase had been in early 1929, the stock market crash and the following depression would have decreased the value of your portfolio drastically.

Was the investment decision a good one? It certainly could have been if it was made after careful consideration of the available information and thorough deliberation about the goals and possible outcomes. Was the outcome a good one? For the 1929 investor, the answer is no. This example illustrates the difference between a good decision and a lucky outcome: You can make a good decision but still have an unlucky outcome. Of course, you may prefer to have lucky outcomes rather than make good decisions! Although decision analysis cannot improve your luck, it can help you to understand better the problems you face and thus make better decisions. That understanding must

include the structure of the problem as well as the uncertainty and trade-offs inherent in the alternatives and outcomes. You may then improve your chances of enjoying a better outcome; more important, you will be less likely to experience unpleasant surprises in the form of unlucky outcomes that were either unforeseen or not fully understood. In other words, you will be making a decision with your eyes open.

The preceding discussion suggests that decision analysis allows people to make effective decisions more consistently. This idea itself warrants discussion. Decision analysis is intended to help people deal with *difficult* decisions. It is a "prescriptive approach designed for normally intelligent people who want to think hard and systematically about some important real problems" (Keeney and Raiffa 1976, p. vii).

This prescriptive view is the most appropriate way to think about decision analysis. It gets across the idea that although we are not perfect decision makers, we can do better through more structure and guidance. We will see that decision analysis is not an idealized theory designed for superrational and omniscient beings. Nor does it describe how people actually make decisions. In fact, ample experimental evidence from psychology shows that people generally do not process information and make decisions in ways that are consistent with the decision-analysis approach. (If they did, then there would be no need for decision analysis; why spend a lot of time studying decision analysis if it suggests that you do what you already do?) Instead, using some fundamental principles, and informed by what we know about human frailties in judgment and decision making, decision analysis offers guidance to normal people working on hard decisions.

Although decision analysis provides structure and guidance for systematic thinking in difficult situations, it does not claim to recommend an alternative that must be blindly accepted. Indeed, after the hard thinking that decision analysis fosters, there should be no need for blind acceptance; the decision maker should understand the situation thoroughly. Instead of providing solutions, decision analysis is perhaps best thought of as simply an information source, providing insight about the situation, uncertainty, objectives, and trade-offs, and possibly yielding a recommended course of action. Thus, decision analysis does not usurp the decision maker's job. According to another author,

[t]he basic presumption of decision analysis is not at all to replace the decision maker's intuition, to relieve him or her of the obligations in facing the problem, or to be, worst of all, a competitor to the decision maker's personal style of analysis, but to complement, augment, and generally work alongside the decision maker in exemplifying the nature of the problem. Ultimately, it is of most value if the decision maker has actually learned something about the problem and his or her own decision-making attitude through the exercise (Bunn 1984, p. 8).

We have been discussing decision analysis as if it were always used to help an individual make a decision. Indeed, this is what it is designed for, but its techniques have many other uses. For example, one might use decision-analysis methods to solve complicated inference problems (that is, answering questions such as "What conclusions can be drawn from the available evidence?"). Structuring a decision problem may be useful for understanding its precise nature, for generating alternative courses of action, and for identifying important objectives and trade-offs. Understanding trade-offs can be crucial for making progress in negotiation settings. Finally, decision analysis can be used to justify why a previously chosen action was appropriate.

Subjective Judgments and Decision Making

Personal judgments about uncertainty and values are important inputs for decision analysis. It will become clear through this text that discovering and developing these judgments involves thinking hard and systematically about important aspects of a decision.

Managers and policy makers frequently complain that analytical procedures from management science and operations research ignore subjective judgments. Such procedures often purport to generate "optimal" actions on the basis of purely objective inputs. But the decision-analysis approach allows the inclusion of subjective judgments. In fact, decision analysis *requires* personal judgments; they are important ingredients for making good decisions.

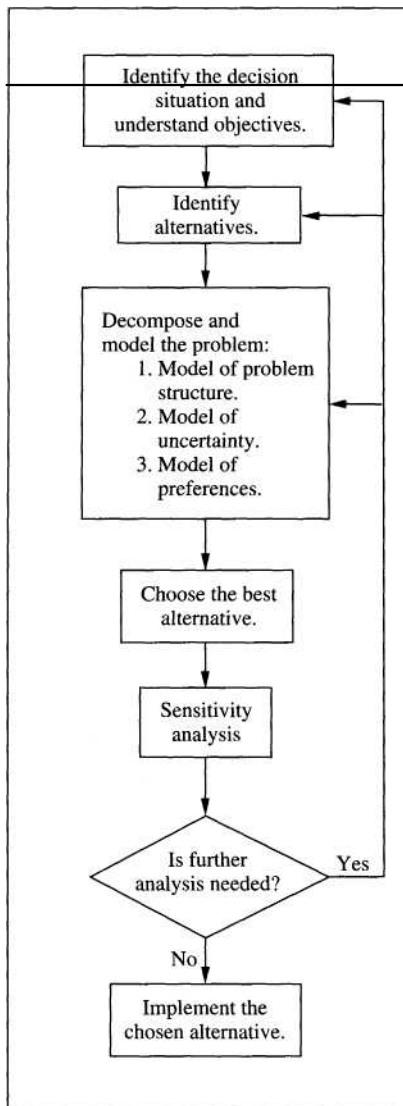
At the same time, it is important to realize that human beings are imperfect information processors. Personal insights about uncertainty and preferences can be both limited and misleading, even while the individual making the judgments may demonstrate an amazing overconfidence. An awareness of human cognitive limitations is critical in developing the necessary judgmental inputs, and a decision maker who ignores these problems can magnify rather than adjust for human frailties. Much current psychological research has a direct bearing on the practice of decision-analysis techniques. In the chapters that follow, many of the results from this research will be discussed and related to decision-analysis techniques. The spirit of the discussion is that understanding the problems people face and carefully applying decision-analysis techniques can lead to better judgments and improved decisions.

The Decision-Analysis Process

Figure 1.1 shows a flowchart for the decision-analysis process. The first step is for the decision maker to identify the decision situation and to understand his or her objectives in that situation. Although we usually do not have trouble finding decisions to make or problems to solve, we do sometimes have trouble identifying the exact problem, and thus we sometimes treat the wrong problem. Such a mistake has been called an "error of the third kind." Careful identification of the decision at hand is always important. For example, perhaps a surface problem hides the real issue. For example, in the gypsy moth case, is the decision which insecticide to use to control the insects, or is it how to mollify a vocal and ecologically minded minority?

Understanding one's objectives in a decision situation is also an important first step and involves some introspection. What is important? What are the objectives? Minimizing cost? Maximizing profit or market share? What about minimizing risks? Does risk mean the chance of a monetary loss, or does it refer to conditions potentially damaging to health and the environment? Getting a clear understanding of the crucial objectives in a decision situation must be done before much more can be accomplished. In the next step, knowledge of objectives can help in identifying

Figure 1.1
A decision-analysis process flowchart.



alternatives, and beyond that the objectives indicate how outcomes must be measured and what kinds of uncertainties should be considered in the analysis.

Many authors argue that the first thing to do is to identify the problem and then to figure out the appropriate objectives to be used in addressing the problem. But Keeney (1992) argues the opposite; it is far better, he claims, to spend a lot of effort understanding one's central values and objectives, and then looking for ways — decision opportunities — to achieve those objectives. The debate notwithstanding, the fact is that decisions come in many forms. Sometimes we are lucky enough to shape our decision-making future in the way Keeney suggests, and other times we find ourselves in diffi-

cult situations that we may not have anticipated. In either case, establishing the precise nature of the decision situation (which we will later call the *decision context*) goes hand in hand with identifying and understanding one's objectives in that situation.

With the decision situation and pertinent objectives established, we turn to the discovery and creation of alternatives. Often a careful examination and analysis of objectives can reveal alternatives that were not obvious at the outset. This is an important benefit of a decision-analysis approach. In addition, research in the area of creativity has led to a number of techniques that can improve the chance of finding new alternatives.

The next two steps, which might be called "modeling and solution," form the heart of most textbooks on decision analysis, including this one. Much of this book will focus on decomposing problems to understand their structures and measure uncertainty and value; indeed, decomposition is the key to decision analysis. The approach is to "divide and conquer." The first level of decomposition calls for structuring the problem in smaller and more manageable pieces. Subsequent decomposition by the decision maker may entail careful consideration of elements of uncertainty in different parts of the problem or careful thought about different aspects of the objectives.

The idea of *modeling* is critical in decision analysis, as it is in most quantitative or analytical approaches to problems. As indicated in Figure 1.1, we will use models in several ways. We will use influence diagrams or decision trees to create a representation or model of the decision problem. Probability will be used to build models of the uncertainty inherent in the problem. Hierarchical and network models will be used to understand the relationships among multiple objectives, and we will assess utility functions in order to model the way in which decision makers value different outcomes and trade off competing objectives. These models are mathematical and graphical in nature, allowing one to find insights that may not be apparent on the surface. Of course, a key advantage from a decision-making perspective is that the mathematical representation of a decision can be subjected to analysis, which can indicate a "preferred" alternative.

Decision analysis is typically an iterative process. Once a model has been built, *sensitivity analysis* is performed. Such analysis answers "what if" questions: "If we make a slight change in one or more aspects of the model, does the optimal decision change?" If so, the decision is said to be sensitive to these small changes, and the decision maker may wish to reconsider more carefully those aspects to which the decision is sensitive. Virtually any part of a decision is fair game for sensitivity analysis. The arrows in Figure 1.1 show that the decision maker may return even to the identification of the problem. It may be necessary to refine the definition of objectives or include objectives that were not previously included in the model. New alternatives may be identified, the model structure may change, and the models of uncertainty and preferences may need to be refined. The term *decision-analysis cycle* best describes the overall process, which may go through several iterations before a satisfactory solution is found.

In this iterative process, the decision maker's perception of the problem changes, beliefs about the likelihood of various uncertain eventualities may develop and change, and preferences for outcomes not previously considered may mature as more time is spent in reflection. Decision analysis not only provides a structured way to think about decisions, but also more fundamentally provides a structure

within which a decision maker can develop beliefs and feelings, those subjective judgments that are critical for a good solution.

Requisite Decision Models

Phillips (1982, 1984) has introduced the term *requisite decision modeling*. This marvelous term captures the essence of the modeling process in decision analysis. In Phillips's words, "a model can be considered requisite only when no new intuitions emerge about the problem" (1984, p. 37), or when it contains everything that is essential for solving the problem. That is, a model is requisite when the decision maker's thoughts about the problem, beliefs regarding uncertainty, and preferences are fully developed. For example, consider a first-time mutual-fund investor who finds high, overall long-term returns appealing. Imagine, though, that in the process of researching the funds the investor begins to understand and become wary of highly volatile stocks and mutual funds. For this investor, a decision model that selected a fund by maximizing the average return in the long run would not be requisite. A requisite model would have to incorporate a trade-off between long-term returns and volatility.

A careful decision maker may cycle through the process shown in Figure 1.1 several times as the analysis is refined. Sensitivity analysis at appropriate times can help the decision maker choose the next modeling steps to take in developing a requisite model. Successful decision analysts artistically use sensitivity analysis to manage the iterative development of a decision model. An important goal of this book is that you begin to acquire this artistic ability through familiarity and practice with the concepts and tools of decision analysis.

Where Is Decision Analysis Used?

Decision analysis is widely used in business and government decision making. Perusing the literature reveals applications that include managing research-and-development programs, negotiating for oil and gas leases, forecasting sales for new products, understanding the world oil market, deciding whether to launch a new product or new venture, and developing ways to respond to environmental risks, to name a few. And some of the largest firms make use of decision analysis, including General Motors, Chevron, and Eli Lilly. A particularly important arena for decision-analysis applications has been in public utilities, especially electric power generation. In part this is because the problems utilities face (e.g., site selection, power-generation methods, waste cleanup and storage, pollution control) are particularly appropriate for treatment with decision-analysis techniques; they involve long time frames and hence a high degree of uncertainty. In addition, multiple objectives must be considered when a decision affects many different stakeholder groups.

In the literature, many of the reported applications relate to public-policy problems and relatively few to commercial decisions, partly because public-policy problems are of interest to such a wide audience. It is perhaps more closely related

to the fact that commercial applications often are proprietary; a good decision analysis can create a competitive advantage for a firm, which may not appreciate having its advantage revealed in the open literature. Important public-policy applications have included regulation in the energy (especially nuclear) industry and standard setting in a variety of different situations ranging from regulations for air and water pollution to standards for safety features on new cars.

Another important area of application for decision analysis has been in medicine. Decision analysis has helped doctors make specific diagnoses and individuals to understand the risks of different treatments. Institutional-level studies have been done such as studying the optimal inventory or usage of blood in a blood bank or the decision of a firm regarding different kinds of medical insurance to provide its employees. On a grander scale, studies have examined policies such as widespread testing for various forms of cancer or the impact on society of different treatment recommendations.

This discussion is by no means exhaustive; the intent is only to give you a feel for the breadth of possible applications of decision analysis and a glimpse at some of the things that have been done. Many other applications are described in cases and examples throughout the book; by the time you have finished, you should have a good understanding of how decision analysis can be (and is) used in many different arenas. And if you feel the need for more, articles by Ulvila and Brown (1982) and Corner and Kirkwood (1991) describe many different applications.

Where Are We Going from Here?

This book is divided into three main sections. The first is titled "Modeling Decisions," and it introduces influence diagrams and decision trees as methods for building models of decision problems. The process is sometimes called *structuring* because it specifies the elements of the decision and how the elements are interrelated (Chapters 2 and 3). We also introduce ways to organize a decision maker's values into hierarchies and networks; doing so is useful when multiple objectives must be considered. We will find out how to analyze our decision models (Chapter 4) and how to conduct sensitivity analysis (Chapter 5). In Chapter 6 we discuss creativity and decision making.

The second section is "Modeling Uncertainty." Here we delve into the use of probability for modeling uncertainty in decision problems. First we review basic probability concepts (Chapter 7). Because subjective judgments play a central role in decision analysis, subjective assessments of uncertainty are the topic of Chapter 8. Other ways to use probability include theoretical probability models (Chapter 9), data-based models (Chapter 10), and simulation (Chapter 11). Chapter 12 closes the section with a discussion of information and how to value it in the context of a probability model of uncertainty within a decision problem.

"Modeling Preferences" is the final section. Here we turn to the development of a mathematical representation of a decision maker's preferences, including the identification of desirable objectives and trade-offs between conflicting objectives. A fundamental issue that we often must confront is how to trade off riskiness and expected value. Typically, if we want to increase our chances at a better outcome, we must accept a

simultaneous risk of loss. Chapters 13 and 14 delve into the problem of modeling a decision maker's attitude toward risk. Chapters 15 and 16 complete the section with a treatment of other conflicting objectives. In these chapters we will complete the discussion of multiple objectives begun in Section 1, showing how to construct a mathematical model that reflects subjective judgments of relative importance among competing objectives.

By the end of the book, you will have learned all of the basic techniques and concepts that are central to the practice of modern decision analysis. This does not mean that your hard decisions will suddenly become easy! But with the decision-analysis framework, and with tools for modeling decisions, uncertainty, and preferences, you will be able to approach your hard decisions systematically. The understanding and insight gained from such an approach will give you confidence in your actions and allow for better decisions in difficult situations. That is what the book is about—an approach that will help you to make hard decisions.

SUMMARY

The purpose of decision analysis is to help a decision maker think systematically about complex problems and to improve the quality of the resulting decisions. In this regard, it is important to distinguish between a good decision and a lucky outcome. A good decision is one that is made on the basis of a thorough understanding of the problem and careful thought regarding the important issues. Outcomes, on the other hand, may be lucky or unlucky, regardless of decision quality.

In general, decision analysis consists of a framework and a tool kit for dealing with difficult decisions. The incorporation of subjective judgments is an important aspect of decision analysis, and to a great extent mature judgments develop as the decision maker reflects on the decision at hand and develops a working model of the problem. The overall strategy is to decompose a complicated problem into smaller chunks that can be more readily analyzed and understood. These smaller pieces then can be brought together to create an overall representation of the decision situation. Finally, the decision-analysis cycle provides the framework within which a decision maker can construct a requisite decision model, one that contains the essential elements of the problem and from which the decision maker can take action.

QUESTIONS AND PROBLEMS

- 1.1 Give an example of a good decision that you made in the face of some uncertainty. Was the outcome lucky or unlucky? Can you give an example of a poorly made decision whose outcome was lucky?
- 1.2 Explain how modeling is used in decision analysis. What is meant by the term "requisite decision model"?
- 1.3 What role do subjective judgments play in decision analysis?
- 1.4 At a dinner party, an acquaintance asks whether you have read anything interesting lately, and you mention that you have begun to read a text on decision analysis. Your

friend asks what decision analysis is and why anyone would want to read a book about it, let alone write one! How would you answer?

- 1.5 Your friend in Question 1.4, upon hearing your answer, is delighted! "This is marvelous," she exclaims. "I have this very difficult choice to make at work. I'll tell you the facts, and you can tell me what I should do!" Explain to her why you cannot do the analysis for her.
- 1.6 Give an example in which a decision was complicated because of difficult preference trade-offs. Give one that was complicated by uncertainty.
- 1.7 In the gypsy moth example, what are some of the issues that you would consider in making this decision? What are the alternative courses of action? What issues involve uncertainty, and how could you get information to help resolve that uncertainty? What are the values held by opposing groups? How might your decision trade off these values?
- 1.8 Can you think of some different alternatives that the ODA might consider for controlling the gypsy moths?
- 1.9 Describe a decision that you have had to make recently that was difficult. What were the major issues? What were your alternatives? Did you have to deal with uncertainty? Were there important trade-offs to make?
- 1.10 "Socially responsible investing" first became fashionable in the 1980s. Such investing involves consideration of the kinds of businesses that a firm engages in and selection of investments that are as consistent as possible with the investor's sense of ethical and moral business activity. What trade-offs must the socially responsible investor make? How are these trade-offs more complicated than those that we normally consider in making investment decisions?
- 1.11 Many decisions are simple, preprogrammed, or already solved. For example, retailers do not have to think long to decide how to deal with a new customer. Some operations-research models provide "ready-made" decisions, such as finding an optimal inventory level using an order-quantity formula or determining an optimal production mix using linear programming. Contrast these decisions with unstructured or strategic decisions, such as choosing a career or locating a nuclear power plant. What kinds of decisions are appropriate for a decision-analysis approach? Comment on the statement, "Decision making is what you do when you don't know what to do." (For more discussion, see Howard 1980.)
- 1.12 The argument was made that beliefs and preferences can change as we explore and learn. This even holds for learning about decision analysis! For example, what was your impression of this book before reading the first chapter? Have your beliefs about the value of decision analysis changed? How might this affect your decision about reading more of the book?

C A S E S T U D I E S

DR. JOYCELYN ELDERS AND THE WAR ON DRUGS

After the Nancy Reagan slogan, "Just Say No," and 12 years of Republican administration efforts to fight illegal drug use and trafficking, on December 7, 1993, then Surgeon General Dr. Joycelyn Elders made a startling statement. In response to a reporter's question, she indicated that, based on the experiences of other countries, the

crime rate in the United States might actually decrease if drugs were legalized. She conceded that she did not know all of the ramifications and suggested that perhaps some studies should be done.

The nation and especially the Clinton administration were shocked to hear this statement. What heresy after all the efforts to control illegal drugs! Of course, the White House immediately went on the defensive, making sure that everyone understood that President Clinton was not in favor of legalizing drugs. And Dr. Elders had to clarify her statement; it was her personal opinion, not a statement of administration policy.

Questions

- 1 What decision situation did Dr. Elders identify? What specific values would be implied by choosing to study the legalization of drugs?
- 2 From a decision-making perspective, which makes more sense: Nancy Reagan's "Just Say No" policy or Elders's suggestion that the issue of legalization be studied? Why?
- 3 Consider Elders's decision to suggest studying the legalization of drugs. Was her decision to respond to the reporter the way she did a good decision with a bad outcome? Or was it a bad decision in the first place?
- 4 Why was Elders's suggestion a political hot potato for Clinton's administration? What, if any, are the implications for decision analysis in political situations?

LLOYD BENTSEN FOR VICE PRESIDENT?

In the summer of 1988, Michael Dukakis was the Democratic Party's presidential nominee. The son of Greek immigrants, his political career had flourished as governor of Massachusetts, where he had demonstrated excellent administrative and fiscal skills. He chose Lloyd Bentsen, U.S. Senator from Texas, as his running mate. In an analysis of Dukakis's choice, E. J. Dionne of *The New York Times* (July 13, 1988) made the following points:

- 1 The main job of the vice presidential nominee is to carry his or her home state. Could Bentsen carry Texas? The Republican presidential nominee was George Bush, whose own adopted state was Texas. Many people thought that Texas would be very difficult for Dukakis to win, even with Bentsen's help. If Dukakis could win Texas's 29 electoral votes, however, the gamble would pay off dramatically, depriving Bush of one of the largest states that he might have taken for granted.
- 2 Bentsen was a conservative Democrat. Jesse Jackson had run a strong race and had assembled a strong following of liberal voters. Would the Jackson supporters be disappointed in Dukakis's choice? Or would they ultimately come back to the fold and be faithful to the Democratic Party?

- 3 Bentsen's ties with big business were unusual for a Democratic nominee. Would Democratic voters accept him? The other side of this gamble was that Bentsen was one of the best fund raisers around and might be able to eliminate or even reverse the Republicans' traditional financial advantage. Even if some of the more liberal voters were disenchanted, Bentsen could appeal to a more business-oriented constituency.
- 4 The safer choice for a running mate would have been Senator John Glenn from Ohio. The polls suggested that with Glenn as his running mate, Dukakis would have no trouble winning Ohio and its 23 electoral votes.

Questions

- 1 Why is choosing a running mate a hard decision?
- 2 What objectives do you think a presidential nominee should consider in making the choice?
- 3 What elements of risk are involved?
- 4 The title of Dionne's article was "Bentsen: Bold Choice or Risky Gamble?" In what sense was Dukakis's decision a "bold choice," and in what sense was it a "risky gamble"?

DU PONT AND CHLOROFLUOROCARBONS

Chlorofluorocarbons (CFCs) are chemicals used as refrigerants in air conditioners and other cooling appliances, propellants in aerosol sprays, and in a variety of other applications. Scientific evidence has been accumulating for some time that CFCs released into the atmosphere can destroy ozone molecules in the ozone layer 15 miles above the earth's surface. This layer shields the earth from dangerous ultraviolet radiation. A large hole in the ozone layer above Antarctica has been found and attributed to CFCs, and a 1988 report by 100 scientists concluded that the ozone shield above the mid-Northern Hemisphere had shrunk by as much as 3% since 1969. Moreover, depletion of the ozone layer appears to be irreversible. Further destruction of the ozone layer could lead to crop failures, damage to marine ecology, and possibly dramatic changes in global weather patterns.

Environmentalists estimate that approximately 30% of the CFCs released into the atmosphere come from aerosols. In 1978, the U.S. government banned their use as aerosol propellants, but many foreign governments still permit them.

Some \$2.5 billion of CFCs are sold each year, and Du Pont Chemical Corporation is responsible for 25% of that amount. In early 1988, Du Pont announced that the company would gradually phase out its production of CFCs and that replacements would be developed. Already Du Pont claims to have a CFC substitute for automobile air conditioners, although the new substance is more expensive.

Questions

Imagine that you are a Du Pont executive charged with making the decision regarding continued production of CFCs.

- 1 What issues would you take into account?
- 2 What major sources of uncertainty do you face?
- 3 What corporate objectives would be important for you to consider? Do you think that Du Pont's corporate objectives and the way the company views the problem might have evolved since the mid-1970s when CFCs were just beginning to be come an issue?

Sources: "A Gaping Hole in the Sky," *Newsweek*, July 11, 1988, pp. 21-23, A. M. Louis (1988) "Du Pont to Ban Products that Harm Ozone," *San Francisco Chronicle*, March 25, p. 1.

REFERENCES

The decision-analysis view is distinctly *prescriptive*. That is, decision analysis is interested in helping people make better decisions; in contrast, a *descriptive* view of decision making focuses on how people actually make decisions. Keeney and Raiffa (1976) explain the prescriptive view as well as anyone. For an excellent summary of the descriptive approach, see Hogarth (1987). Bell, Raiffa, and Tversky (1988) provide many readings on these topics.

A fundamental element of the prescriptive approach is discerning and accepting the difference between a good decision and a lucky outcome. This issue has been discussed by many authors, both academics and practitioners. An excellent recent reference is Vlek et al. (1984).

Many other books and articles describe the decision-analysis process, and each seems to have its own twist. This chapter has drawn heavily from Ron Howard's thoughts; his 1988 article summarizes his approach. Other books worth consulting include Behn and Vaupel (1982), Bunn (1984), Holloway (1979), Keeney (1992), Lindley (1985), Raiffa (1968), Samson (1988), and von Winterfeldt and Edwards (1986).

Phillips's (1982, 1984) idea of a requisite decision model is a fundamental concept that we will use throughout the text. For a related view, see Watson and Buede (1987).

Behn, R. D., and J. D. Vaupel (1982) *Quick Analysis for Busy Decision Makers*. New York: Basic Books.

Bell, D., H. Raiffa, and A. Tversky (1988) *Decision Making: Descriptive, Normative, and Prescriptive Interactions*. Cambridge, MA: Cambridge University Press.

Bunn, D. (1984) *Applied Decision Analysis*. New York: McGraw-Hill.

Corner, J. L., and C. W. Kirkwood (1991) "Decision Analysis Applications in the Operations Research Literature, 1970-1989." *Operations Research*, 39, 206-219.

Hogarth, R. (1987) *Judgement and Choice*, 2nd ed. New York: Wiley.

Holloway, C. A. (1979) *Decision Making under Uncertainty: Models and Choices*. Englewood Cliffs, NJ: Prentice-Hall.

Howard, R. A. (1980) "An Assessment of Decision Analysis." *Operations Research*, 28, 4-27.

- Howard, R. A. (1988) "Decision Analysis: Practice and Promise," *Management Science*, 34, 679-695.
- Keeney, R. (1992) *Value-Focused Thinking*. Cambridge, MA: Harvard University Press.
- Keeney, R., and H. Raiffa (1976) *Decisions with Multiple Objectives*. New York: Wiley.
- Lindley, D. V. (1985) *Making Decisions*, 2nd ed. New York: Wiley.
- Phillips, L. D. (1982) "Requisite Decision Modelling." *Journal of the Operational Research Society*, 33, 303-312.
- Phillips, L. D. (1984) "A Theory of Requisite Decision Models." *Acta Psychologica*, 56, 29-48.
- Raiffa, H. (1968) *Decision Analysis*. Reading, MA: Addison-Wesley.
- Samson, D. (1988) *Managerial Decision Analysis*. Homewood, IL: Irwin.
- Ulvila, J. W., and R. V. Brown (1982) "Decision Analysis Comes of Age." *Harvard Business Review*, September-October 1982, 130-141.
- Vlek, C., W. Edwards, I. Kiss, G. Majone, and M. Toda (1984) "What Constitutes a Good Decision?" *Acta Psychologica*, 56, 5-27.
- von Winterfeldt, D., and W. Edwards (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press.
- Watson, S., and D. Buede (1987) *Decision Synthesis*. Cambridge: Cambridge University Press.

EPILOGUE

What did the ODA decide? Its directors decided to use only BT on all 227,000 acres, which were sprayed on three separate occasions in late spring and early summer 1985. At the time, this was the largest gypsy moth-control program ever attempted in Oregon. In 1986, 190,000 acres were sprayed, also with BT. Most of the areas sprayed the second year had not been treated the first year because ODA had found later that the gypsy moth infestation was more widespread than first thought. In the summer of 1986, gypsy moth traps throughout the area indicated that the population was almost completely controlled. In the spring of 1987, the ODA used BT to spray only 7500 acres in 10 isolated pockets of gypsy moth populations on the fringes of the previously sprayed areas. By 1988, the spray program was reduced to a few isolated areas near Eugene, and officials agreed that the gypsy moth population was under control.

S E C T I O N
1

Modeling Decisions

T

his first section is about modeling decisions. Chapter 2 presents a short discussion on the elements of a decision. Through a series of simple examples, the basic elements are illustrated: values and objectives, decisions to be made, upcoming uncertain events, and consequences. The focus is on identifying the basic elements. This skill is necessary for modeling decisions as described in Chapters 3, 4, and 5.

In Chapter 3, we learn how to create graphical structures for decision models. First we consider values and objectives, discussing in depth how multiple objectives can be organized in hierarchies and networks that can provide insight and help to generate creative alternatives. We also develop both influence diagrams and decision trees as graphical modeling tools for representing the basic structure of decisions. An influence diagram is particularly useful for developing the structure of a complex decision problem because it allows many aspects of a problem to be displayed in a compact and intuitive form. A decision-tree representation provides an alternative picture of a decision in which more of the details can be displayed. Both graphical techniques can be used to represent single-objective decisions, but we show how they can be used in multiple-objective situations as well. We end Chapter 3 with a discussion of measurement, presenting concepts and techniques which can be used to ensure that we can adequately measure achievement of our objectives, whether those

objectives are straightforward (e.g., maximizing dollars or saving time) or more difficult to quantify (e.g., minimizing environmental damage).

Chapters 4 and 5 present the basic tools available to the decision maker for analyzing a decision model. Chapter 4 shows how to solve decision trees and influence diagrams. The basic concept presented is *expected value*. When we are concerned with monetary outcomes, we call this *expected monetary value* and abbreviate it as EMV. In analyzing a decision, EMV is calculated for each of the available alternatives. In many decision situations it is reasonable to choose the alternative with the highest EMV. In addition to the EMV criterion, Chapter 4 also looks briefly at the idea of risk analysis and the uses of a stochastic-dominance criterion for making decisions. Finally, we show how expected value and risk analysis can be used in multiple-objective decisions.

In Chapter 5 we learn how to use sensitivity-analysis tools in concert with EMV calculations in the iterative decision-structuring and analysis process. After an initial basic model is built, sensitivity analysis can tell which of the input variables really matter in the decision and deserve more attention in the model. Thus, with Chapter 5 we bring the discussion of modeling decisions full circle, showing how structuring and analysis are intertwined in the decision-analysis process.

Finally, Chapter 6 delves into issues relating to creativity and decision making. One of the critical aspects of constructing a model of a decision is the determination of viable alternatives. When searching for alternative actions in a decision situation, though, we are subject to a variety of creative blocks that hamper our search for new and different possibilities. Chapter 6 describes these blocks to creativity, discusses creativity from a psychological perspective, and shows how a careful understanding of one's objectives can aid the search for creative alternatives. Several creativity-enhancing techniques are described.

Elements of Decision Problems

Given a complicated problem, how should one begin? A critical first step is to identify the elements of the situation. We will classify the various elements into (1) values and objectives, (2) decisions to make, (3) uncertain events, and (4) consequences. In this chapter, we will discuss briefly these four basic elements and illustrate them in a series of examples.

Values and Objectives

Imagine a farmer whose trees are laden with fruit that is nearly ripe. Even without an obvious problem to solve or decision to make, we can consider the farmer's objectives. Certainly one objective is to harvest the fruit successfully. This may be important because the fruit can then be sold, providing money to keep the farm operating and a profit that can be spent for the welfare of the family. The farmer may have other underlying objectives as well, such as maximizing the use of organic farming methods.

Before we can even talk about making decisions, we have to understand *values* and *objectives*. "Values" is an overused term that can be somewhat ambiguous; here we use it in a general sense to refer to things that matter to you. For example, you may want to learn how to sail and take a trip around the world. Or you may have an objective of learning how to speak Japanese. A scientist may be interested in resolving a specific scientific question. An investor may want to make a lot of money or

gain a controlling interest in a company. A manager, like our farmer with the orchard, may want to earn a profit.

An *objective* is a specific thing that you want to achieve. All of the examples in the previous paragraph refer to specific objectives. As you can tell from the examples, some objectives are related. The farmer may want to earn a profit because it will provide the means to purchase food for the family or to take a trip. The scientist may want to find an answer to an important question in order to gain prestige in the scientific community; that prestige may in turn lead to a higher salary and more research support at a better university.

An individual's objectives taken together make up his or her values. They define what is important to that person in making a decision. We can make an even broader statement: A person's values are the reason for making decisions in the first place! If we did not care about anything, there would not be a reason to make decisions at all, because we would not care how things turned out. Moreover, we would not be able to choose from among different alternatives. Without objectives, it would not be possible to tell which alternative would be the best choice.

Making Money: A Special Objective

In modern western society, most adults work for a living, and if you ask them why, they will all include in their answers something about the importance of making money. It would appear that making money is an important objective, but a few simple questions (Why is money important? What would you do if you had a million dollars?) quickly reveal that money is important because it helps us do things that we want to do. For many people, money is important because it allows us to eat, afford housing and clothing, travel, engage in activities with friends, and generally live comfortably. Many people spend money on insurance because they have an objective of avoiding risks. For very few individuals is money important in and of itself. Unlike King Midas, most of us do not want to earn money simply to have it; money is important because it provides the means by which we can work toward more basic objectives.

Money's role as a trading mechanism in our economy puts it in a special role. Although it is typically not one of our basic objectives, it can serve as a proxy objective in many situations. For example, imagine a young couple who wants to take a vacation. They will probably have to save money for some period of time before achieving this goal, and they will face many choices regarding just how to go about saving their money. In many of these decisions, the main concern will be how much money they will have when they are ready to take their holiday. If they are considering investing their money in a mutual fund, say, they will have to balance the volatility of the fund's value against the amount they can expect to earn over the long run, because most investment decisions require a trade-off between risk and return.

For corporations, money is often a primary objective, and achievement of the objective is measured in terms of increase in the shareholders' wealth through divi-

dends and increased company value. The shareholders themselves can, of course, use their wealth for their own welfare however they want. Because the shareholders have the opportunity to trade their wealth to achieve specific objectives, the company need not be concerned with those objectives but can focus on making its shareholders as wealthy as possible.

Although making money is indeed a special objective, it is important to realize that many situations require a trade-off between making money and some other objective. In many cases, one can *price out* the value of different objectives. When you purchase a car, how much more would you pay to have air conditioning? How much more to get the color of your choice? These questions may be difficult to answer, but we all make related decisions all the time as we decide whether a product or service is worth the price that is asked. In other cases, though, it may not be reasonable to convert everything to dollars. For example, consider the ethical problems faced by a hospital that performs organ transplants. Wealthy individuals can pay more for their operations, and often are willing to do so in order to move up in the queue. The additional money may permit the hospital to purchase new equipment or perform more transplants for needy individuals. But moving the wealthy patient up in the queue will delay surgery for other patients, perhaps with fatal consequences. What if the other patients include young children? Pricing out the lives and risks to the other patients seems like a cold-hearted way to make this decision; in this case, the hospital will probably be better off thinking in terms of its fundamental objectives and how to accomplish them with or without the wealthy patient's fee.

Values and the Current Decision Context

Suppose you have carefully thought about all of your objectives. Among other things you want to do what you can to reduce homelessness in your community, learn to identify birds, send your children to college, and retire at age 55. Having spent the morning figuring out your objectives, you have become hungry and are ready for a good meal. Your decision is where to go for lunch, and it is obvious that the large-scale, overall objectives that you have spent all morning thinking about will not be much help.

You can still think hard about your objectives, though, as you consider your decision. It is just that different objectives are appropriate for this particular decision. Do you want a lot to eat or a little? Do you want to save money? Are you interested in a particular type of ethnic food, or would you like to try a new restaurant? If you are going out with friends, what about their preferences? What about a picnic instead of a restaurant meal?

Each specific decision situation calls for specific objectives. We call the setting in which the decision occurs the *decision context*. In one case, a decision context might be deciding where to go for lunch, in which case the appropriate objectives involve satisfying hunger, spending time with friends, and so on. In another case, the

context might be what to choose for a career, which would call for consideration of more global objectives. What do you want to accomplish in your life?

Values and decision context go hand in hand. On one hand, it is worthwhile to think about your objectives in advance to be prepared for decisions when they arise or so that you can identify new decision opportunities that you might not have thought about before. On the other hand, every decision situation involves a specific context, and that context determines what objectives need to be considered. The idea of a requisite model comes into play here. A requisite decision model includes all of the objectives that matter, and only those that matter, in the decision context at hand. Without all of the appropriate objectives considered, you will be left with the gnawing concern that "something is missing" (which would be true), and considering superfluous or inappropriate objectives can distract you from the truly important issues. When the decision context is specified and appropriate objectives aligned with the context, the decision maker knows what the situation is and exactly why he or she cares about making a decision in that situation.

Finding realistic examples in which individuals or companies use their objectives in decision making is easy. In the following example, the Boeing Company found itself needing to acquire a new supercomputer.

BOEING'S SUPERCOMPUTER

As a large-scale manufacturer of sophisticated aircraft, Boeing needs computing power for tasks ranging from accounting and word processing to computer-aided design, inventory control and tracking, and manufacturing support. When the company's engineering department needed to expand its high-power computing capacity by purchasing a supercomputer, the managers faced a huge task of assembling and evaluating massive amounts of information. There were systems requirements and legal issues to consider, as well as price and a variety of management issues. (*Source*: D. Barnhart, (1993) "Decision Analysis Software Helps Boeing Select Supercomputer." *OR/MS Today*, April, 62—63.)

Boeing's decision context is acquiring supercomputing capacity for its engineering needs. Even though the company undoubtedly has global objectives related to aircraft production, maximizing shareholder wealth, and providing good working conditions for its employees, in the current decision context the appropriate objectives are specific to the company's computing requirements.

Organizing all of Boeing's objectives in this decision context is complex because of the many different computer users involved and their needs. With careful thought, though, management was able to specify five main objectives: minimize costs, maximize performance, satisfy user needs, satisfy organizational needs, and satisfy management issues. Each of these main objectives can be further broken down into different aspects, as shown in Figure 2.1.

Figure 2.1
Objectives for Boeing's
supercomputer.

Supercomputer Objectives				
<u>Cost</u>	<u>Performance</u>	<u>User Needs</u>	<u>Operational Needs</u>	<u>Management Issues</u>
Five-Year Costs	Speed	Installation Date	Square Footage	Vendor Health
Cost of Improved Performance	Throughput	Roll In/Roll Out	Water Cooling	U.S. Ownership
	Memory Size	Ease of Use	Operator Tools	Commitment to Supercomputer
	Disk Size	Software Compatibility	Telecommunications	
	On-Site Performance	Mean Time between Failures	Vendor Support	

Decisions to Make

With the decision context understood and values well in hand, the decision maker can begin to identify specific elements of a decision. Consider our farmer whose fruit crop will need to be harvested soon. If the weather report forecasts mild weather, the farmer has nothing to worry about, but if the forecast is for freezing weather, it might be appropriate to spend some money on protective measures that will save the crop. In such a situation, the farmer has a decision to make, and that decision is whether or not to take protective action. This is a decision that must be made with the available information.

Many situations have as the central issue a decision that must be made right away. There would always be at least two alternatives; if there were no alternatives, then it would not be a matter of making a decision! In the case of the farmer, the alternatives are to take protective action or to leave matters as they are. Of course, there may be a wide variety of alternatives. For example, the farmer may have several strategies for saving the crop, and it may be possible to implement one or more.

Another possibility may be to wait and obtain more information. For instance, if the noon weather report suggests the possibility of freezing weather depending on exactly where a weather system travels, then it may be reasonable to wait and listen to the evening report to get better information. Such a strategy, however, may entail a cost. The farmer may have to pay his hired help overtime if the decision to protect the crop is made late in the evening. Some measures may take time to set up; if the farmer waits, there may not be enough time to implement some of these procedures.

Other possible alternatives are taking out insurance or hedging. For example, the farmer might be willing to pay the harvesting crew a small amount to be available at night if quick action is needed. Insurance policies also may be available to protect

against crop loss (although these typically are not available at the last minute). Any of these alternatives might give the farmer more flexibility but would probably cost something up front.

Identifying the immediate decision to make is a critical step in understanding a difficult decision situation. Moreover, no model of the decision situation can be built without knowing exactly what the decision problem at hand is. In identifying the central decision, it is important also to think about possible alternatives. Some decisions will have specific alternatives (protect the crop or not), while others may involve choosing a specific value out of a range of possible values (deciding on an amount to bid for a company you want to acquire). Other than the obvious alternative courses of action, a decision maker should always consider the possibilities of doing nothing, of waiting to obtain more information, or of somehow hedging against possible losses.

Sequential Decisions

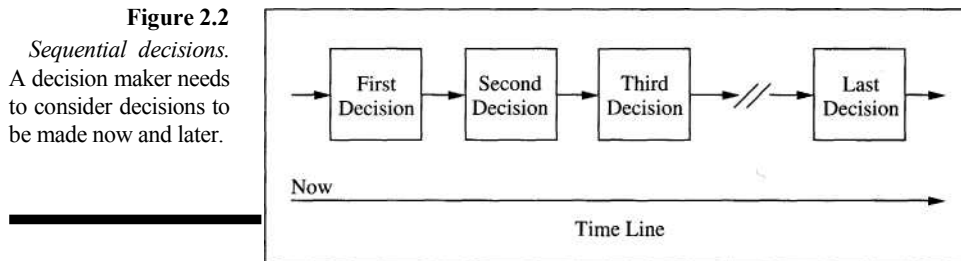
In many cases, there simply is no single decision to make, but several sequential decisions. The orchard example will demonstrate this. Suppose that several weeks of the growing season remain. Each day the farmer will get a new weather forecast, and each time there is a forecast of adverse weather, it will be necessary to decide once again whether to protect the crop.

The example shows clearly that the farmer has a number of decisions to make, and the decisions are ordered sequentially. If the harvest is tomorrow, then the decision is fairly easy, but if several days or weeks remain, then the farmer really has to think about the upcoming decisions. For example, it might be worthwhile to adopt a policy whereby the amount spent on protection is less than the value of the crop. One good way to do this would be not to protect during the early part of the growing season; instead, wait until the harvest is closer, and then protect whenever the weather forecast warrants such action. In other words, "If we're going to lose the crop, let's lose it early."

It is important to recognize that in many situations one decision leads eventually to another in a sequence. The orchard example is a special case because the decisions are almost identical from one day to the next: Take protective action or not. In many cases, however, the decisions are radically different. For example, a manufacturer considering a new product might first decide whether or not to introduce it. If the decision is to go ahead, the next decision might be whether to produce it or subcontract the production. Once the production decision is made, there may be marketing decisions about distribution, promotion, and pricing.

When a decision situation is complicated by sequential decisions, a decision maker will want to consider them when making the immediate decision. Furthermore, a future decision may depend on exactly what happened before. For this reason, we refer to these kinds of problems as *dynamic* decision situations. In identifying elements of a decision situation, we want to know not only what specific decisions are to be made, but the sequence in which they will arise. Figure 2.2 shows graphically a sequence of decisions, represented by squares, mapped along a time line.

Figure 2.2
Sequential decisions.
 A decision maker needs to consider decisions to be made now and later.



Uncertain Events

In Chapter 1 we saw that decision problems can be complicated because of uncertainty about what the future holds. Many important decisions have to be made without knowing exactly what will happen in the future or exactly what the ultimate outcome will be from a decision made today. A classic example is that of investing in the stock market. An investor may be in a position to buy some stock, but in which company? Some share prices will go up and others down, but it is difficult to tell exactly what will happen. Moreover, the market as a whole may move up or down, depending on economic forces. The best the investor can do is think very carefully about the chances associated with each different security's prices as well as the market as a whole.

The possible things that can happen in the resolution of an uncertain event are called *outcomes*. In the orchard example above, the key uncertain event is the weather, with outcomes of crop damage or no crop damage. With some uncertain events, such as with the orchard, there are only a few possible outcomes. In other cases, such as the stock market, the outcome is a value within some range. That is, next year's price of the security bought today for \$50 per share may be anywhere between, say, \$0 and \$100. (It certainly could never be worth less than zero, but the upper limit is not so well defined: Different individuals might consider different upper limits for the same stock.) The point is that the outcome of the uncertain event that we call "next year's stock price" comes from a range of possible values and may fall anywhere within that range.

Many different uncertain events might be considered in a decision situation, but only some are relevant. How can you tell which ones are relevant? The answer is straightforward; the outcome of the event must have some impact on at least one of your objectives. That is, it should matter to you what actually comes to pass. Although this seems like common sense, in a complex decision situation it can be all too easy to concentrate on uncertain events that we can get information about rather than those that really have an impact in terms of our objectives. One of the best examples comes from risk analysis of nuclear power plants; engineers can make judgments about the chance that a power-plant accident will release radioactive material into the atmosphere, but what may really matter is how local residents react to siting the plant in their neighborhood and to subsequent accidents if they occur.

Of course, a decision situation often involves more than one uncertain event. The larger the number of uncertain but relevant events in a given situation, the more

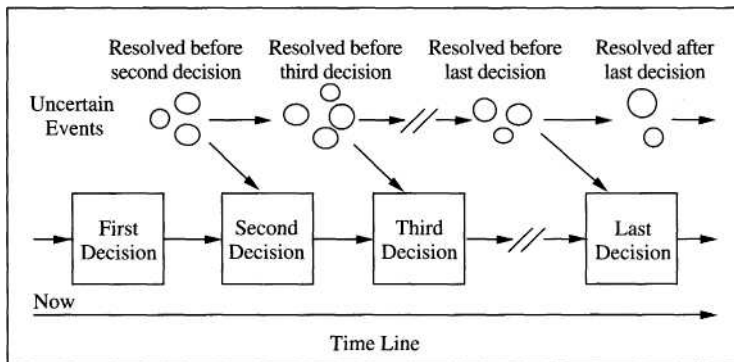
complicated the decision. Moreover, some uncertain events may depend on others. For example, the price of the specific stock purchased may be more likely to go up if the economy as a whole continues to grow or if the overall stock market increases in value. Thus there may be interdependencies among the uncertain events that a decision maker must consider.

How do uncertain events relate to the decisions in Figure 2.2? They must be dovetailed with the time sequence of the decisions to be made; it is important to know at each decision exactly what information is available and what remains unknown. At the current time ("Now" on the time line), all of the uncertain events are just that; their outcomes are unknown, although the decision maker can look into the future and specify which uncertainties will be resolved prior to each upcoming decision. For example, in the dynamic orchard decision, on any given day the farmer knows what the weather has been in the past but not what the weather will be in the future.

Sometimes an uncertain event that is resolved before a decision provides information relevant for future decisions. Consider the stock market problem. If the investor is considering investing in a company that is involved in a lawsuit, one alternative might be to wait until the lawsuit is resolved. Note that the sequence of decisions is (1) wait or buy now, and (2) if waiting, then buy or do not buy after the lawsuit. The decision to buy or not may depend crucially on the outcome of the lawsuit that occurs between the two decisions.

What if there are many uncertain events that occur between decisions? There may be a natural order to the uncertain events, or there may not. If there is, then specifying that order during modeling of the decision problem may help the decision maker. But the order of events between decisions is not nearly as crucial as the dovetailing of decisions and events to clarify what events are unknown and what information is available for each decision in the process. It is the time sequence of the decisions that matters, along with the information available at each decision. In Figure 2.3, uncertain events, represented by circles, are dovetailed with a sequence of decisions. An arrow from a group of uncertain events to a decision indicates that the outcomes of those events are known at the time the decision is made. Of course, the decision maker is like the proverbial elephant and never forgets what has happened. For upcoming decisions, he or she should be able to recall (possibly with the aid of notes and documents) everything that happened (decisions and event outcomes) up to that point.

Figure 2.3
Dovetailing uncertain events and sequential decisions.



Consequences

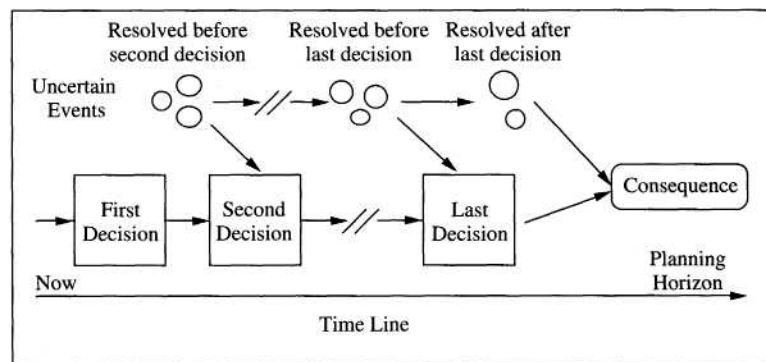
After the last decision has been made and the last uncertain event has been resolved, the decision maker's fate is finally determined. It may be a matter of profit or loss as in the case of the farmer. It may be a matter of increase in value of the investor's portfolio. In some cases the final consequence may be a "net value" figure that accounts for both cash outflows and inflows during the time sequence of the decisions. This might happen in the case of the manufacturer deciding about a new product; certain costs must be incurred (development, raw materials, advertising) before any revenue is obtained.

If the decision context requires consideration of multiple objectives, the consequence is what happens with respect to each of the objectives. For example, consider the consequence of a general's decision to storm a hill. The consequence might be good because the army succeeds in taking the hill (a specific objective), but it may be bad at the same time if many lives are lost.

In our graphical scheme, we must think about the consequence at the end of the time line after all decisions are made and all uncertain events are resolved. For example, the consequence for the farmer after deciding whether to protect and then experiencing the weather might be a profit of \$15,000, a loss of \$3400, or some other dollar amount. For the general it might be "gain the hill, 10 men killed, 20 wounded" or "don't gain the hill, two men killed, five wounded." Thus, the end of the time line is when the decision maker finds out the results. Looking forward from the current time and current decision, the end of the time line is called the *planning horizon*. Figure 2.4 shows how the consequence fits into our graphical scheme.

What is an appropriate planning horizon? For the farmer, the answer is relatively easy; the appropriate planning horizon is at the time of the harvest. But for the general, this question is not so simple. Is the appropriate horizon the end of the next day when he will know whether his men were able to take the hill? Or is it at the end of the war? Or is it sometime in between — say, the end of next month? For the investor, how far ahead should the planning horizon be? A week? A month? Several years? For individuals planning for retirement, the planning horizon may be years in the future. For speculators making trades on the floor of a commodity exchange, the planning horizon may be only minutes into the future.

Figure 2.4
Including the
consequence.



Thus, one of the fundamental issues with which a decision maker must come to grips is how far into the future to look. It is always possible to look farther ahead; there will always be more decisions to make, and earlier decisions may have some effect on the availability of later alternatives. Even death is not an obvious planning horizon because the decision maker may be concerned with effects on future generations; environmental policy decisions provide perfect examples. At some point the decision maker has to stop and say, "My planning horizon is there. It's not worthwhile for me to think beyond that point in time." For the purpose of constructing a requisite model, the idea is to choose a planning horizon such that the events and decisions that would follow after are not essential parts of the immediate decision problem. To put it another way, choose a planning horizon that is consistent with your decision context and the relevant objectives.

Once the dimensions of the consequences and the planning horizon have been determined, the next step is to figure out how to value the consequences. As mentioned, in many cases it will be possible to work in terms of monetary values. That is, the only relevant objective in the decision context is to make money, so all that matters at the end is profit, cost, or total wealth. Or it may be possible to price out nonmonetary objectives as discussed above. For example, a manager might be considering whether to build and run a day care center for the benefit of the employees. One objective might be to enhance the goodwill between the company and the workforce. Enhanced goodwill would in turn have certain effects on the operations of the company, including reduced absenteeism, improved ability to recruit, and a better image in the community. Some of these, such as the reduced absenteeism and improved recruiting, could easily be translated into dollars. The image may be more difficult to translate, but the manager might assess its value subjectively by estimating how much money it would cost in terms of public relations work to improve the firm's image by the same amount.

In some cases, however, it will be difficult to determine exactly how the different objectives should be traded off. In the hospital case discussed earlier, how should the administrator trade off the risks to patients who would be displaced in the queue versus the fee paid by a wealthy patient? How many lives should the general be willing to sacrifice in order to gain the hill? How much damage to the environment are we willing to accept in order to increase the U.S. supply of domestic oil? How much in the way of health risks are we willing to accept in order to have blemish-free fruits and vegetables? Many decisions, especially governmental policy decisions, are complicated by trade-offs like these. Even personal decisions, such as taking a job or purchasing a home, require a decision maker to think hard about the trade-offs involved.

The Time Value of Money: A Special Kind of Trade-Off

One of the most common consequences in personal and business decisions is a stream of cash flows. For example, an investor may spend money on a project (an initial cash outflow) in order to obtain revenue in the future (cash inflows) over a period of years. In such a case, there is a special kind of trade-off: spending dollars today to obtain dollars tomorrow. If a dollar today were worth the same as a dollar

next year, there would be no problem. However, this is not the case. A dollar today can be invested in a savings account or other interest-bearing security; at the end of a year, one dollar invested now would be worth one dollar plus the interest paid.

Trade-offs between current and future dollars (and between future dollars at different points in time) refer to the fact that the value of a dollar depends on when it is available to the decision maker. Because of this, we often refer to the "time value of money." Fortunately, there is a straightforward way to collapse a stream of cash flows into a single number. This number is called the *present value*, or value in present dollars, of the stream of cash flows.

Suppose, for example, you have \$100 in your pocket. If you put that money into a savings account that earns 10% per year, paid annually, then you would have $\$100 \times 1.1 = \110 at the end of the year. At the end of two years, the balance in the account would be \$110 plus another 10%, or $\$110 \times 1.1 = \121 . In fact, you can see that the amount you have is just the original \$100 multiplied by 1.1 twice: $\$121 = \$100 \times 1.1 \times 1.1 = \100×1.1^2 . If you keep the money in the account for five years, say, then the interest compounds for five years. The account balance would be $\$100 \times 1.1^5 = \161.05 .

We are going to use this idea of interest rates to work backward. Suppose, for example, that someone promises that you can have \$110 next year. What is this worth to you right now? If you have available some sort of investment like a savings account that pays 10% per year, then you would have to invest \$100 in order to get \$110 next year. Thus, the present value of the \$110 that arrives next year is just $\$110/1.1 = \100 . Similarly, the present value of \$121 dollars promised at the end of two years is $\$121/(1.1^2) = \100 .

In general, we will talk about the present value of an amount x that will be received at the end of n time periods. Of course, we must know the appropriate interest rate. Let r denote the interest rate per time period in decimal form; that is, if the interest rate is 10%, then $r = 0.10$. With this notation, the formula for calculating present value (PV) is

$$PV(x, n, r) = \frac{x}{(1 + r)^n}$$

The denominator in this formula is a number greater than 1. Thus, dividing x by $(1 + r)^n$ will give a present value that is less than x . For this reason, we often say that we "discount" x back to the present. You can see that if you had the discounted amount now and could invest it at the interest rate r , then after n time periods (days, months, years, and so on) the value of the investment would be the discounted amount times $(1 + r)^n$, which is simply x .

Keeping the interest rate consistent with the time periods is important. For example, a savings account may pay 10% "compounded monthly." Thus, a year is really 12 time periods, and so $n = 12$. The monthly interest rate is $10\%/12$, or 0.8333%. Thus, the value of \$100 deposited in the account and left for a year would be $\$100 \times (1.00833)^{12} = \110.47 . Notice that compounding helps because the interest itself earns interest during each time period. Thus, if you have a choice among savings accounts that have the same interest rate, the one that compounds more frequently will end up having a higher eventual payoff.

We can now talk about the present value of a stream of cash flows. Suppose that a friend is involved in a business deal and offers to let you in on it. For \$425 paid to him now, he says, you can have \$110.00 next year, \$121.00 the following year, \$133.10 the third year, and \$146.41 at the end of Year 4. This is a great deal, he says, because your payments will total \$510.51.

What is the present value of the stream of payments? (You probably can guess already!) Let us suppose you put your money into a savings account at 10%, compounded annually. Then we would calculate the present value of the stream of cash flows as the sum of the present values of the individual cash flows:

$$\begin{aligned} \text{PV} &= \frac{110.00}{1.1} + \frac{121.00}{(1.1)^2} + \frac{133.10}{(1.1)^3} + \frac{146.41}{(1.1)^4} \\ &= \$100 + \$100 + \$100 + \$100 = \$400 \end{aligned}$$

Thus, the deal is not so great. You would be paying \$425 for a stream of cash flows that has a present value of only \$400. The *net present value* (NPV) of the cash flows is the present value of the cash flows (\$400) minus the cost of the deal (\$425), or -\$25; you would be better off keeping your \$425 and investing it in the savings account.

The formula for calculating NPV for a stream of cash flows x_0, \dots, x_n over n periods at interest rate r is

$$\begin{aligned} \text{NPV} &= \frac{x_0}{(1+r)^0} + \frac{x_1}{(1+r)^1} + \dots + \frac{x_n}{(1+r)^n} \\ &= \sum_{i=0}^n \frac{x_i}{(1+r)^i} \end{aligned}$$

In general, we can have both outflows (negative numbers) and inflows. In the example, we could include the cash outflow of \$425 as a negative number in calculating NPV:

$$\begin{aligned} \text{NPV} &= \frac{-425.00}{(1.1)^0} + \frac{110.00}{(1.1)^1} + \frac{121.00}{(1.1)^2} + \frac{133.10}{(1.1)^3} + \frac{146.41}{(1.1)^4} \\ &= -\$425 + \$400 \\ &= -\$25 \end{aligned}$$

[Recall that raising any number to the zero power is equal to 1, and so $(1.1)^0 = 1$.] Clearly, we could deal with any stream of cash flows. There could be one big inflow and then a bunch of outflows (such as with a loan), or there could be a large outflow (buying a machine), then some inflows (revenue), another outflow (maintenance costs), and so on. When NPV is calculated, it reveals the value of the stream of cash flows. A negative NPV for a project indicates that the money would be better invested to earn interest rate r .

We began our discussion by talking about trade-offs. You can see how calculating present values establishes trade-offs between dollars at one point in time and dollars at another. That is, you would be indifferent between receiving \$1 now or $\$1(1+r)$ at the end of the next time period. More generally, \$1 now is worth $\$1(1+r)^n$ at the end of n time periods. NPV works by using these trade-off rates to discount all the cash flows back to the present.

Knowing the interest rate is the key in using present-value analysis. What is the appropriate interest rate? In general, it is the interest rate that you could get for investing your money in the next best opportunity. Often we use the interest rate from a savings account, a certificate of deposit, or short-term (money market) securities. For a corporation, the appropriate interest rate to use might be the interest rate they would have to pay in order to raise money by issuing bonds. Often the interest rate is called the *hurdle rate*, indicating that an acceptable investment must earn more than this rate.

We have talked about the elements of decision problems: objectives, decisions to make, uncertain events, and consequences. The discussion of the time value of money showed how a consequence that is a stream of cash flows can be valued through the trade-offs implicit in interest rates. Now it is time to put all of this together and try it out in an example. Imagine the problems that an oil company might face in putting together a plan for dealing with a major oil spill. Here are managers in the fictitious "Larkin Oil" struggling with this situation.

LARKIN OIL

Pat Mills was restless. The Oil Spill Contingency Plan Committee was supposed to come up with a concrete proposal for the top management of Larkin Oil, Inc. The committee had lots of time; the CEO had asked for recommendations within three months. This was their first meeting.

Over the past hour, Sandy Wilton and Marty Kelso had argued about exactly what level of resources should be committed to planning for a major oil spill in the company's main shipping terminal bay.

"Look," said Sandy, "We've been over this so many times. When, and if, an oil spill actually occurs, we will have to move fast to clean up the oil. To do that, we have to have equipment ready to go."

"But having equipment on standby like that means tying up a lot of capital," Chris Brown replied. As a member of the financial staff, Chris was sensitive to committing capital for equipment that would be idle all the time and might actually have to be replaced before it was ever used. "We'd be better off keeping extensive records, maybe just a long list of equipment that would be useful in a major cleanup. We need to know where it is, what it's capable of, what its condition is, and how to transport it."

"Come to think of it, our list will also have to include information on transportation equipment and strategies," Leslie Taylor added.

Pat finally stirred. "You know what bothers me? We're talking about these alternatives, and the fact that we need to do thus and so in order to accomplish such and such. We're getting the cart before the horse. We just don't have our hands on the problem yet. I say we go back to basics. First, how could an oil spill happen?"

"Easy," said Sandy. "Most likely something would happen at the pipeline terminal. Something goes wrong with a coupling, or someone just doesn't pay attention while loading oil on the ship. The other possibility is that a tanker's hull fails for some reason, probably from running aground because of weather."

"Weather may not be the problem," suggested Leslie. "What about incompetence? What if the pilot gets drunk?"

Marty Kelso always was able to imagine the unusual scenarios. "And what about the possibility of sabotage? What if a terrorist decides to wreak environmental havoc?"

"OK," said Pat, "In terms of the actual cleanup, the more likely terminal spill would require a different kind of response than the less likely event of a hull failure. In planning for a terminal accident, we need to think about having some equipment at the terminal. Given the higher probability of such an accident, we should probably spend some money on cleanup equipment that would be right there and available."

"I suppose so," conceded Chris. "At least we would be spending our money on the right kind of thing."

"You know, there's another problem that we're not really thinking about," Leslie offered. "An oil spill at the terminal can be easily contained with relatively little environmental damage. On the other hand, if we ever have a hull failure, we have to act fast. If we don't, and mind you, we may not be able to because of the weather, Larkin Oil will have a terrible time trying to clean up the public relations as well as the beaches. And think about the difference in the PR problem if the spill is due to incompetence on the part of a pilot rather than weather or sabotage."

"Even if we act fast, a huge spill could still be nearly impossible to contain," Pat pointed out. "So what's the upshot? Sounds to me like we need someone who could make a decision immediately about how to respond. We need to recover as much oil as possible, minimize environmental damage, and manage the public relations problem."

"And do this all efficiently," growled Chris Brown. "We still have to do it without having tied up all of the company's assets for years waiting for something to happen."

The committee at Larkin Oil has a huge problem on its hands. The effects of its work now and the policy that is eventually implemented for coping with future accidents will substantially affect the company resources and possibly the environment. We cannot solve the problem entirely, but we can apply the principles discussed so far in the chapter. Let us look at the basic elements of the decision situation.

First, what is the committee's decision context, and what are Larkin's objectives? The context is making recommendations regarding plans for possible future oil spills, and the immediate decision is what policy to adopt for dealing with oil spills. Exactly what alternatives are available is not clear. The company's objectives are well stated by Pat Mills and Chris Brown at the end of the example: (1) recover as much oil as possible, (2) minimize environmental damage, (3) minimize damage to Larkin's public image, and (4) minimize cost. Recovering as much oil as possible is perhaps best viewed as a means to minimize environmental damage as well as the impact on Larkin's image. It also appears that a fundamental issue is how much of the company's resources should be committed to standby status waiting for an accident to occur. In general, the more resources committed, the faster the company could respond and the less damage would be done. Having these objectives out on the table immediately and understanding the inherent trade-offs will help the committee organize their efforts as they explore potential policy recommendations.

Is this a sequential decision problem? Based on Pat's last statement, the immediate decision must anticipate future decisions about responses to specific accident

situations. Thus, in figuring out an appropriate policy to adopt now, they must think about possible appropriate future decisions and what resources must be available at the time so that the appropriate action can be taken.

The scenario is essentially about uncertain events. Of course, the main uncertain event is whether an oil spill will ever occur. From Chris Brown's point of view, an important issue might be how long the cleanup equipment sits idle, requiring periodic maintenance, until an accident occurs. Also important are events such as the kind of spill, the location, the weather, the cause, and the extent of the damage. At the present time, imagining the first accident, all of these are unknowns, but if and when a decision must be made, some information will be available (location, current weather, cause), while other factors — weather conditions for the cleanup, extent of the eventual damage, and total cleanup cost — probably will not be known.

What is an appropriate planning horizon for Larkin? No indication is given in the case, but the committee members may want to consider this. How far into the future should they look? How long will their policy recommendations be active? They may wish to specify that at some future date (say three years from the present) another committee be charged with reviewing and updating the policy in light of scientific and technological advances.

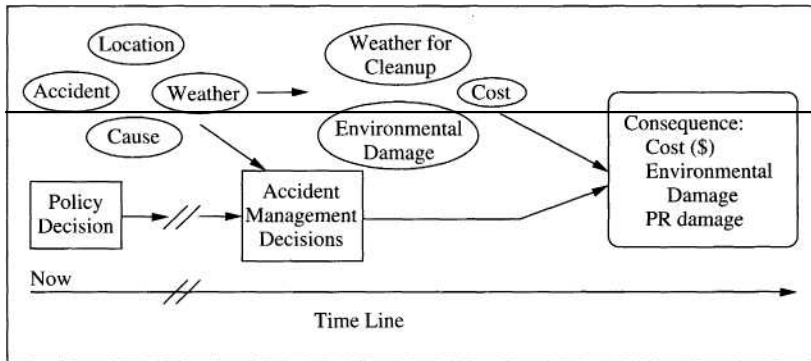
The problem also involves fundamental issues about how the different consequences are valued. As indicated, the fundamental trade-off is whether to save money by committing fewer resources or to provide better protection against future possible accidents. In other words, just how much is insurance against damage worth to Larkin Oil? In talking about consequences, the committee can imagine some possible ones and the overall "cost" (in generic terms) to the company: (1) committing substantial resources and never needing them; (2) committing a lot of resources and using them effectively to contain a major spill; (3) committing few resources and never needing them (the best possible outcome); and (4) committing few resources and not being able to clean up a spill effectively (the worst possible outcome).

Just considering the dollars spent, there is a time-value-of-money problem that Chris Brown eventually will want the committee to address. To some extent, dollars can be spent for protection now instead of later on. Alternative financing schemes can be considered to pay for the equipment required. Different strategies for acquiring and maintaining equipment may have different streams of cash flows. Calculating the present value of these different strategies for providing protection may be an important aspect of the decision.

Finally, the committee members also need to think about exactly how to allocate resources in terms of the other objectives stated by Pat Mills. They need to recover oil, minimize environmental damage, and handle public relations problems. Of course, recovering oil and minimizing environmental damage are linked to some extent. Overall, though, the more resources committed to one of these objectives, the less available they are to satisfy the others. The committee may want to specify some guidelines for resource allocation in its recommendations, but for the most part this allocation will be made at the time of future decisions that are in turn made in response to specific accidents.

Can we put all of this together? Figure 2.5 shows the sequence of decisions and uncertain events. This is only a rough picture, intended to capture the elements discussed

Figure 2.5
A graphical representation of Larkin Oil's situation.



here, a first step toward the development of a requisite decision model. Different decision makers most likely would have different representations of the situation, although most would probably agree on the essential elements of the values, decisions, uncertain events, and consequences.

SUMMARY

Hard decisions often have many different aspects. The basic elements of decision situations include values and objectives, decisions to be made, uncertain events, and consequences. This chapter discussed identification of the immediate decision at hand as well as subsequent decisions. We found that uncertain future events must be dovetailed with the sequence of decisions, showing exactly what is known before each decision is made and what uncertainties still remain. We discussed valuing consequences in some depth, emphasizing the specification of a planning horizon and the identification of relevant trade-offs. The discussion about the time value of money showed how interest rates imply a special kind of trade-off between cash flows at different points in time. Finally, the Larkin Oil example served to illustrate the identification of the basic elements of a major (and messy) decision problem.

QUESTIONS AND PROBLEMS

- 2.1 Suppose you are in the market for a new car, the primary use for which would be commuting to work, shopping, running errands, and visiting friends. **a** What are your objectives in this situation? What are some different alternatives? **b** Suppose you broaden the decision context. Instead of deciding on a car for commuting purposes, you are interested in having transportation for getting around your community. In this new decision context, how would you describe your objectives? What are some alternatives that you might not have considered in the narrower decision context?

- c How might you broaden the decision context further? (There are many ways to do this!) In this broader context, what new objectives must you consider? What new alternatives are available?
 - d Does your planning horizon change when you broaden the decision context in question b? Question c?
- 2.2 Explain in your own words why it is important in some situations to consider future decisions as well as the immediate decision at hand. Can you give an example from your own experience of an occasion in which you had to make a decision while explicitly anticipating a subsequent decision? How did the immediate decision affect the subsequent one?
 - 2.3 Sometimes broadening the decision context can change the planning horizon. For example, many companies face specific technical problems. Framed in a narrow decision context, the question is how to solve the specific problem, and a reasonable solution may be to hire a consultant. On the other hand, if the decision context is broadened to include solving related problems as well as the current one, the company might want to develop in-house expertise by hiring one or more permanent employees or training an existing employee in the required skills. What is the planning horizon in each case, and why does it change with the broader context? What objectives must be considered in the broader context that can be ignored in the narrower one?
 - 2.4 Explain in your own words why it is important to keep track of what information is known and what events are still uncertain for each decision.
 - 2.5 What alternatives other than specific protection strategies might Larkin Oil consider (for example, insurance)?
 - 2.6 Imagine the difficulties of an employer whose decision context is choosing a new employee from a set of applicants whom he will interview. What do you think the employer's objectives should be? Identify the employer's specific decisions to make and uncertainties, and describe the relevant uncertain events. How does the problem change if the employer has to decide whether to make an offer on the spot after each interview?
 - 2.7 Identify the basic elements of a real-estate investor's decision situation. What are the investor's objectives? Is the situation dynamic (that is, are there sequential decisions)? What are some of the uncertainties that the investor faces? What are the crucial trade offs? What role does the time value of money play for this investor?
 - 2.8 Describe a decision problem that you have faced recently (or with which you are currently struggling). Describe the decision context and your objectives. What were the specific decisions that you faced, and what were the relevant uncertainties? Describe the possible consequences.
 - 2.9 Calculate the net present value of a business deal that costs \$2500 today and will return \$1500 at the end of this year and \$1700 at the end of the following year. Use an interest rate of 13%.
 - 2.10 Find the net present value of a project that has cash flows of - \$ 12,000 in Year 1, +\$5000 in Years 2 and 3, -\$2000 in Year 4, and +\$6000 in Years 5 and 6. Use an interest rate of 12%. Find the interest rate that gives a net present value of zero.
 - 2.11 A friend asks you for a loan of \$1000 and offers to pay you back at the rate of \$90 per month for 12 months.

- a** Using an annual interest rate of 10%, find the net present value (to you) of loaning your friend the money. Repeat, using an interest rate of 20%.
- b** Find an interest rate that gives a net present value of 0. The interest rate for which $NPV = 0$ is often called the *internal rate of return*.
- 2.12** Terry Martinez is considering taking out a loan to purchase a desk. The furniture store manager rarely finances purchases, but will for Terry "as a special favor." The rate will be 10% per year, and because the desk costs \$600, the interest will come to \$60 for a one-year loan. Thus, the total price is \$660, and Terry can pay it off in 12 installments of \$55 each.
- a** Use the interest rate of 10% per year to calculate the net present value of the loan. (Remember to convert to a monthly interest rate.) Based on this interest rate, should Terry accept the terms of the loan? **b** Look at this problem from the store manager's perspective. Using the interest rate of 10%, what is the net present value of the loan to the manager? **c** What is the net present value of the loan to the manager if an interest rate of 18% is used? What does this imply for the real rate of interest that Terry is being charged for the loan?
- This kind of financing arrangement was widely practiced at one time, and you can see why from your answers to (c). By law, lenders in the United States now must clearly state the actual annual percentage rate in the loan contract.
- 2.13** Lynn Rasmussen is deciding what sports car to purchase. In reflecting about the situation, it becomes obvious that after a few years Lynn may elect to trade in the sports car for a new one, although the circumstances that might lead to this choice are uncertain. Should trading in the car count as an uncertain event or a future decision? What are the implications for building a requisite model of the current car-purchase decision if trading in the car is treated as an uncertain event? As a decision?

CASE STUDIES

THE VALUE OF PATIENCE

Robin Briggs, a wealthy private investor, had been approached by Union Finance Company on the previous day. It seemed that Union Finance was interested in loaning money to one of its larger clients, but the client's demands were such that Union could not manage the whole thing. Specifically, the client wanted to obtain a loan for \$385,000, offering to repay Union Finance \$100,000 per year over seven years.

Union Finance made Briggs the following proposition. Since it was bringing Briggs business, its directors argued, they felt that it was only fair for Briggs to put up a proportionately larger share of the money. If Briggs would put up 60% of the money (\$231,000), then Union would put up the remaining 40% (\$154,000). They

would split the payments evenly, each getting \$50,000 at the end of each year for the next seven years.

Questions

- 1 Union Finance can usually earn 18% on its money. Using this interest rate, what is the net present value of the client's offer to Union?
- 2 Robin Briggs does not have access to the same investments as Union. In fact, the best available alternative is to invest in a security earning 10% over the next seven years. Using this interest rate, what is Briggs's net present value of the offer made by Union? Should Briggs accept the offer?
- 3 What is the net present value of the deal to Union if Briggs participates as proposed?
- 4 The title of this case study is "The Value of Patience." Which of these two investors is more patient? Why? How is this difference exploited by them in coming to an agreement?

EARLY BIRD, INC.

The directors of Early Bird, Inc., were considering whether to begin a sales promotion for their line of specialty coffees earlier than originally planned. "I think we should go ahead with the price cuts," Tracy Brandon said. "After all, it couldn't hurt! At the very worst, we'll sell some coffee cheap for a little longer than we had planned, and on the other side we could beat New Morning to the punch."

"That's really the question, isn't it?" replied Jack Santorini. "If New Morning really is planning their own promotion, and we start our promotion now, we would beat them to the punch. On the other hand, we might provoke a price war. And you know what a price war with that company means. We spend a lot of money fighting with each other. There's no real winner. We both just end up with less profit."

Janice Wheeler, the finance VP for Early Bird, piped up, "The consumer wins in a price war. They get to buy things cheaper for a while. We ought to be able to make something out of that."

Ira Press, CEO for Early Bird, looked at the VP thoughtfully. "You've shown good horse sense in situations like these, Janice. How do you see it?"

Janice hesitated. She didn't like being put on the spot like this. "You all know what the projections are for the six-week promotion as planned. The marketing group tells us to expect sales of 10 million dollars. The objective is to gain at least two percentage points of market share, but our actual gain could be anywhere from nothing to three points. Profits during the promotion are expected to be down by 10 percent, but after the promotion ends, our increased market share should result in more sales and more profits."

Tracy broke in. "That's assuming New Morning doesn't come back with their own promotion in reaction to ours. And you know what our report is from Pete. He says that he figures New Morning is up to something."

"Yes, Pete did say that. But you have to remember that Pete works for our advertising agent. His incentive is to sell advertising. And if he thinks he can talk us into spending more money, he will. Furthermore, you know, he isn't always right. Last time he told us that New Morning was going to start a major campaign, he had the dates right, but it was for a different product line altogether."

Ira wouldn't let Janice off the hook. "But Janice, if New Morning does react to our promotion, would we be better off starting it early?"

Janice thought for a bit. If she were working at New Morning and saw an unexpected promotion begin, how would she react? Would she want to cut prices to match the competition? Would she try to stick with the original plans? Finally she said, "Look, we have to believe that New Morning also has some horse sense. They would not want to get involved in a price war if they could avoid it. At the same time, they aren't going to let us walk away with the market. I think that if we move early, there's about a 30 percent chance that they will react immediately, and we'll be in a price war before we know it."

"We don't have to react to their reaction, you know," replied Ira.

"You mean," asked Jack, "we have another meeting like this to decide what to do if they do react?"

"Right."

"So," Janice said, "I guess our immediate options are to start our promotion early or to start it later as planned. If we start it now, we risk a strong reaction from New Morning. If they do react, then we can decide at that point whether we want to cut our prices further."

Jack spoke up. "But if New Morning reacts strongly and we don't, we would probably end up just spending our money for nothing. We would gain no market share at all. We might even lose some market share. If we were to cut prices further, it might hurt profits, but at least we would be able to preserve what market share gains we had made before New Morning's initial reaction."

At this point, several people began to argue among themselves. Sensing that no resolution was immediately forthcoming, Ira adjourned the meeting, asking everyone to sleep on the problem and to call him with any suggestions or insights they had.

Questions

- 1 Based on the information in the case, what are Early Bird's objectives in this situation? Are there any other objectives that you think they should consider?
- 2 Given your answer to the previous question, what do you think Early Bird's planning horizon should be?
- 3 Identify the basic elements (values, decisions, uncertain events, consequences) of Early Bird's decision problem.
- 4 Construct a diagram like Figure 2.5 showing these elements.

REFERENCES

Identifying the elements of decision situations is implicit in a decision-analysis approach, although most textbooks do not explicitly discuss this initial step in decision modeling. The references listed at the end of Chapter 1 are all appropriate for discussions of values, objectives, decisions, uncertain events, and consequences.

The idea of understanding one's values as a prerequisite for good decision making is Ralph Keeney's thesis in his book *Value-Focused Thinking* (1992). A good summary is Keeney (1994). In the conventional approach, espoused by most authors on decision analysis, one finds oneself in a situation that demands a decision, identifies available alternatives, evaluates those alternatives, and chooses the best of those alternatives. Keeney argues persuasively that keeping one's values clearly in mind provides the ability to proactively find new decision opportunities and creative alternatives. Of course, the first step, and sometimes a difficult one, is understanding one's values, which we will explore in depth in Chapter 3.

Dynamic decision situations can be very complicated, and many articles and books have been written on the topic. A basic-level textbook that includes dynamic decision analysis is Buchanan (1982). DeGroot (1970) covers many dynamic decision problems at a somewhat more sophisticated level. Murphy et al. (1985) discuss the orchardist's dynamic decision problem in detail.

The time value of money is a standard topic in finance courses, and more complete discussions of net present value, internal rate of return (the implied interest rate in a sequence of cash flows), and related topics can be found in most basic financial management textbooks. Two good ones are Brigham (1985) and Schall and Haley (1986).

Brigham, E. F. (1985) *Financial Management: Theory and Practice*, 4th ed. Hinsdale, IL: Dryden.

Buchanan, J. T. (1982) *Discrete and Dynamic Decision Analysis*. New York: Wiley.

DeGroot, M. H. (1970) *Optimal Statistical Decisions*. New York: McGraw-Hill.

Keeney, R. L. (1992) *Value-Focused Thinking*. Cambridge, MA: Harvard University Press.

Keeney, R. L. (1994) "Creativity in Decision Making with Value-Focused Thinking." *Sloan Management Review*, Summer, 33-41.

Murphy, A. H., R. W. Katz, R. L. Winkler, and W.-R. Hsu (1985) "Repetitive Decision Making and the Value of Forecasts in the Cost-Loss Ratio Situation: A Dynamic Model." *Monthly Weather Review*, 113, 801-813.

Schall, L. D., and C. W. Haley (1986) *Introduction to Financial Management*, 4th ed. New York: McGraw-Hill.

EPILOGUE

On March 24, 1989, the Exxon *Valdez* tanker ran aground on a reef in Prince William Sound after leaving the Valdez, Alaska, pipeline terminal. Over 11 million gallons of oil spilled into Prince William Sound, the largest spill in the United States. In the aftermath, it was revealed that Aleyeska, the consortium of oil companies responsible for constructing and managing the pipeline, had instituted an oil spill contingency plan that

was inadequate to the task of cleaning up a spill of such magnitude. As a result of the inadequate plan and the adverse weather immediately after the spill, little oil was recovered. Hundreds of miles of environmentally delicate shoreline were contaminated. Major fisheries were damaged, leading to specific economic harm to individuals who relied on fishing for a livelihood. In addition, the spill proved an embarrassment for all of the major oil companies and sparked new interest in environmental issues, especially upcoming leases for offshore oil drilling. Even though the risk of a major oil spill was very small, in retrospect one might conclude that the oil companies would have been better off with a much more carefully thought out contingency plan and more resources invested in it. {Source: "Dead Otters and Silent Ducks," *Newsweek*, April 24, 1989, p. 70.)

Structuring Decisions

Having identified the elements of a decision problem, how should one begin the modeling process? Creating a decision model requires three fundamental steps. First is identifying and structuring the values and objectives. Structuring values requires identifying those issues that matter to the decision maker, as discussed in Chapter 2. Simply listing objectives, however, is not enough; we also must separate the values into fundamental objectives and means objectives, and we must specify ways to measure accomplishment of the objectives.

The second step is structuring the elements of the decision situation into a logical framework. To do this we have two tools: influence diagrams and decision trees. These two approaches have different advantages for modeling difficult decisions. Both approaches are valuable and, in fact, complement one another nicely. Used in conjunction with a carefully developed value structure, we have a complete model of the decision that shows all of the decision elements: relevant objectives, decisions to make, uncertainties, and consequences.

The final step is the refinement and precise definition of all of the elements of the decision model. For example, we must be absolutely clear on the precise decisions that are to be made and the available alternatives, exactly what the uncertain events are, and how to measure the consequences in terms of the objectives that have been specified. Although many consequences are easily measured on a natural scale (for example, NPV can be measured in dollars), nonquantitative objectives such as increasing health or minimizing environmental impact are more problematic. We will discuss ways to create formal scales to measure achievement of such objectives.

Structuring Values

Our first step is to structure values. In Chapter 2 we discussed the notion of objectives. In many cases, a single objective drives the decision; a manager might want to maximize profits next year, say, or an investor might want to maximize the financial return of an investment portfolio. Often, though, there are multiple objectives that conflict; for example, the manager might want to maximize profits but at the same time minimize the chance of losing money. The investor might want to maximize the portfolio's return but minimize the volatility of the portfolio's value.

If a decision involves a single objective, that objective is often easily identified. Careful thought may be required, however, to define the objective in just the right way. For example, you might want to calculate NPV over three years, using a particular interest rate. The discussion of value structuring that follows can help in the identification and clarification of the objective in a single-objective decision situation.

Even though many pages in this book are devoted to the analysis of single-objective decisions, for many decisions the real problem lies in balancing multiple conflicting objectives. The first step in dealing with such a situation is to understand just what the objectives are. Specifying objectives is not always a simple matter, as we will see in the following example.

Suppose you are an employer with an opening for a summer intern in your marketing department. Under the supervision of a senior employee, the intern would assist in the development of a market survey relating to a line of your company's consumer products.

HIRING A SUMMER INTERN

Many businesses hire students for short-term assignments. Such jobs often are called *internships*, and the employee — or intern — gets a chance to see what a particular kind of job and a specific company are like. Likewise, the company gets to try out a new employee without making a long-term commitment.

In this example, the fictional PeachTree Consumer Products has an opening for a summer intern. Working under the supervision of a senior employee in the marketing group, the intern would focus primarily on the development of a market survey for certain of the company's products. The problem is how to find an appropriate individual to fill this slot. Where should the company go to locate good candidates, which ones should be interviewed, and on the basis of what criteria should a particular candidate be chosen?

Imagine that you are the manager charged with finding an appropriate intern for PeachTree. Your first step is to create a long list of all the things that matter to you in this decision context. What objectives would you want to accomplish in filling this position? Certainly you would want the market survey to be done well. You might

want to use the summer as a trial period for the intern, with an eye toward a permanent job for the individual if the internship worked out. You might want to establish or cement a relationship with a college or university placement service. Table 3.1 shows a list of objectives (in no special order) that an employer might write down.

How would you go about generating a list like Table 3.1? Keeney (1994) gives some ideas. For example, think about some possible alternatives and ask what is good or bad about them. Or think about what you would like if you could have anything. Table 3.2 gives eight suggestions for generating your list of objectives.

Table 3.1
Objectives for hiring
summer intern.

Maximize quality of market survey.
Sell more consumer products.
Build market share.
Identify new market niches for company's products.
Minimize cost of survey design.
Try out prospective permanent employee.
Establish relationship with local college.
Provide assistance for senior employee.
Free up an employee to be trained for new assignment
Learn updated techniques from intern:
 Self
 Supervisor
 Market research department
 Entire company
Expose intern to real-world business experience.
Maximize profit.
Improve company's working environment by bringing in new and youthful energy.
Provide financial assistance for college student

Table 3.2
Techniques for identifying objectives.

1. **Develop a wish list.** What do you want? What do you value? What should you want?
2. **Identify alternatives.** What is a perfect alternative, a terrible alternative, some reasonable alternative? What is good or bad about each? «
3. **Consider problems and shortcomings.** What is wrong or right with your organization? What needs fixing?
4. **Predict consequences.** What has occurred that was good or bad? What might occur that you care about?
5. **Identify goals, constraints, and guidelines*** What are your aspirations? What limitations are placed on you?
6. **Consider different perspectives.** What would your competitor or your constituency be concerned about? At some time in the future, what would concern you?
7. **Determine strategic objectives*** What are your ultimate objectives? What are your values that are absolutely fundamental?
8. **Determine generic objectives.** What objectives do you have for your customers, your employees, your shareholders, yourself? What environmental, social, economic, or health and safety objectives are important?

Source: Keeney, R. L. (1994) "Creativity in Decision Making with Value-Focused Thinking," *Sloan Management Review*, Summer, 33-41.

Once you have a list of objectives, what do you do? Structuring the objectives means organizing them so that they describe in detail what you want to achieve and can be incorporated in an appropriate way into your decision model. We start by separating the list into items that pertain to different kinds of objectives. In the summer-intern example, objectives can be sorted into several categories:

- Business performance (sell more products, maximize profit, increase market share, identify market niches)
- Improve the work environment (bring in new energy, assist senior employee)
- Improve the quality and efficiency of marketing activities (maximize survey quality, minimize survey cost)
- Personnel and corporate development (learn updated techniques, free up employee for new assignment, try out prospective employee)
- Community service (financial aid, expose intern to real world, relationship with local college)

Of course, there are other ways to organize these objectives; the idea is to create categories that reflect the company's overall objectives.

Before continuing with the value structuring, we must make sure that the objectives are appropriate for the decision context. Recall that the decision context is hiring a summer intern for the marketing department. This is a relatively narrow context for which some of the listed objectives are not especially relevant. For example, selling more consumer products and maximizing profit, although indeed important objectives, are too broad to be essential in the current decision context. Although hiring the best individual should have a positive impact on overall company performance, more crucial in the specific context of hiring the best intern are the objectives of enhancing marketing activities, personnel development, community service, and enhancing the work environment. These are the areas that hiring an intern may directly affect.

Fundamental and Means Objectives

With a set of objectives that is consistent with the decision context, the next step is to separate *means* from *fundamental* objectives. This is a critical step, because here we indicate those objectives that are important because they help achieve other objectives and those that are important simply because they reflect what we really want to accomplish. For example, working fewer hours may appear to be an important objective, but it may be important only because it would allow an individual to spend more time with his or her family or to pursue other activities that represent fundamental interests, things that are important simply because they are important. Thus, "minimize hours worked" is a means objective, whereas "maximize time with family" is a fundamental objective.

Fundamental objectives are organized into *hierarchies*. The upper levels in a hierarchy represent more general objectives, and the lower levels explain or describe important elements of the more general levels. For example, in the context of defining vehicle regulations, a higher-level fundamental objective might be "Maximize Safety," below which one might find "Minimize Loss of Life," "Minimize Serious Injuries," and "Minimize Minor Injuries." The three lower-level objectives are fundamental objectives that explain what is meant by the higher-level objective "Maximize Safety." The three lower-level objectives are also fundamental; each one describes a specific aspect of safety, and as such each one is inherently important. This hierarchy could be expanded by including another level. For example, we might include the objectives "Minimize Loss of Child Lives" and "Minimize Loss of Adult Lives" as aspects of the loss-of-life objective and similarly distinguish between serious injuries to children and adults. Figure 3.1 displays the hierarchy.

Means objectives, on the other hand, are organized into *networks*. In the vehicle-safety example, some means objectives might be "Minimize Accidents" and "Maximize Use of Vehicle-Safety Features." Both of these are important because they help to maximize safety. Beyond these two means objectives might be other means objectives such as "Maximize Driving Quality," "Maintain Vehicles Properly," and "Maximize Purchase of Safety Features on Vehicles." Figure 3.2 shows a means-objectives network that includes still more means objectives. A key difference between this network and the fundamental-objectives hierarchy in Figure 3.1 is that means objectives can be connected to several objectives, indicating that they help accomplish these objectives. For example, "Have Reasonable Traffic Laws" affects both "Maximize Driving Quality" and "Maintain Vehicles Properly."

Structuring the fundamental-objectives hierarchy is crucial for developing a multiple-objective decision model. As we will see, the lowest-level fundamental objectives will be the basis on which various consequences will be measured. Distinguishing means and fundamental objectives is important at this stage of the game primarily so that the decision maker is certain that the appropriate objectives—fundamental, not means—are specified in the decision model. But the means network has other uses as well. We will see in the last portion of the chapter that an easily measured means objective can sometimes substitute for a fundamental objective that is

Figure 3.1
A fundamental-objectives hierarchy.
 Source: Keeney (1992, p. 70).


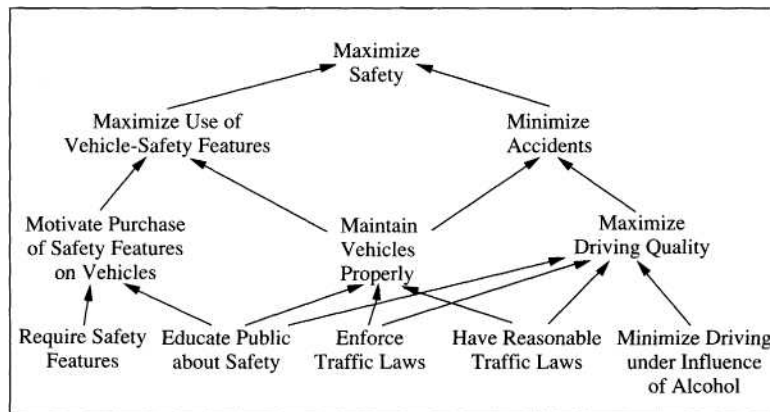


Figure 3.2
A means-objectives
network.
Source: Keeney,
(1992, p. 70).



more difficult to measure. And in Chapter 6 we will see how the means-objectives network provides an important basis for generating creative new alternatives.

How do we first separate means and fundamental objectives and then construct the fundamental-objectives hierarchy and the means-objectives network? A number of guiding questions are used to accomplish these tasks.

The first question to ask regarding any objective is, "*Why Is That Important?*" Known as the WITI test, this question does two things: distinguishes between means and fundamental objectives and reveals connections among the objectives. If the answer to the question is, "This objective is important because it helps accomplish X," then you know that the original objective is a means objective and that it has an impact on X. Moreover, a decision maker can continue by asking, "Why is X important?" By continuing to ask why the next objective is important, we can trace out the connections from one means objective to the next until we arrive at an objective for which the answer is, "This objective is important just because it is important. It is one of the fundamental reasons why I care about this decision." In this case, we have identified a fundamental objective.

As an example, look again at Figure 3.2. We might ask, for example, "Why is it important to maintain vehicles properly?" The answer is that doing so helps to minimize accidents and maximize the use of vehicle-safety features. Asking why minimizing accidents is important reveals that it helps maximize safety. The same is true if we ask why maximizing use of safety features is important. Finally, why is safety important? Maximizing safety is fundamentally important; it is what we care about in the context of establishing regulations regarding vehicle use. The answers to the questions trace out the connections among these four objectives and appropriately identify "Maximize Safety" as a fundamental objective.

The WITI test is useful for moving from means objectives to fundamental objectives. What about going the other way? The obvious question to ask is, "How can this objective be achieved?" For example, in the vehicle-regulation context we would ask, "How can we maximize safety?" The answer might give any of the upstream means

objectives that appear in Figure 3.2. Sequentially asking "How can this objective be achieved?" can help to identify means objectives and establish links among them.

What about constructing the fundamental-objectives hierarchy? Starting at the top of the hierarchy, the question to ask is, "What do you mean by that?" In our vehicle example, we would ask, "What does maximize safety mean?" The answer is that we mean minimizing lives lost, serious injuries, and minor injuries. In turn we could ask, "What do you mean by minimizing lives lost?" The answer might be minimizing child lives lost and adult lives lost; that is, it might be useful in this decision context to consider safety issues for children and adults separately, perhaps because different kinds of regulations would apply to these two groups.

Finally, we can work upward in the fundamental-objectives hierarchy, starting at a lower-level objective. Ask the question, "Of what more general objective is this an aspect?" For example, if we have identified saving adult lives as a fundamental objective — it is a fundamental reason we care about vehicle regulations — then we might ask, "Is there a more general objective of which saving adult lives is an aspect?" The answer would be the more general objective of saving lives, and asking the question again with respect to saving lives would lead us to the overall fundamental objective of maximizing safety.

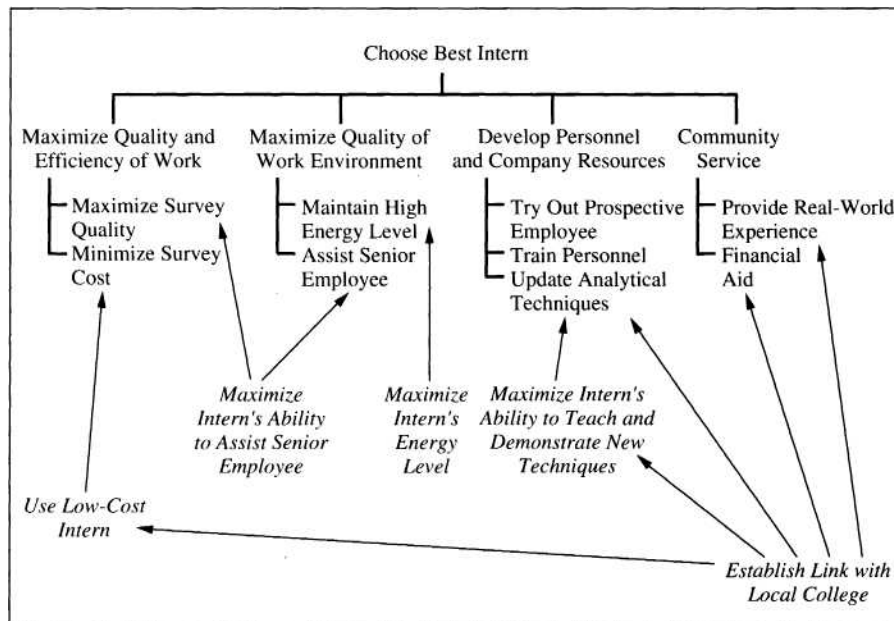
Figure 3.3 summarizes these four techniques for organizing means and fundamental objectives. It is important to realize that one might ask these questions in any order, mixing up the sequence, jumping from the means network to the fundamental-objectives hierarchy and back again. Be creative and relaxed in thinking about your values!

Let us look again at PeachTree's summer-intern decision. Figure 3.4 shows both a fundamental-objectives hierarchy and a means network with appropriate connections between them. The means objectives are shown in italics. Note that some objectives have been added, especially criteria for the intern, such as ability to work with the senior employee, ability to demonstrate new techniques to the staff, and a high level of energy. In the decision context, choosing the best intern for the summer

Figure 3.3
How to construct mean-objectives networks and fundamental-objectives hierarchies.

	Fundamental Objectives	Means Objectives
To Move: Ask:	<i>Downward in the Hierarchy:</i> "What do you mean by that?"	<i>Away from Fundamental Objectives:</i> "How could you achieve this?"
To Move: Ask:	<i>Upward in the Hierarchy:</i> "Of what more general objective is this an aspect?"	<i>Toward Fundamental Objectives:</i> "Why is that important?" (WITI)

Figure 3.4
 Fundamental and means objectives for PeachTree's summer-intern decision.



position, these criteria help define what "best" means in terms that relate directly to the company's fundamental objectives.

Insights can be gleaned from Figure 3.4. First, the means objectives give some guidance about what kind of intern to hire; up-to-date technical skills, good "people skills" for working with the senior employee, an ability (and willingness) to demonstrate new techniques for the firm, and a high energy level. In addition, establishing a link with the local college is a very important step. Although this is a means objective and hence not important in and of itself, it has an impact on many other objectives, both means and fundamental.

The fundamental-objectives hierarchy and the means-objectives network can provide a lot of insight even at this initial level. The fundamental objectives tell you why you care about the decision situation and what criteria you should be looking at in evaluating options. For the summer-intern situation, the company cares about the four main-level fundamental objectives, and the lower-level objectives provide more detail. Having sorted out the means objectives, we can rest assured that we will be able to evaluate candidates (and perhaps even develop a strategy for finding good candidates) whose qualities are consistent with the company's concerns. Finally, as we mentioned above, the means network can suggest creative new alternatives. For example, a great strategy would be to become acquainted with professors or career

counselors at the local college and to explain to them exactly what the company is looking for in a summer intern.

Getting the Decision Context Right

Recall that the context for PeachTree's decision has been to hire the best intern. What would happen if we were to broaden the context? Suppose we were to set the context as enhancing the company's marketing activities. First, we would want to consider far more options than just hiring an intern. The broader context also suggests looking for permanent new hires or training current employees in new methods. One of the results would be that the means objectives might change; some of the means objectives might broaden from optimizing characteristics of the intern to optimizing characteristics of the marketing group as a whole. "Maximize Intern's Energy Level" might become "maximize marketing group's energy level," which might suggest means objectives such as hiring new high-energy employees or sending employees to a workshop or retreat. You can see that as we broaden the decision context, the objectives change in character somewhat. The more the context is broadened, the greater the change. If we were to go all the way to a *strategic* — broadest possible — context of "maximize profit" or "build market share," for example, then many of the fundamental objectives in Figure 3.4 would become means objectives, and alternatives affecting all parts of the company would have to be considered.

At this point you may be wondering how you know when you have identified the appropriate decision context and its corresponding fundamental-objectives hierarchy and means network. As in Chapter 2, we can invoke the notion of a requisite model to ensure that all appropriate but no superfluous objectives have been included, given the decision context. The real question, though, is the decision context itself. How do you know how broad or narrow to make the context? This question is absolutely fundamental, and unfortunately there is no simple answer. As a decision maker, you must choose a context that fits three criteria. The first is straightforward; ask whether the context you have set really captures the situation at hand. Are you addressing the right problem? For example, searching for a job of the same type as your present one but with a different company is the wrong decision context if your real problem is that you do not enjoy the kind of work required in that job; you should broaden the context to consider different kinds of jobs, careers, or lifestyles. On the other hand, if you really love what you do but are dissatisfied with your current job for reasons related to that particular position or your firm (low salary, poor working conditions, conflicts with fellow workers, and so on), then looking for another similar job with another firm is just the right context.

The second criterion might be called *decision ownership*. Within organizations especially, the broader the decision context, the higher up the organizational ladder are the authority to make the decision and the responsibility for its consequences. Do

you have the authority to make decisions within the specified context (or will you be reporting the results of your analysis to someone with that authority)? If you conclude that you do not have this authority, then look for a narrower context that matches the authority you do have.

Feasibility is the final issue; in the specified context, will you be able to do the necessary study and analysis in the time allotted with available resources? Broader contexts often require more careful thought and more extensive analysis; addressing a broad decision context with inadequate time and resources can easily lead to dissatisfaction with the decision process (even though good consequences may result from lucky outcomes). It would be better in such a situation to narrow the context in some way until the task is manageable.

Like most aspects of decision analysis, setting the context and structuring objectives may not be a once-and-for-all matter. After initially specifying objectives, you may find yourself refining the context and modifying the objectives. Refining the context several times and iterating through the corresponding sets of objectives are not signs of poor decision making; instead, they indicate that the decision situation is being taken seriously, and that many different possibilities and perspectives are being considered.

Structuring Decisions: Influence Diagrams

With the fundamental objectives specified, structured, and sorted out from the means objectives, we can turn now to the process of structuring the various decision elements — decisions and alternatives, uncertain events and outcomes, and consequences. We begin with influence diagrams, which can provide simple graphical representations of decision situations. Different decision elements show up in the influence diagram as different shapes. These shapes are then linked with arrows in specific ways to show the relationships among the elements.

In an influence diagram, rectangles represent decisions, and ovals represent chance events. A rectangle with rounded corners is used to represent a mathematical calculation or a constant value; these rounded rectangles will have a variety of uses, but the most important is to represent consequences. The three shapes are generally referred to as *nodes*: decision nodes, chance nodes, and consequence or calculation nodes. Nodes are put together in a *graph*, connected by arrows, or *arcs*. We call a node at the beginning of an arc a *predecessor* and a node at the end of an arc a *successor*.

Consider a venture capitalist's situation in deciding whether to invest in a new business. For the moment, let us assume that the capitalist has only one objective in this context — to make money (not an unreasonable objective for a person in this line of work). The entrepreneur seeking the investment has impeccable qualifications and has generally done an excellent job of identifying the market, assembling a skilled management and production team, and constructing a suitable business plan. In fact,


it is clear that the entrepreneur will be able to obtain financial backing from some source whether the venture capitalist decides to invest or not. The only problem is that the proposed project is extremely risky — more so than most new ventures. Thus, the venture capitalist must decide whether to invest in this highly risky undertaking. If she invests, she may be able to get in on the ground floor of a very successful business. On the other hand, the operation may fail altogether. Clearly, the dilemma is whether the chance of getting in on the ground floor of something big is worth the risk of losing the investment entirely. If she does not invest in this project, she may leave her capital in the stock market or invest in other less risky ventures. Her investment situation appears as an influence diagram in Figure 3.5.

Note that both "Invest?" and "Venture Succeeds or Fails" are predecessors of the consequence "Return on Investment." The implication is that the consequence depends on both the decision and the chance event. In general, consequences depend on what happens or what is decided in the nodes that are predecessors of the consequence node. Moreover, as soon as the decision is made *and* the uncertain event is resolved, the consequence is determined; there is no uncertainty about the consequence at this point, and so the use of the rounded rectangle is appropriate. Note also that no arc points from the chance node to the decision node. The absence of an arc indicates that when the decision is made, the venture capitalist does not know whether the project will succeed. She may have some feeling for the chance of success, and this information would be included in the influence diagram as probabilities of possible levels of success or failure. Thus, the influence diagram as drawn captures the decision maker's current state of knowledge about the situation.

Also note that no arc points from the decision to the uncertain event. The absence of this arrow has an important and subtle meaning. The uncertainty node is about the success of the venture. The absence of the arc from "Invest?" to "Venture Succeeds or Fails" means that the venture's chances for success are not affected by the capitalist's decision. In other words, the capitalist need not concern herself with her impact on the venture.

It is possible to imagine situations in which the capitalist may consider different levels of investment as well as managerial involvement. For example, she may be willing to invest \$100,000 and leave the entrepreneur alone. But if she invests \$500,000, she may wish to be more active in running the company. If she believes her involvement would improve the company's chance of success, then it would be appropriate to include an arrow from the decision node to the chance

Figure 3.5
*Influence diagram of
venture capitalist's
decision.*



node; her investment decision — the level of investment and the concomitant level of involvement — would be relevant for determining the company's chance of success. In our simple and stylized example, however, we are assuming that her choice simply is whether to invest and that she has no impact on the company's chance of success.

Influence Diagrams and the Fundamental-Objectives Hierarchy

Suppose the venture capitalist actually has multiple objectives. For example, she might wish to focus on a particular industry, such as personal computers, obtaining satisfaction by participating in the growth of this industry. Thus, in addition to the objective of making money, she would have an objective of investing in the personal-computer industry.

Figure 3.6 shows a simple two-level objectives hierarchy and the corresponding influence diagram for the venture capitalist's decision. You can see in this figure how the objectives hierarchy is reflected in the pattern of consequence nodes in the influence diagram; two consequence nodes labeled "Invest in Computer Industry" and "Return on Investment" represent the lower-level objectives and in turn are connected to the "Overall Satisfaction" consequence node. This structure indicates that in some situations the venture capitalist may have to make a serious trade-off between these two objectives, especially when comparing a computer-oriented busi-

Figure 3.6
The venture capitalist's decision with two objectives.

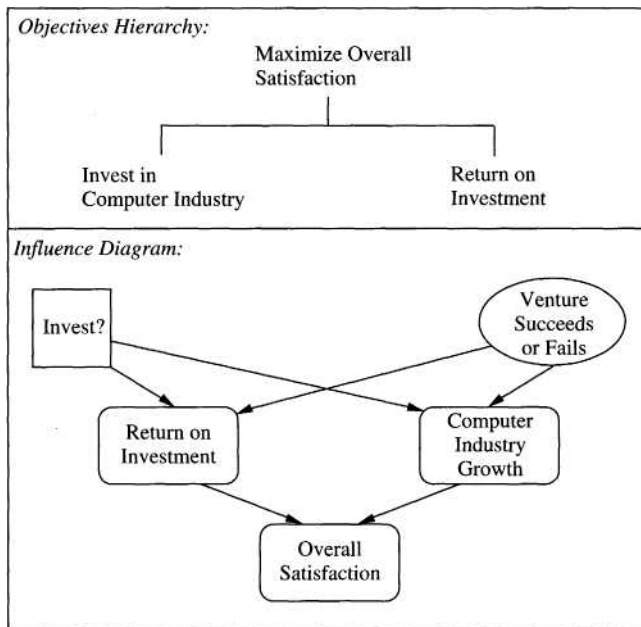
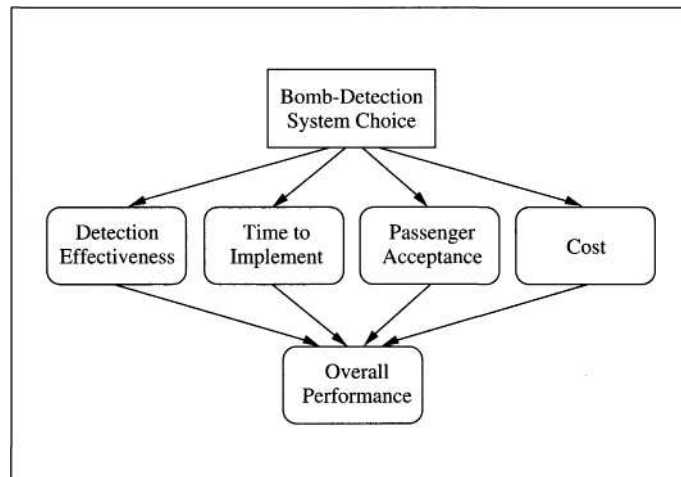


Figure 3.7
Multiple objectives in
selecting a bomb-
detection system.



ness startup with a noncomputer business that has more potential to make money. As before, the rounded rectangles for "Computer Industry" and "Return on Investment" are appropriate; after the decision is made and the venture's level of success is determined, these consequences are known. The rounded rectangle for "Overall Satisfaction" indicates that once the two individual consequences are known, then the overall consequence can be determined.

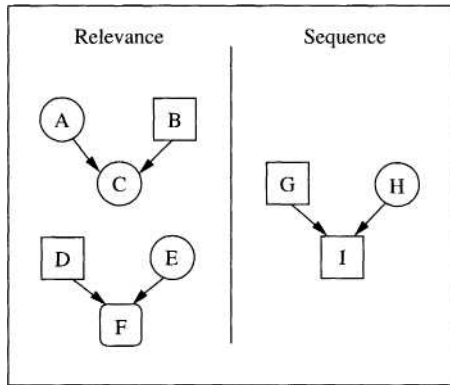
Figure 3.7 shows the influence diagram for another multiple-objective decision. In this situation, the Federal Aviation Administration (FAA) must choose from among a number of bomb-detection systems for commercial air carriers (Ulvila and Brown, 1982). In making the choice, the agency must try to accomplish several objectives. First, it would like the chosen system to be as effective as possible at detecting various types of explosives. The second objective is to implement the system as quickly as possible. The third is to maximize passenger acceptance, and the fourth is to minimize cost. To make the decision and solve the influence diagram, the FAA would have to score each candidate system on how well it accomplishes each objective. The measurements of time and cost would naturally be made in terms of days and dollars, respectively. Measuring detection effectiveness and passenger acceptance might require experiments or surveys and the development of an appropriate measuring device. The "Overall Performance" node would contain a formula that aggregates the individual scores, incorporating the appropriate trade-offs among the four objectives. Assessing the trade-off rates and constructing the formula to calculate the overall score is demonstrated in an example in Chapter 4 and is discussed thoroughly in Chapters 15 and 16.

Using Arcs to Represent Relationships

The rules for using arcs to represent relationships among the nodes are shown in Figure 3.8. In general, an arc can represent either *relevance* or *sequence*. The context of the arrow indicates the meaning. For example, an arrow pointing into a chance

Figure 3.8

Representing influence with arrows. Arrows into chance and consequence nodes represent relevance, and arrows into decision nodes represent sequence.



node designates relevance, indicating that the predecessor is relevant for assessing the chances associated with the uncertain event. In Figure 3.8 the arrow from Event A to Event C means that the chances (probabilities) associated with C may be different for different outcomes of A. Likewise, an arrow pointing from a decision node to a chance node means that the specific chosen decision alternative is relevant for assessing the chances associated with the succeeding uncertain event. For instance, the chance that a person will become a millionaire depends to some extent on the choice of a career. In Figure 3.8 the choice taken in Decision B is relevant for assessing the chances associated with Event C's possible outcomes.

Relevance arcs can also point into consequence or calculation nodes, indicating that the consequence or calculation depends on the specific outcome of the predecessor node. In Figure 3.8, consequence F depends on both Decision D and Event E. Relevance arcs in Figure 3.6 point into the "Computer Industry Growth" and "Return on Investment" nodes; the decision made and the success of the venture are relevant for determining these two consequences. Likewise, relevance arcs point from the two individual consequence nodes into the "Overall Satisfaction" node.

When the decision maker has a choice to make, that choice would normally be made on the basis of information available at the time. What information is available? Everything that happens before the decision is made. Arrows that point to decisions represent information available at the time of the decision and hence represent *sequence*. Such an arrow indicates that the decision is made knowing the outcome of the predecessor node. An arrow from a chance node to a decision means that, from the decision maker's point of view, all uncertainty associated with a chance event is resolved and the outcome known when the decision is made. Thus, information is available to the decision maker regarding the event's outcome. This is the case with Event H and Decision I in Figure 3.8; the decision maker is waiting to learn the outcome of H before making Decision I. An arrow from one decision to another decision simply means that the first decision is made before the second, such as Decisions G and I in Figure 3.8. Thus, the sequential ordering of decisions is shown in an influence diagram by the path of arcs through the decision nodes.

The nature of the arc — relevance or sequence — can be ascertained by the context of the arc within the diagram. To reduce the confusion of overabundant notation, all arcs have the same appearance in this book. For our purposes, the rule for determining the nature of the arcs is simple; an arc pointing to a decision represents sequence, and all others represent relevance.

Properly constructed influence diagrams have no *cycles*; regardless of the starting point, there is no path following the arrows that leads back to the starting point. For example, if there is an arrow from A to B, there is no path, however tortuous, that leads back to A from B. Imagine an insect traveling from node to node in the influence diagram, always following the direction of the arrows. In a diagram without cycles, once the insect leaves a particular node, it has no way to get back to that node.

Some Basic Influence Diagrams

In this section, several basic influence diagrams are described. Understanding exactly how these diagrams work will provide a basis for understanding more complex diagrams.

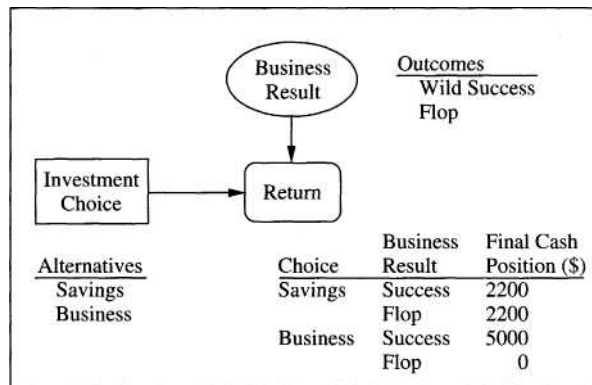
The Basic Risky Decision

This is the most elementary decision under uncertainty that a person can face. The venture-capital example above is a basic risky decision; there is one decision to make and one uncertain event.

Many decision situations can be reduced to a basic risky decision. For example, imagine that you have \$2000 to invest, with the objective of earning as high a return on your investment as possible. Two opportunities exist, investing in a friend's business or keeping the money in a savings account with a fixed interest rate. If you invest in the business, your return depends on the success of the business, which you figure could be wildly successful, earning you \$3000 beyond your initial investment (and hence leaving you with a total of \$5000), or a total flop, in which case you will lose all your money and have nothing. On the other hand, if you put the money into a savings account, you will earn \$200 in interest (leaving you with a total of \$2200) regardless of your friend's business.

The influence diagram for this problem is shown in Figure 3.9. This figure also graphically shows details underlying the decision, chance, and consequence nodes. The decision node includes the choice of investing in either the business or the savings account. The chance node represents the uncertainty associated with the business and shows the two possible outcomes. The consequence node includes information on the dollar return for different decisions (business investment versus savings) and the outcome of the chance event. This table shows clearly that if you invest in the business, your return depends on what the business does. However, if you put your money into savings, your return is the same regardless of what happens with the business.

Figure 3.9
*Basic risky decision
 with displayed
 choices, outcomes,
 and consequences.*



You can see that the essential question in the basic risky decision is whether the potential gain in the risky choice (the business investment) is worth the risk that must be taken. The decision maker must, of course, make the choice by comparing the risky and riskless alternatives. Variations of the basic risky choice exist. For example, instead of having just two possible outcomes for the chance event, the model could include a range of possible returns, a much more realistic scenario. The structure of the influence diagram for this range-of-risk dilemma, though, would look just the same as the influence diagram in Figure 3.9; the difference lies in the details of the chance event, which are not shown explicitly in the structure of the diagram.

Imperfect Information

Another basic kind of influence diagram reflects the possibility of obtaining imperfect information about some uncertain event that will affect the eventual payoff. This might be a forecast, an estimate or diagnosis from an acknowledged expert, or information from a computer model. In the investment example, you might subscribe to a service that publishes investment advice, although such services can never predict market conditions perfectly.

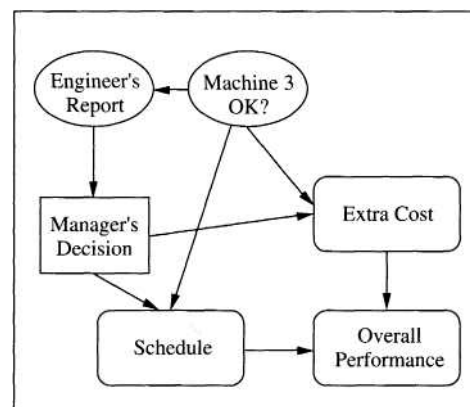
Imagine a manufacturing-plant manager who faces a string of defective products and must decide what action to take. The manager's fundamental objectives are to solve this problem with as little cost as possible and to avoid letting the production schedule slip. A maintenance engineer has been dispatched to do a preliminary inspection on Machine 3, which is suspected to be the source of the problem. The preliminary check will provide some indication as to whether Machine 3 truly is the culprit, but only a thorough and expensive series of tests — not possible at the moment — will reveal the truth. The manager has two alternatives. First, a replacement for Machine 3 is available and could be brought in at a certain cost. If Machine 3 is the problem, then work can proceed and the production schedule will not fall behind. If Machine 3 is not the source of the defects, the problem will still exist, and the workers will have to change to another product while the problem is

tracked down. Second, the workers could be changed immediately to the other product. This action would certainly cause the production schedule for the current product to fall behind but would avoid the risk (and cost) of unnecessarily replacing Machine 3.

Without the engineer's report, this problem would be another basic risky decision; the manager would have to decide whether to take the chance of replacing Machine 3 based on personal knowledge regarding the chance that Machine 3 is the source of the defective products. However, the manager is able to wait for the engineer's preliminary report before taking action. Figure 3.10 shows an influence diagram for the manager's decision problem, with the preliminary report shown as an example of imperfect information. The diagram shows that the consequences depend on the choice made (replace Machine 3 or change products) and whether Machine 3 actually turns out to be defective. There is no arrow from "Engineer's Report" to the consequence nodes because the report does not have a direct effect on the consequence.

The arrow from "Engineer's Report" to "Manager's Decision" is a sequence arc; the manager will hear from the engineer before deciding. Thus, the engineer's preliminary report is information available at the time of the decision, and this influence diagram represents the situation while the manager is waiting to hear from the engineer. Analyzing the influence diagram will tell the manager how to interpret this information; the appropriate action will depend not only on the engineer's report but also on the extent to which the manager believes the engineer to be correct. The manager's assessment of the engineer's accuracy is reflected in the chances associated with the "Engineer's Report" node. Note that a relevance arc points from "Machine 3 OK?" to "Engineer's Report," indicating that Machine 3's state is relevant for assessing the chances associated with the engineer's report. For example, if the manager believes the engineer is very good at diagnosing the situation, then when Machine 3 really is OK, the chances should be near 100% that the engineer will say so. Likewise, if Machine 3 is causing the defective products, the engineer should be very likely to indicate 3 is the problem. On the other hand, if the manager does not think the engineer is very good at diagnosing the problem—because of lack

Figure 3.10
*Influence diagram for
manufacturing-plant
manager's imperfect
information.*



of familiarity with this particular piece of equipment, say — then there might be a substantial chance that the engineer makes a mistake.

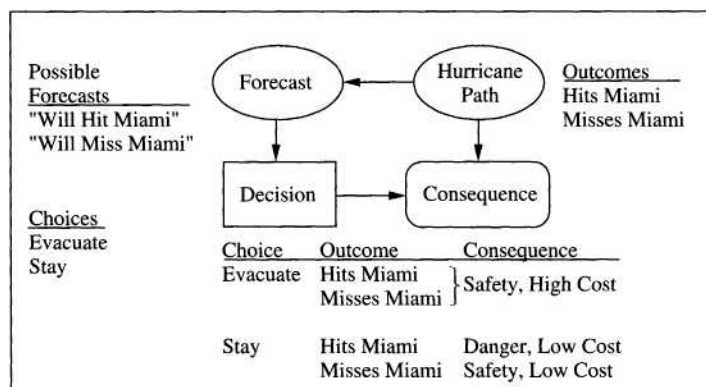
Weather forecasting provides another example of imperfect information. Suppose you live in Miami. A hurricane near the Bahama Islands threatens to cause severe damage; as a result, the authorities recommend that everyone evacuate. Although evacuation is costly, you would be safe. On the other hand, staying is risky. You could be injured or even killed if the storm comes ashore within 10 miles of your home. If the hurricane's path changes, however, you would be safe without having incurred the cost of evacuating. Clearly, two fundamental objectives are to maximize your safety and to minimize your costs.

Undoubtedly, you would pay close attention to the weather forecasters who would predict the course of the storm. These weather forecasters are not perfect predictors, however. They can provide some information about the storm, but they may not perfectly predict its path because not everything is known about hurricanes.

Figure 3.11 shows the influence diagram for the evacuation decision. The relevance arc from "Hurricane Path" to "Forecast" means that the actual weather situation is relevant for assessing the uncertainty associated with the forecast. If the hurricane is actually going to hit Miami, then the forecaster (we hope) is more likely to predict a hit rather than a miss. Conversely, if the hurricane really will miss Miami, the forecaster should be likely to predict a miss. In either case, though, the forecast may be incorrect because the course of a hurricane is not fully predictable. In this situation, although the forecast actually precedes the hurricane's landfall, it is relatively straightforward to think about the forecaster's tendency to make a mistake conditioned on what direction the hurricane goes. (The modeling choice is up to you, though! If you would feel more confident in assessing the chance of the hurricane hitting Miami by conditioning on the forecast — that is, have the arrow pointing the other way — then by all means do so!)

The consequence node in Figure 3.11 encompasses both objectives of minimizing cost and maximizing safety. An alternative representation might explicitly include both consequences as separate nodes as in Figure 3.10. Moreover, these two

Figure 3.11
Influence diagram for
the evacuation
decision.



objectives are somewhat vaguely defined, as they might be in an initial specification of the decision. A more complete specification would define these objectives carefully, giving levels of cost (probably in dollars) and a scale for the level of danger. In addition, uncertainty about the possible outcomes — ranging from no injury to death — could be included in the influence diagram. You will get a chance in Problem 3.14 to modify and improve on Figure 3.11.

As with the manufacturing example, the influence diagram in Figure 3.11 is a snapshot of your situation as you wait to hear from the forecaster. The sequence arc from "Forecast" to the decision node indicates that the decision is made knowing the imperfect weather forecast. You might imagine yourself waiting for the 6 p.m. weather report on the television, and as you wait, you consider what the forecaster might say and what you would do in each case. The sequence of events, then, is that the decision maker hears the forecast, decides what to do, and then the hurricane either hits Miami or misses. As with the manufacturing example, analyzing this model will result in a strategy that recommends a particular decision for each of the possible statements the forecaster might make.

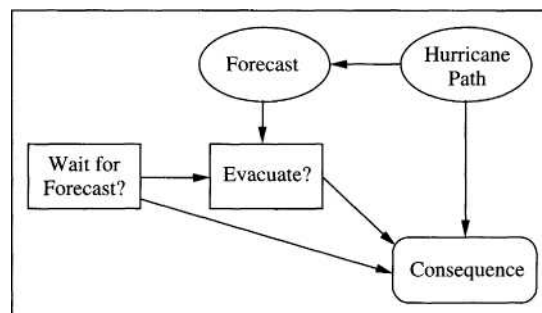
Sequential Decisions

The hurricane-evacuation decision above can be thought of as part of a larger picture. Suppose you are waiting anxiously for the forecast as the hurricane is bearing down. Do you wait for the forecast or leave immediately? If you wait for the forecast, what you decide to do may depend on that forecast. In this situation, you face a sequential decision situation as diagrammed in Figure 3.12.

The order of the events is implied by the arcs. Because there is no arc from "Forecast" to "Wait for Forecast," but there is one to "Evacuate," it is clear that the sequence is first to decide whether to wait or leave immediately. If you wait, the forecast is revealed, and finally you decide, based on the forecast, whether to evacuate.

In an influence diagram sequential decisions are strung together via sequence arcs, in much the same way that we did in Chapter 2. (In fact, now you can see that the figures in Chapter 2 use essentially the same graphics as influence diagrams!) For another example, let us take the farmer's decision from Chapter 2

Figure 3.12
*A sequential version of
the evacuation
decision.*

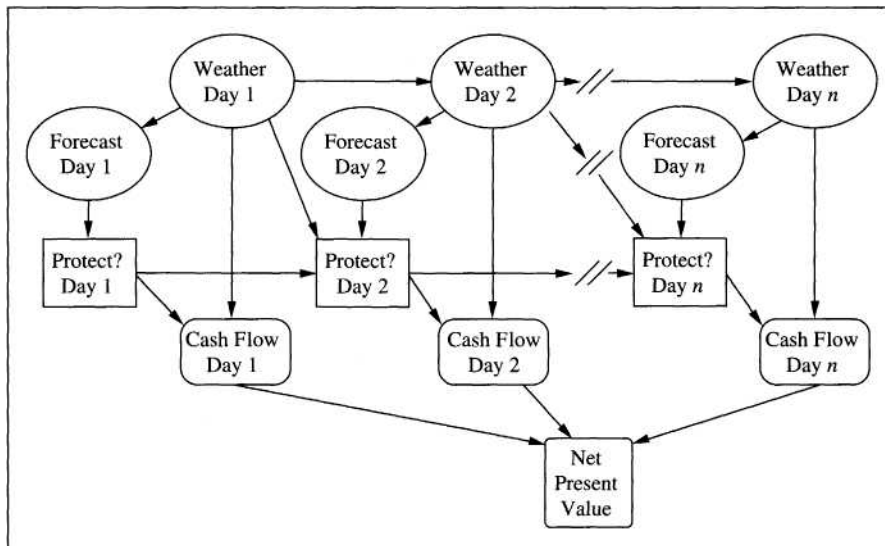


about protecting his trees against adverse weather. Recall that the farmer's decision replayed itself each day; based on the next day's weather forecast, should the fruit crop be protected? Let us assume that the farmer's fundamental objective is to maximize the NPV of the investment, including the costs of protection. Figure 3.13 shows that the influence diagram essentially is a series of imperfect-information diagrams strung together. Between decisions (to protect or not) the farmer observes the weather and obtains the forecast for the next day. The arcs from one decision to the next show the time sequence.

The arrows among the weather and forecast nodes from day to day indicate that the observed weather and the forecast both have an effect. That is, yesterday's weather is relevant for assessing the chance of adverse weather today. Not shown explicitly in the influence diagram are arcs from forecast and weather nodes before the previous day. Of course, the decision maker observed the weather and the forecasts for each prior day. These arcs are not included in the influence diagram but are implied by the arcs that connect the decision nodes into a time sequence. The missing arcs are sometimes called *no-forgetting* arcs to indicate that the decision maker would not forget the outcomes of those previous events. Unless the no-forgetting arcs are critical in understanding the situation, it is best to exclude them because they tend to complicate the diagram.

Finally, although we indicated that the farmer has a single objective, that of maximizing NPV, Figure 3.13 represents the decision as a multiple-objective one, the objectives being to maximize the cash inflow (and hence minimize outflows or costs) each day. The individual cash flows, of course, are used to calculate the farmer's NPV. As indicated in Chapter 2, the interest rate defines the trade-off between earlier and later cash flows.

Figure 3.13
Influence diagram for
farmer's sequential
decision problem.



Intermediate Calculations

In some cases it is convenient to include an additional node that simply aggregates results from certain predecessor nodes. Suppose, for example, that a firm is considering introducing a new product. The firm's fundamental objective is the profit level of the enterprise, and so we label the consequence node "Profit." At a very basic level, both cost and revenue may be uncertain, and so a first version of the influence diagram might look like the one shown in Figure 3.14.

On reflection, the firm's chief executive officer (CEO) realizes that substantial uncertainty exists for both variable and fixed costs. On the revenue side, there is uncertainty about the number of units sold, and a pricing decision will have to be made. These considerations lead the CEO to consider a somewhat more complicated influence diagram, which is shown in Figure 3.15.

Figure 3.15 is a perfectly adequate influence diagram. Another representation is shown in Figure 3.16. Intermediate nodes have been included in Figure 3.16 to calculate cost on one hand and revenue on the other; we will call these *calculation nodes*, because they calculate cost and revenue given the predecessors. (In many discussions of influence diagrams, the term *deterministic node* is used to denote a node that represents an intermediate calculation or a constant, and in graphical representation in the influence diagram such a node is shown as a circle with a double outline. The use of the rounded rectangle, the same as the consequence node, is consistent with the repre-

Figure 3.14
Simple influence diagram for new product decision.

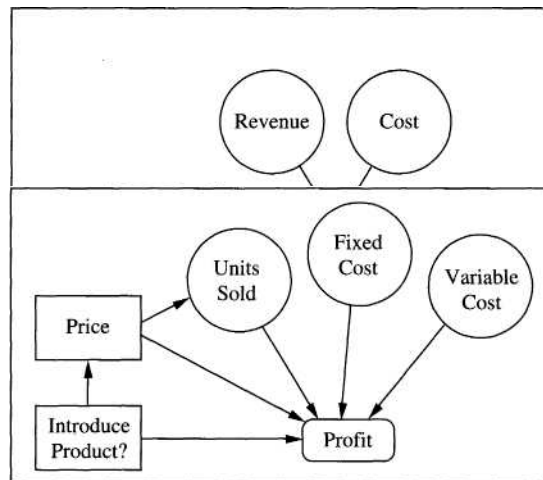



Figure 3.15
New product decision with additional detail



sentation in the computer program DPL and the discussion of these nodes in its documentation.)

Calculation nodes behave just like consequence nodes; given the inputs from the predecessor nodes, the value of a calculation node can be found immediately. No uncertainty exists after the conditioning variables — decisions, chance events, or other calculation nodes — are known. Of course, there is no uncertainty only in a conditional sense; the decision maker can look forward in time and know what the calculation node will be for any possible combination of the conditioning variables. Before the conditioning variables are known, though, the value that the node will eventually have is uncertain.

In general, calculation nodes are useful for emphasizing the structure of an influence diagram. Whenever a node has a lot of predecessors, it may be appropriate to include one or more intermediate calculations to define the relationships among the predecessors more precisely. In Figure 3.16, the calculation of cost and revenue is represented explicitly, as is the calculation of profit from cost and revenue. The pricing decision is clearly related to the revenue side, uncertainty about fixed and variable costs are clearly on the cost side, and uncertainty about sales is related to both.

Another example is shown in Figure 3.17. In this situation, a firm is considering building a new manufacturing plant that may create some incremental pollution. The profitability of the plant depends on many things, of course, but highlighted in Figure 3.17 are the impacts of other pollution sources. The calculation node "Regional Pollution Level" uses information on the number of cars and local industry growth to determine a pollution-level index. The pollution level in turn has an impact on the chances that the new plant will be licensed and that new regulations (either more or less strict) will be imposed.

With the basic understanding of influence diagrams provided above, you should be able to look at any influence diagram (including any that you find in this book) and understand what it means. Understanding an influence diagram is an important decision-analysis skill. On the other hand, actually creating an influence diagram from scratch is considerably more difficult and takes much practice. The following optional section gives an example of the construction process for an influence dia-

Figure 3.16
New product decision
with calculation
nodes for intermediate
calculation of cost
and revenue.

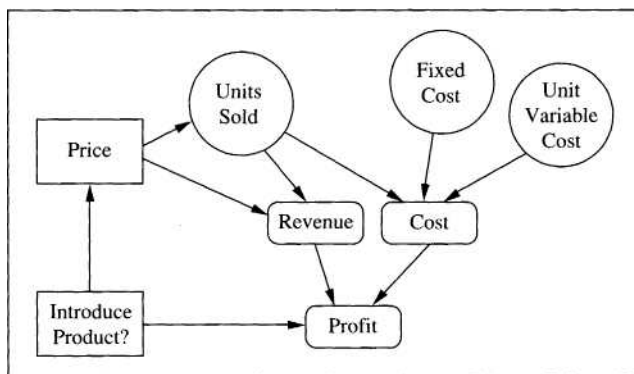
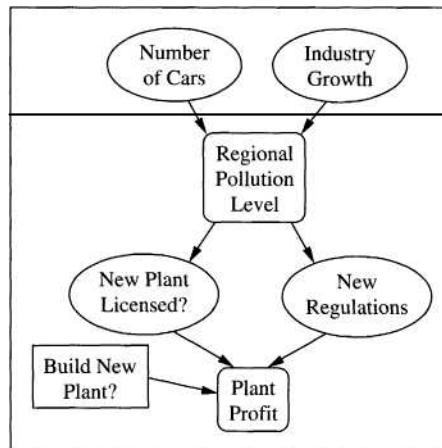


Figure 3.17
Using a calculation
node to determine
pollution level.



gram and discusses some common mistakes. If you wish to become proficient in constructing influence diagrams, the next section is highly recommended. Working through the reading and exercises, however, is just one possibility; in fact, practice with an influence-diagram program (like DPL, which is discussed later in this chapter) is the best way to develop skill in constructing influence diagrams.

Constructing an Influence Diagram (Optional)

There is no set strategy for creating an influence diagram. Because the task is to structure a decision that may be complicated, the best approach may be to put together a simple version of the diagram first and add details as necessary until the diagram captures all of the relevant aspects of the problem. In this section, we will demonstrate the construction of an influence diagram for the classic toxic-chemical problem.

TOXIC CHEMICALS AND THE EPA

The Environmental Protection Agency (EPA) often must decide whether to permit the use of an economically beneficial chemical that may be carcinogenic (cancer-causing). Furthermore, the decision often must be made without perfect information about either the long-term benefits or health hazards. Alternative courses of action are to permit the use of the chemical, restrict its use, or to ban it altogether. Tests can be run to learn something about the carcinogenic potential of the material, and survey data can give an indication of the extent to which people are exposed when they do use the chemical. These pieces of information are both important in making the decision. For example, if

the chemical is only mildly toxic and the exposure rate is minimal, then restricted use may be reasonable. On the other hand, if the chemical is only mildly toxic but the exposure rate is high, then banning its use may be imperative.

The first step should be to identify the decision context and the objectives. In this case, the context is choosing an allowed level of use, and the fundamental objectives are to maximize the economic benefits from the chemicals and at the same time to minimize the risk of cancer. These two objectives feed into an overall consequence node ("Net Value") that aggregates "Economic Value" and "Cancer Cost" as shown in Figure 3.18.

Now let us think about what affects "Economic Value" and "Cancer Cost" other than the usage decision. Both the uncertain carcinogenic character of the chemical and the exposure rate have an effect on the cancer cost that could occur, thus yielding the diagram shown in Figure 3.19. Because "Carcinogenic Potential" and "Exposure Rate" jointly determine the level of risk that is inherent in the chemical, their effects are aggregated in an intermediate calculation node labeled "Cancer Risk." Different values of the predecessor nodes will determine the overall level of "Cancer Risk."

Note that no arrow runs from "Usage Decision" to "Exposure Rate," even though such an arrow might appear to make sense. "Exposure Rate" refers to the extent of contact when the chemical is actually used and would be measured in terms of an amount of contact per unit of time (e.g., grams of dust inhaled per hour). The rate is unknown, and the usage decision does not influence our beliefs concerning the likelihood of various possible rates when the chemical is used.

The influence diagram remains incomplete, however, because we have not incorporated the test for carcinogenicity or the survey on exposure. Presumably, results

Figure 3.18
Beginning the toxic-chemical influence diagram.





Figure 3.19
Intermediate influence diagram for the toxic-chemical decision.



from both the test (called a *bioassay*) and the survey would be available to EPA at the time the usage decision is made. Furthermore, it should be clear that the actual degrees of carcinogenic potential and exposure will influence the test and survey results, and thus "Carcinogenic Potential" and "Exposure Rate" are connected to "Test" and "Survey," respectively, in Figure 3.20. Note that "Test" and "Survey" each represent imperfect information; each one provides some information regarding carcinogenicity or exposure. These two nodes are connected to the decision node. These are sequence arcs, indicating that the information is available when the decision is made. This completes the influence diagram.

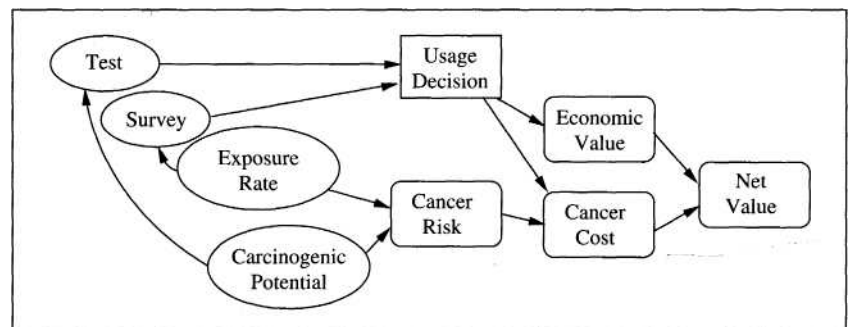
This example demonstrates the usefulness of influence diagrams for structuring decisions. The toxic-chemicals problem is relatively complex, and yet its influence diagram is compact and, more important, understandable. Of course, the more complicated the problem, the larger the influence diagram. Nevertheless, influence diagrams are useful for creating easily understood overviews of decision situations.

Some Common Mistakes

First, an easily made mistake in understanding and constructing influence diagrams is to interpret them as flowcharts, which depict the sequential nature of a particular process where each node represents an event or activity. For example, Figure 1.1 is a flowchart of a decision-analysis system, displaying the different things a decision analyst does at each stage of the process.

Even though they look a little like flowcharts, influence diagrams are very different. An influence diagram is a snapshot of the decision situation at a particular time, one that must account for all the decision elements that play a part in the immediate decision. Putting a chance node in an influence diagram means that although the decision maker is not sure exactly what will happen, he or she has some idea of how likely the different possible outcomes are. For example, in the toxic-chemical problem, the carcinogenic potential of the chemical is unknown, and in fact will never be known for sure. That uncertainty, however, can be modeled using probabilities for different

Figure 3.20
Complete influence diagram for the toxic-chemical decision.



levels of carcinogenic potential. Likewise, at the time the influence diagram is created, the results of the test are not known. The uncertainty surrounding the test results also can be modeled using probabilities. The informational arrow from "Test" to "Usage Decision," however, means that the decision maker will learn the results of the test before the decision must be made.

The metaphor of a picture of the decision that accounts for all of the decision elements also encompasses the possibility of upcoming decisions that must be considered. For example, a legislator deciding how to vote on a given issue may consider upcoming votes. The outcome of the current issue might affect the legislator's future voting decisions. Thus, at the time of the immediate decision, the decision maker foresees future decisions and models those decisions with the knowledge on hand.

A second common mistake, one related to the perception of an influence diagram as a flowchart, is building influence diagrams with many chance nodes having arrows pointing to the primary decision node. The intention usually is to represent the uncertainty in the decision environment. The problem is that the arrows into the decision node are sequence arcs and indicate that the decision maker is waiting to learn the outcome of these uncertain events, which may not be the case. The solution is to think carefully when constructing the influence diagram. Recall that only sequence arcs are used to point to a decision node. Thus, an arrow into a decision node means that the decision maker will have a specific bit of information when making the decision; something will be known for sure, with no residual uncertainty. Before drawing an arrow into a decision node, ask whether the decision maker is waiting for the event to occur *and* will learn the information before the decision is made. If so, the arrow is appropriate. If not, don't draw the arrow!

So how should you include information that the decision maker has about the uncertainty in the decision situation? The answer is simple. Recall that the influence diagram is a snapshot of the decision maker's understanding of the decision situation at a particular point in time. When you create a chance node and connect it appropriately in the diagram, you are explicitly representing the decision maker's uncertainty about that event and showing how that uncertainty relates to other elements of the decision situation.

A third mistake is the inclusion of cycles (circular paths among the nodes). As indicated previously, a properly constructed influence diagram contains no cycles. Cycles are occasionally included in an attempt to denote feedback among the chance and decision nodes. Although this might be appropriate in the case of a flowchart, it is inappropriate in an influence diagram. Again, think about the diagram as a picture of the decision that accounts for all of the decision elements at an instant in time. There is no opportunity for feedback at a single point in time, and hence there can be no cycles.

Influence diagrams provide a graphical representation of a decision's structure, a snapshot of the decision environment at one point in time. All of the details (alternatives, outcomes, consequences) are present in tables that are contained in the nodes, but usually this information is suppressed in favor of a representation that shows off the decision's structure.

Multiple Representations and Requisite Models

Even though your influence diagram may be technically correct in the sense that it contains no mistakes, how do you know whether it is the "correct" one for your decision situation? This question presupposes that a unique correct diagram exists, but for most decision situations, there are many ways in which an influence diagram can appropriately represent a decision. Consider the decision modeled in Figures 3.14, 3.15, and 3.16; these figures represent three possible approaches. With respect to uncertainty in a decision problem, several sources of uncertainty may underlie a single chance node. For example, in Figure 3.16, units sold may be uncertain because the CEO is uncertain about the timing and degree of competitive reactions, the nature of consumer tastes, the size of the potential market, the effectiveness of advertising, and so on. In many cases, and certainly for a first-pass representation, the simpler model may be more appropriate. In other cases, more detail may be necessary to capture all of the essential elements of a situation. In the farmer's problem, for example, a faithful representation of the situation may require consideration of the sequence of decisions rather than looking at each decision as being independent and separate from the others. Thus, different individuals may create different influence diagrams for the same decision problem, depending on how they view the problem. The real issue is determining whether a diagram is appropriate. Does it capture and accurately reflect the elements of the decision problem that the decision maker thinks are important?

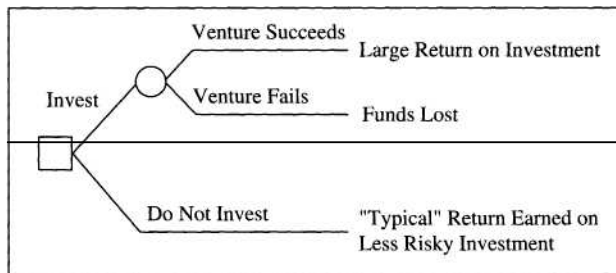
How can you tell whether your influence diagram is an appropriate one? The representation that is the most appropriate is the one that is *requisite* for the decision maker along the lines of our discussion in Chapter 1. That is, a requisite model contains everything that the decision maker considers important in making the decision. Identifying all of the essential elements may be a matter of working through the problem several times, refining the model on each pass. The only way to get to a requisite decision model is to continue working on the decision until all of the important concerns are fully incorporated. Sensitivity analysis (Chapter 5) will be a great help in determining which elements are important.

Structuring Decisions: Decision Trees

Influence diagrams are excellent for displaying a decision's basic structure, but they hide many of the details. To display more of the details, we can use a *decision tree*. As with influence diagrams, squares represent decisions to be made, while circles represent chance events. The branches emanating from a square correspond to the choices available to the decision maker, and the branches from a circle represent the possible outcomes of a chance event. The third decision element, the consequence, is specified at the ends of the branches.

Again consider the venture-capital decision (Figure 3.5). Figure 3.21 shows the decision tree for this problem. The decision tree flows from left to right, and so the

Figure 3.21
Decision-tree
representation of
venture-capital
decision.



immediate decision is represented by the square at the left side. The two branches represent the two alternatives, invest or not. If the venture capitalist invests in the project, the next issue is whether the venture succeeds or fails. If the venture succeeds, the capitalist earns a large return. However, if the venture fails, then the amount invested in the project will be lost. If the capitalist decides not to invest in this particular risky project, then she would earn a more typical return on another less risky project. These outcomes are shown at the ends of the branches at the right.

The interpretation of decision trees requires explanation. First, the options represented by branches from a decision node must be such that the decision maker can choose only one option. For example, in the venture-capital decision, the decision maker can either invest or not, but not both. In some instances, combination strategies are possible. If the capitalist were considering two separate projects (A and B), for instance, it may be possible to invest in Firm A, Firm B, both, or neither. In this case, each of the four separate alternatives would be modeled explicitly, yielding four branches from the decision node.

Second, each chance node must have branches that correspond to a set of *mutually exclusive* and *collectively exhaustive* outcomes. *Mutually exclusive* means that only one of the outcomes can happen. In the venture-capital decision, the project can either succeed or fail, but not both. *Collectively exhaustive* means that no other possibilities exist; one of the specified outcomes has to occur. Putting these two specifications together means that when the uncertainty is resolved, one and only one of the outcomes occurs.

Third, a decision tree represents all of the possible paths that the decision maker might follow through time, including all possible decision alternatives and outcomes of chance events. Three such paths exist for the venture capitalist, corresponding to the three branches at the right-hand side of the tree. In a complicated decision situation with many sequential decisions or sources of uncertainty, the number of potential paths may be very large.

Finally, it is sometimes useful to think of the nodes as occurring in a time sequence. Beginning on the left side of the tree, the first thing to happen is typically a decision, followed by other decisions or chance events in chronological order. In the venture-capital problem, the capitalist decides first whether to invest, and the second step is whether the project succeeds or fails.

As with influence diagrams, the dovetailing of decisions and chance events is critical. Placing a chance event before a decision means that the decision is made conditional on the specific chance outcome having occurred. Conversely, if a chance

node is to the right of a decision node, the decision must be made in anticipation of the chance event. The sequence of decisions is shown in a decision tree by order in the tree from left to right. If chance events have a logical time sequence between decisions, they may be appropriately ordered. If no natural sequence exists, then the order in which they appear in the decision tree is not critical, although the order used does suggest the conditioning sequence for modeling uncertainty. For example, it may be easier to think about the chances of a stock price increasing given that the Dow Jones average increases rather than the other way around.

Decision Trees and the Objectives Hierarchy

Including multiple objectives in a decision tree is straightforward; at the end of each branch, simply list all of the relevant consequences. An easy way to do this systematically is with a *consequence matrix* such as Figure 3.22, which shows the FAA's bomb-detection decision in decision-tree form. Each column of the matrix represents a fundamental objective, and each row represents an alternative, in this case a candidate detection system. Evaluating the alternatives requires "filling in the boxes" in the matrix; each alternative must be measured on every objective. Thus every detection system must be evaluated in terms of detection effectiveness, implementation time, passenger acceptance, and cost.

Figure 3.23 shows a decision-tree representation of the hurricane example. The initial "Forecast" branch at the left indicates that the evacuation decision would be made conditional on the forecast made — recall the imperfect-information decision situation shown in the influence diagram in Figure 3.11. This figure demonstrates that consequences must be considered for every possible endpoint at the right side of the decision tree, regardless of whether those endpoints represent a decision alternative or an uncertain outcome. In addition, Figure 3.23 shows clearly the nature of the risk that the decision to stay entails, and that the decision maker must make a fundamental trade-off between the sure safety of evacuating and the cost of doing so. Finally, the extent of the risk may depend strongly on what the forecast turns out to be!

Figure 3.22
Decision-tree representation of FAA's multiple-objective bomb-detection decision.

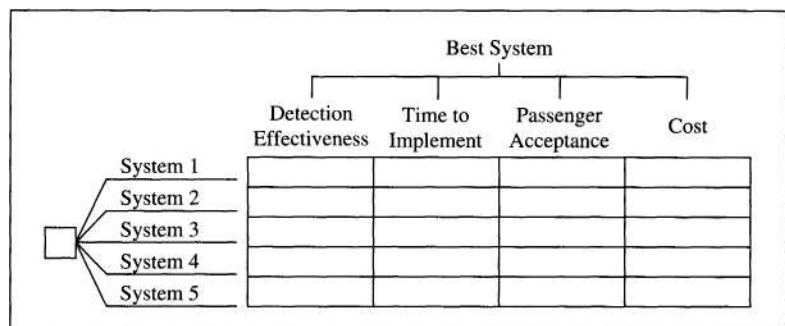
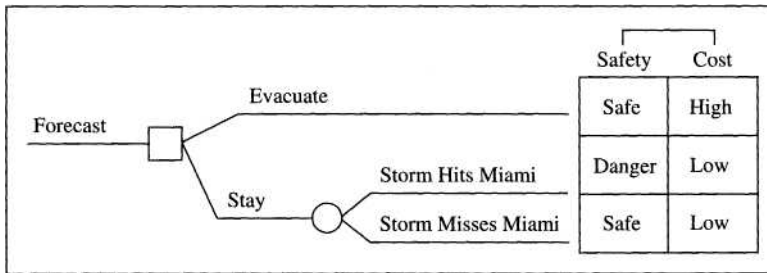


Figure 3.23
Decision-tree representation of evacuation decision.



Some Basic Decision Trees

In this section we will look at some basic decision-tree forms. Many correspond to the basic influence diagrams discussed above.


The Basic Risky Decision

Just as the venture-capital decision was the prototypical basic risky decision in our discussion of influence diagrams, so it is here as well. The capitalist's dilemma is whether the potential for large gains in the proposed project is worth the additional risk. If she judges that it is not, then she should not invest in the project.

Figure 3.24 shows the decision-tree representation of the investment decision given earlier in influence-diagram form in Figure 3.9. In the decision tree you can see how the sequence of events unfolds. Beginning at the left side of the tree, the choice is made whether to invest in the business or savings. If the business is chosen, then the outcome of the chance event (wild success or a flop) occurs, and the consequence — the final cash position — is determined. As before, the essential question is whether the chance of wild success and ending up with \$5000 is worth the risk of losing everything, especially in comparison to the savings account that results in a bank balance of \$2200 for sure.

For another example, consider a politician's decision. The politician's fundamental objectives are to have a career that provides leadership for the country and representation for her constituency, and she can do so to varying degrees by serving in congress. The politician might have the options of (1) running for reelection to her

Figure 3.24
The investor's basic risky decision.



U.S. House of Representatives seat, in which case reelection is virtually assured, or (2) running for a Senate seat. If the choice is to pursue the Senate seat, there is a chance of losing, in which case she could return to her old job as a lawyer (the worst possible outcome). On the other hand, winning the Senate race would be the best possible outcome in terms of her objective of providing leadership and representation. Figure 3.25 diagrams the decision. The dilemma in the basic risky decision arises because the riskless alternative results in an outcome that, in terms of desirability, falls between the outcomes for the risky alternatives. (If this were not the case, there would be no problem deciding!) The decision maker's task is to figure out whether the chance of "winning" in the risky alternative is great enough relative to the chance of "losing" to make the risky alternative more valuable than the riskless alternative. The more valuable the riskless alternative, the greater the chance of winning must be for the risky alternative to be preferred.

A variation of the basic risky decision might be called the *double-risk decision dilemma*. Here the problem is deciding between two risky prospects. On one hand, you are "damned if you do and damned if you don't" in the sense that you could lose either way. On the other hand, you could win either way. For example, the political candidate may face the decision represented by the decision tree in Figure 3.26, in which she may enter either of two races with the possibility of losing either one.

In our discussion of the basic risky decision and influence diagrams, we briefly mentioned the *range-of-risk dilemma*, in which the outcome of the chance event can take on any value within a range of possible values. For example, imagine an individual who has sued for damages of \$450,000 because of an injury. The insurance company has offered to settle for \$100,000. The plaintiff must decide whether to accept the settlement or go to court; the decision tree is shown as Figure 3.27. The crescent shape

Figure 3.25
The politician's basic risky decision.




Figure 3.26
Double-risk decision dilemma.

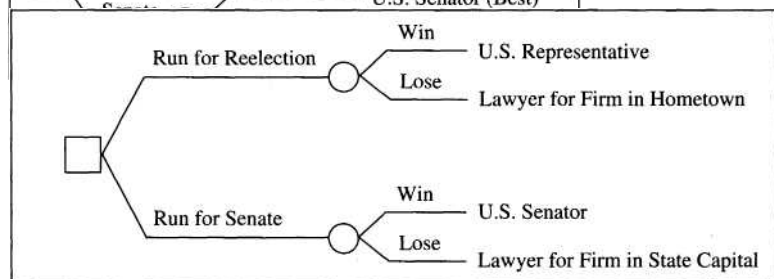
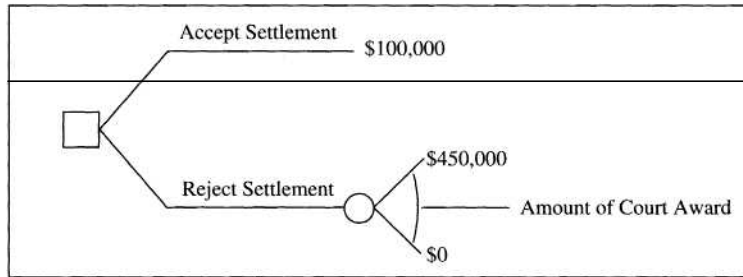


Figure 3.27
A range-of-risk
decision dilemma.



indicates that the uncertain event — the court award — may result in any value between the extremes of zero and \$450,000, the amount claimed in the lawsuit.

Imperfect Information

Representing imperfect information with decision trees is a matter of showing that the decision maker is waiting for information prior to making a decision. For example, the evacuation decision problem is shown in Figure 3.28. Here is a decision tree that begins with a chance event, the forecast. The chronological sequence is clear; the forecast arrives, then the evacuation decision is made, and finally the hurricane either hits or misses Miami.

Sequential Decisions

As we did in the discussion of influence diagrams, we can modify the imperfect-information decision tree to reflect a sequential decision situation in which the first choice is whether to wait for the forecast or evacuate now. Figure 3.29 shows this decision tree.

At this point, you can imagine that representing a sequential decision problem with a decision tree may be very difficult if there are many decisions and chance events because the number of branches can increase dramatically under such conditions. Although full-blown decision trees work poorly for this kind of problem, it is possible to use a *schematic* approach to depict the tree.

Figure 3.28
Evacuation decision
represented by
decision tree.

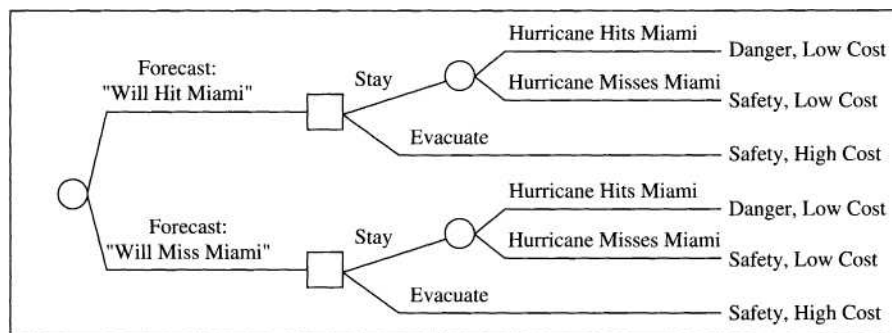


Figure 3.29
 Sequential version of
 evacuation decision in
 decision-tree form.

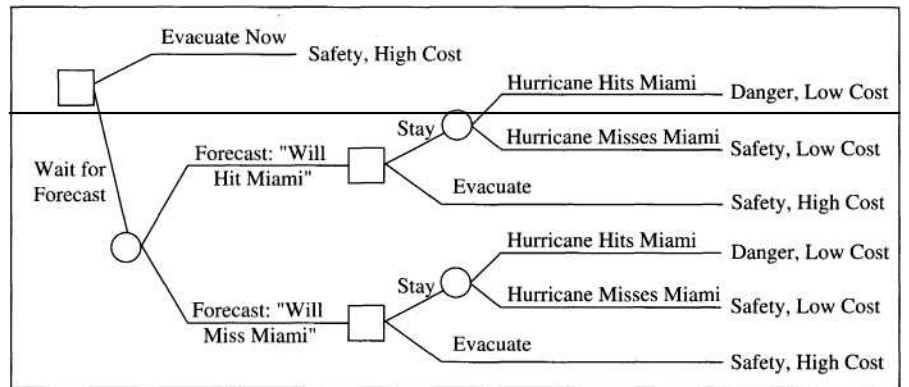


Figure 3.30 shows a schematic version of the farmer's sequential decision problem. This is the decision-tree version of Figure 3.13. Even though each decision and chance event has only two branches, we are using the crescent shape to avoid having the tree explode into a bushy mess. With only the six nodes shown, there would be 2^6 , or 64, branches.

We can string together the crescent shapes sequentially in Figure 3.30 because, regardless of the outcome or decision at any point, the same events and decisions follow in the rest of the tree. This ability is useful in many kinds of situations. For example, Figure 3.31 shows a decision in which the immediate decision is whether to invest in an entrepreneurial venture to market a new product or invest in the stock

Figure 3.30
 Schematic version of
 farmer's sequential
 decision: decision-tree
 form.


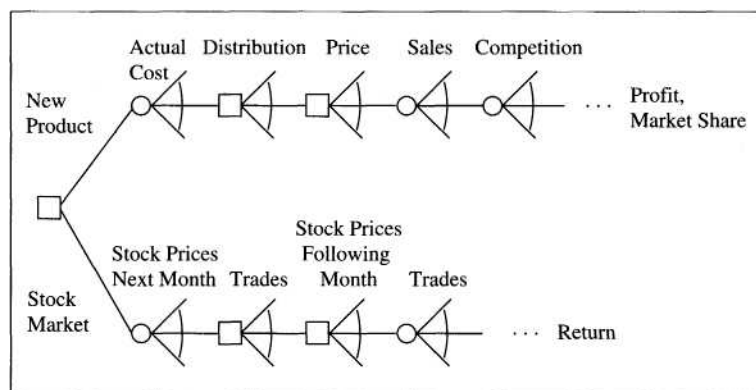


Figure 3.31
 An investment decision
 in schematic form.



market. Each alternative leads to its own set of decisions and chance events, and each set can be represented in schematic form.

Decision Trees and Influence Diagrams Compared

It is time to step back and compare decision trees with influence diagrams. The discussion and examples have shown that, on the surface at least, decision trees display considerably more information than do influence diagrams. It should also be obvious, however, that decision trees get "messy" much faster than do influence diagrams as decision problems become more complicated. One of the most complicated decision trees we constructed was for the sequential decision in Figure 3.31, and it really does not show all of the intricate details contained in the influence-diagram version of the same problem. The level of complexity of the representation is not a small issue. When it comes time to present the results of a decision analysis to upper-level managers, their understanding of the graphical presentation is crucial. Influence diagrams are superior in this regard; they are especially easy for people to understand regardless of mathematical training.

Should you use decision trees or influence diagrams? Both are worthwhile, and they complement each other well. Influence diagrams are particularly valuable for the structuring phase of problem solving and for representing large problems. Decision trees display the details of a problem. The ultimate decision made should not depend on the representation, because influence diagrams and decision trees are *isomorphic*, any properly built influence diagram can be converted into a decision tree, and vice versa, although the conversion may not be easy. One strategy is to start by using an influence diagram to help understand the major elements of the situation and then convert to a decision tree to fill in details.

Influence diagrams and decision trees provide two approaches for modeling a decision. Because the two approaches have different advantages, one may be more appropriate than the other, depending on the modeling requirements of the particular situation. For example, if it is important to communicate the overall structure of a model to other people, an influence diagram may be more appropriate. Careful reflection and sensitivity analysis on specific probability and value inputs may work better in the context of a decision tree. Using both approaches together may prove useful; the goal, after all, is to make sure that the model accurately represents the decision situation. Because the two approaches have different strengths, they should be viewed as complementary techniques rather than as competitors in the decision-modeling process.

Decision Details: Defining Elements of the Decision

With the overall structure of the decision understood, the next step is to make sure that all elements of the decision model are clearly defined. Beginning efforts to structure decisions usually include some rather loose specifications. For example,

when the EPA considers regulating the use of a potentially cancer-causing substance, it would have a fundamental objective of minimizing the social cost of the cancers. (See, for example, Figure 3.20 and the related discussion.) But how will cancer cost be measured? In incremental lives lost? Incremental cases of cancer, both treatable and fatal? In making its decision, the EPA would also consider the rate at which people are exposed to the toxin while the chemical is in use. What are possible levels of exposure? How will we measure exposure? Are we talking about the number of people exposed to the chemical per day or per hour? Does exposure consist of breathing dust particles, ingesting some critical quantity, or skin contact? Are we concerned about contact over a period of time? Exactly how will we know if an individual has had a high or low level of exposure? The decision maker must give unequivocal answers to these questions before the decision model can be used to resolve the EPA's real-world policy problem.

Much of the difficulty in decision making arises when different people have different ideas regarding some aspect of the decision. The solution is to refine the conceptualizations of events and variables associated with the decision enough so that it can be made. How do we know when we have refined enough? The *clarity test* (Howard, 1988) provides a simple and understandable answer. Imagine a clairvoyant who has access to all future information: newspapers, instrument readings, technical reports, and so on. Would the clairvoyant be able to determine unequivocally what the outcome would be for any event in the influence diagram? No interpretation or judgment should be required of the clairvoyant. Another approach is to imagine that, in the future, perfect information will be available regarding all aspects of the decision. Would it be possible to tell exactly what happened at every node, again with no interpretation or judgment? The decision model passes the clarity test when these questions are answered affirmatively. At this point, the problem should be specified clearly enough so that the various people involved in the decision are thinking about the decision elements in exactly the same way. There should be no misunderstandings regarding the definitions of the basic decision elements.

The clarity test is aptly named. It requires absolutely clear definitions of the events and variables. In the case of the EPA considering toxic substances, saying that the exposure rate can be either high or low fails the clarity test; what does "high" mean in this case? On the other hand, suppose exposure is defined as high if the average skin contact per person-day of use exceeds an average of 10 milligrams of material per second over 10 consecutive minutes. This definition passes the clarity test. An accurate test could indicate precisely whether the level of exposure exceeded the threshold.

Although Howard originally defined the clarity test in terms of only chance nodes, it can be applied to all elements of the decision model. Once the problem is structured and the decision tree or influence diagram built, consider each node. Is the definition of each chance event clear enough so that an outside observer would know exactly what happened? Are the decision alternatives clear enough so that someone else would know exactly what each one entails? Are consequences clearly defined and measurable? All of the action with regard to the clarity test takes place within the tables in an influence diagram, along the individual branches of a decision tree, or in the tree's consequence matrix. These are the places where the critical decision details are specified. Only after every element of the decision model passes the clarity test is

it appropriate to consider solving the influence diagram or decision tree, which is the topic of Chapter 4.

The next two sections explore some specific aspects of decision details that must be included in a decision model. In the first section we look at how chances can be specified by means of probabilities and, when money is an objective, how cash flows can be included in a decision tree. These are rather straightforward matters in many of the decisions we make. However, when we have multiple fundamental objectives, defining ways to measure achievement of each objective can be difficult; it is easy to measure costs, savings, or cash flows in dollars or pounds sterling, but how does one measure damage to an ecosystem? Developing such measurement scales is an important aspect of attaining clarity in a decision model and is the topic of the second section.

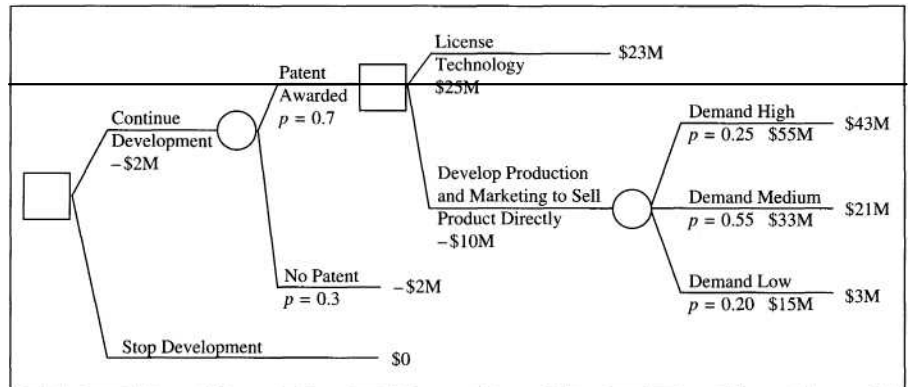
More Decision Details: Cash Flows and Probabilities

Many decision situations, especially business decisions, involve some chance events, one or more decisions to make, and a fundamental objective that can be measured in monetary terms (maximize profit, minimize cost, and so on). In these situations, once the decisions and chance events are defined clearly enough to pass the clarity test, the last step is to specify the final details: specific chances associated with the uncertain events and the cash flows that may occur at different times. What are the chances that a particular outcome will occur? What does it cost to take a given action? Are there specific cash flows that occur at different times, depending on an alternative chosen or an event's outcome?

Specifying the chances for the different outcomes at a chance event requires us to use probabilities. Although probability is the topic of Section 2 of the book, we will use probability in Chapters 4 and 5 as we develop some basic analytical techniques. For now, in order to specify probabilities for outcomes, you need to keep in mind only a few basic rules. First, probabilities must fall between 0 and 1 (or equivalently between 0% and 100%). There is no such thing as a 110% chance that some event will occur. Second, recall that the outcomes associated with a chance event must be such that they are mutually exclusive and collectively exhaustive; only one outcome can occur (you can only go down one path), but one of the set must occur (you must go down some path). The implication is that the probability assigned to any given chance outcome (branch) must be between 0 and 1, and for any given chance node, the probabilities for its outcomes must add up to 1.

Indicating cash flows at particular points in the decision model is straightforward. For each decision alternative or chance outcome, indicate the associated cash flow, either as part of the information in the corresponding influence-diagram node or on the appropriate branch in the decision tree. For example, in the toxic-chemical example, there are certainly economic costs associated with different possible regulatory actions. In the new-product decision (Figure 3.16), different cash inflows are associated with different quantities sold, and different outflows are associated with different costs. All of

Figure 3.32
A research-and-development decision.



these cash flows must be combined (possibly using net present value if the timing of the cash flows is substantially different) at the end of each branch in order to show exactly what the overall consequence is for a specific path through the decision model.

Figure 3.32 shows a decision tree with cash flows and probabilities fully specified. This is a research-and-development decision. The decision maker is a company that must decide whether to spend \$2 million to continue with a particular research project. The success of the project (as measured by obtaining a patent) is not assured, and at this point the decision maker judges only a 70% chance of getting the patent. If the patent is awarded, the company can either license the patent for an estimated \$25 million or invest an additional \$10 million to create a production and marketing system to sell the product directly. If the company chooses the latter, it faces uncertainty of demand and associated profit from sales.

You can see in Figure 3.32 that the probabilities at each chance node add up to 1. Also, the dollar values at the ends of the branches are the net values. For example, if the company continues development, obtains a patent, decides to sell the product directly, and enjoys a high level of demand, the net amount is \$43 million = $(-2) + (-10) + 55$ million. Also, note that cash flows can occur anywhere in the tree, either as the result of a specific choice made or because of a particular chance outcome.

Defining Measurement Scales for Fundamental Objectives

Many of our examples so far (and many more to come!) have revolved around relatively simple situations in which the decision maker has only one easily measured fundamental objective, such as maximizing profit, as measured in dollars. But the world is not always so accommodating. We often have multiple objectives, and some of those objectives are not easily measured on a single, natural numerical scale. What sort of measure should we use when we have fundamental objectives like maximizing our level of physical fitness, enhancing a company's

work environment, or improving the quality of a theatrical production? The answer, not surprisingly, relates back to the ideas embodied in the clarity test; we must find unambiguous ways to measure achievement of the fundamental objectives.

Before going on to the nuts and bolts of developing unambiguous scales, let us review briefly why the measurement of fundamental objectives is crucial in the decision process. The fundamental objectives represent the reasons why the decision maker cares about the decision and, more importantly, how the available alternatives should be evaluated. If a fundamental objective is to build market share, then it makes sense explicitly to estimate how much market share will change as part of the consequence of choosing a particular alternative. The change in market share could turn out to be good or bad depending on the choices made (e.g., bringing a new product to the market) and the outcome of uncertain events (such as whether a competitor launches an extensive promotional campaign). The fact that market share is something which the decision maker cares about, though, indicates that it must be measured.

It is impossible to overemphasize the importance of tying evaluation directly to the fundamental objectives. Too often decisions are based on the wrong measurements because inadequate thought is given to the fundamental objectives in the first place, or certain measurements are easy to make or are made out of habit, or the experts making the measurements have different objectives than the decision maker. An example is trying to persuade the public that high-tech endeavors like nuclear power plants or genetically engineered plants for agricultural use are not risky because few fatalities are expected; the fact is that the public appears to care about many other aspects of these activities as well as potential fatalities! (For example, lay people are very concerned with technological innovations that may have unknown long-term side effects, and they are also concerned with having little personal control over the risks that they may face because of such innovations.) In complex decision situations there may be many objectives that must be considered. The fundamental-objectives hierarchy indicates explicitly what must be accounted for in evaluating potential consequences.

The fundamental-objectives hierarchy starts at the top with an overall objective, and lower levels in the hierarchy describe important aspects of the more general objectives. Ideally, each of the lowest-level fundamental objectives in the hierarchy would be measured. Thus, one would start at the top and trace down as far as possible through the hierarchy. Reconsider the summer-intern example, in which PeachTree Consumer Products is looking for a summer employee to help with the development of a market survey. Figure 3.4 shows the fundamental-objectives hierarchy (as well as the means network). Starting at the top of this hierarchy ("Choose Best Intern"), we would go through "Maximize Quality and Efficiency of Work" and arrive at "Maximize Survey Quality" and "Minimize Survey Cost." Both of the latter require measurements to know how well they are achieved as a result of hiring any particular individual. Similarly, the other branches of the hierarchy lead to fundamental objectives that must be considered. Each of these objectives will be measured on a suitable scale, and that scale is called the objective's *attribute scale* or simply *attribute*.

As mentioned, many objectives have natural attribute scales: hours, dollars, percentage of market. Table 3.3 shows some common objectives with natural attributes. In the intern decision, "Minimize Survey Cost" would be easily measured in terms of dollars. How much must the company spend to complete the survey? In the context

Table 3.3
Some common objectives and their natural attributes.

Objective	Attribute
Maximize profit Maximize revenue Maximize savings Minimize cost	Money (for example, dollars)
Maximize market share Maximize rate of return	Percentage
Maximize proximity Maximize fuel efficiency Maximize time with friends, family Minimize hypertension	Miles, minutes Miles per gallon Days, hours Inches Hg (blood pressure)

of the decision situation, the relevant components of cost are salary, fringe benefits, and payroll taxes. Additional costs to complete the survey may arise if the project remains unfinished when the intern returns to school or if a substantial part of the project must be reworked. (Both of the latter may be important uncertain elements of the decision situation.) Still, for all possible combinations of alternative chosen and uncertain outcomes, it would be possible, with a suitable definition of cost, to determine how much money the company would spend to complete the survey.

While "Minimize Survey Cost" has a natural attribute scale, "Maximize Survey Quality" certainly does not. How can we measure achievement toward this objective? When there is no natural scale, two other possibilities exist. One is to use a different scale as a proxy. Of course, the proxy should be closely related to the original objective. For example, we might take a cue from the means-objectives network in Figure 3.4; if we could measure the intern's abilities in survey design and analysis, that might serve as a reasonable proxy for survey quality. One possibility would be to use the intern's grade point average in market research and statistics courses. Another possibility would be to ask one of the intern's instructors to provide a rating of the intern's abilities. (Of course, this latter suggestion gives the instructor the same problem that we had in the first place: how to measure the student's ability when there is no natural scale!)

The second possibility is to construct an attribute scale for measuring achievement of the objective. In the case of survey quality, we might be able to think of a number of levels in general terms. The best level might be described as follows:

Best survey quality: State-of-the-art survey. No apparent crucial issues left unaddressed. Has characteristics of the best survey projects presented at professional conferences.

On the other hand, the worst level might be:

Worst survey quality: Many issues left unanswered in designing survey. Members of the staff are aware of advances in survey design that could have been incorporated but were not. Not a presentable project.

We could identify and describe fully a number of meaningful levels that relate to survey quality. Table 3.4 shows five possible levels in order from best to worst. You can see that the detailed descriptions define what is meant by quality of the survey and

Table 3.4

A constructed scale for survey quality.

RANK:

Best. State-of-the-art survey. No apparent substantive issues left unaetdressed. Has characteristics of the best survey projects presented at professional conferences.

Better. Excellent survey but not perfect. Methodological techniques were appropriate for the project and similar to previous projects, but in some cases more up-to-date techniques are available. One substantive issue that could have been handled better. Similar to most of the survey projects presented at professional conferences.

Satisfactory. Satisfactory survey. Methodological techniques were appropriate, but superior methods exist and should have been used. Two or three unresolved substantive issues. Project could be presented at a professional conference but has characteristics that would make it less appealing than most presentations.

Worse. Although the survey results will be useful temporarily, a follow-up study must be done to refine the methodology and address substantive issues that were ignored. Occasionally similar projects are presented at conferences, but they are poorly received.

Worst. Unsatisfactory. Survey must be repeated to obtain useful results. Members of the staff are aware of advances in survey design that could have been incorporated but were not. Many substantive issues left unanswered. Not a presentable project.

how to determine whether the survey was well done. According to these defined levels, quality is judged by the extent to which the statistical and methodological techniques were up to date, whether any of the company still has unresolved questions about its consumer products, and a judgmental comparison with similar survey projects presented at professional meetings.

Constructing scales can range from straightforward to complex. The key to constructing a good scale is to identify meaningful levels, including best, worst, and intermediate, and then describe those levels in a way that fully reflects the objective under consideration. The descriptions of the levels must be elaborate enough to facilitate the measurement of the consequences. In thinking about possible results of specific choices made and particular uncertain outcomes, it should be easy to use the constructed attribute scale to specify the corresponding consequences.

The scale in Table 3.4 actually shows two complementary ways to describe a level. First is in terms of specific aspects of survey quality, in this case the methodology and the extent to which the survey successfully addressed the company's concerns about its line of products. The second way is to use a comparison approach; in this case, we compare the survey project overall with other survey projects that have been presented at professional meetings. There is nothing inherently important about the survey's presentability at a conference, but making the comparison can help to measure the level of quality relative to other publicly accessible projects.

Note also from Table 3.4 that we could have extended the fundamental-objectives hierarchy to include "Methodology" and "Address Company's Issues" as branches under the "Maximize Survey Quality" branch, as shown in Figure 3.33. How much detail is included in the hierarchy is a matter of choice, and here the principle of a requisite model comes into play. As long as the scale for "Maximize Survey Quality" can adequately capture the company's concerns regarding this

Figure 3.33
An expanded fundamental-objectives hierarchy for the summer-intern example.

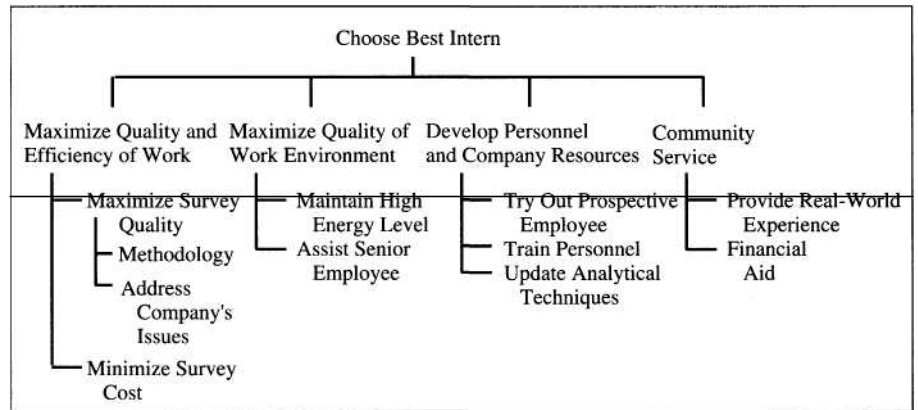


Table 3.5
 A constructed attribute scale for biological impact.

RANK:

- Best**
- Complete loss of 1.0 square mile of land that is entirely in agricultural use or is entirely urbanized; no loss of any "native" biological communities.
 - ◆ Complete loss of 1.0 square mile of primarily (75%) agricultural habitat with loss of 25% of second-growth forest; no measurable loss of wetlands or endangered-species habitat.
 - ◆ Complete loss of 1.0 square mile of land that is 50% farmed and 50% disturbed in some other way (e.g., logged or new second growth); no measurable loss of wetlands or endangered-species habitat.
 - ◆ Complete loss of 1.0 square mile of recently disturbed (for example, logged, plowed) habitat plus disturbance to surrounding previously disturbed habitat within 1.0 mile of site border; or 15% loss of wetlands or endangered-species habitat.
 - ◆ Complete loss of 1.0 square mile of land that is 50% farmed (or otherwise disturbed) and 50% mature second-growth forest or other undisturbed community; 15% loss of wetlands or endangered-species habitat.
 - ◆ Complete loss of 1.0 square mile of land that is primarily (75%) undisturbed mature "desert" community; 15% loss of wetlands or endangered-species habitat.
 - ◆ Complete loss of 1.0 square mile of mature second-growth (but not virgin) forest community; or 50% loss of big game and upland game birds; or 50% loss of wetlands and endangered-species habitat.
 - ◆ Complete loss of 1.0 square mile of mature community or 90% loss of productive wetlands and endangered-species habitat.
- Worst** ◆ Complete loss of 1.0 square mile of mature virgin forest and/or wetlands and/or endangered-species habitat.

Source: Adapted from Keeney (1992, p. 105).

Table 3.6
A constructed
attribute scale for
public attitudes.

RANK:

Best *Support.* No groups are opposed to the facility, and at least one group has organized support for the facility.

- *Neutrality.* All groups are indifferent or uninterested.
- *Controversy.* One or more groups have organized opposition, although no groups have action-oriented opposition (for example, letterwriting, protests, lawsuits). Other groups may either be neutral or support the facility.
- *Action-oriented opposition.* Exactly one group has action-oriented opposition. The other groups have organized support, indifference, or organized opposition.

Worst. *Strong action-oriented opposition.* Two or more groups have action-oriented opposition.

Source: Adapted from Keeney (1992, p. 102).

objective, then there is no need to use more detailed objectives in measuring quality. If, on the other hand, there are many different aspects of quality that are likely to vary separately depending on choices and chance outcomes, then it may be worthwhile to create a more detailed model of the objective by extending the hierarchy and developing attribute scales for the subobjectives.

Developing the ability to construct meaningful attribute scales requires practice. In addition, it is helpful to see examples of scales that have been used in various situations. We have already seen one such scale in Table 3.4 relating to the summer-in-term example. Tables 3.5 and 3.6 show two other constructed attribute scales for biological impacts and public attitudes, respectively, both in the context of selecting a site for a nuclear power generator.

Using Computers for Structuring Decisions

Decision analysis has benefited tremendously from innovations in computers and computer software. Introductions to relevant personal-computer software are included throughout this book. The rapid pace of software development, however, means that even newer programs, not mentioned here, may be available. To get up-to-date information on decision-analysis software, look for recent software reviews in publications like *OR/MS Today*, the official monthly publication of INFORMS, the Institute for Operations Research and Management Science.

Although many of the advances in decision-analysis software have had to do with high-speed calculations that allow elaborate analysis of large and complex decisions, the rise of graphical computer interfaces has permitted the development of important computer tools for structuring decision problems. We have already mentioned the program DPL, which incorporates both influence diagrams and decision trees. A variety of decision-tree programs are also available. In addition, a number of programs are available to help with value structuring. In this section we will take a

quick look at the structuring features of Logical Decisions (a multiple-objective decision program), DPL, and DATA (a decision-tree program). Sources for software described here and in later chapters are listed at the end of Chapter 17.

Logical Decisions

With Logical Decisions, a decision maker can structure a decision situation that involves multiple objectives. The starting point in this program is to create a fundamental-objectives hierarchy. Figure 3.34 shows the screen from a Logical Decisions session in which the hierarchy from the generic safety problem (Figure 3.1) is being created.

None of the existing decision-analysis programs are perfectly consistent with this textbook in the terminology that they use. Logical Decisions is no exception; you can see in Figure 3.34 that the word "Goal" is used to refer to what we call fundamental objectives, and that "Measure" refers to an attribute. The good news is that Logical Decisions permits the user to change these labels. Figure 3.35 shows the fundamental-objectives hierarchy, in Logical Decisions format, for the intern problem (Figure 3.4). You can see that the labels have been modified so that the terms "Objective" and "Attribute" appear in the figure.

Logical Decisions does not provide any facility for creating a means-objectives network. It does, however, permit the user to specify attribute scales, identify alternatives, and to indicate how each alternative is evaluated on each attribute. This is done in a matrix, shown in Figure 3.36 for the FAA bomb-detection problem (Figure 3.23), that is essentially the same as the consequence matrix discussed above for decision trees.

Figure 3.34
 Fundamental-objectives hierarchy for safety decision in logical decisions program.

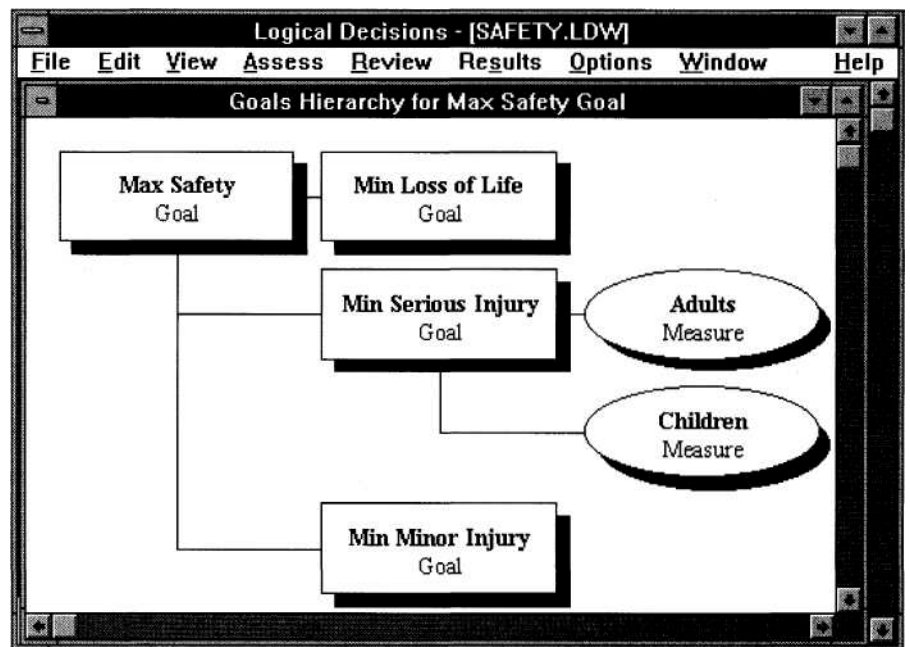


Figure 3.35
Fundamental objectives for intern decision in logical decisions.

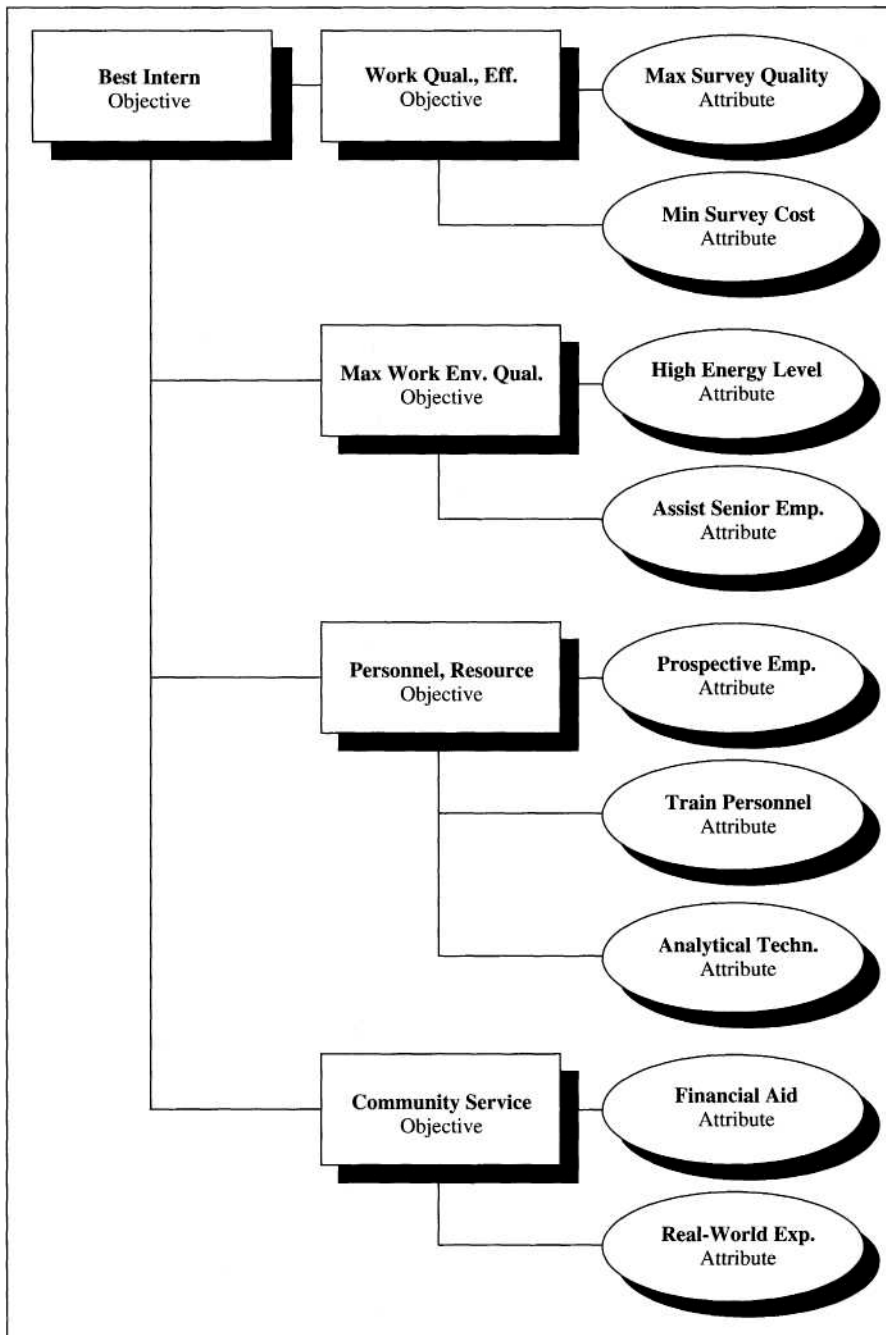


Figure 3.36
 Consequence matrix
 for FAA bomb-
 detection decision in
 logical decisions.

Logical Decisions - [FAA.LDW]				
File Edit View Assess Review Results Options Window Help				
Matrix				
Prob	99	Effectiveness [%] (Per Cent Detected) for System #1		
	Effectiveness [%]	Time (Months)	Passenger Accept	Cost (\$Million)
System #1	99	13	50	20
System #2	95	15	95	18
System #3	87	7	97	55
System #4	92	13	95	36
System #5	95	14	75	22

DPL (Decision Programming Language)

DPL is a unique program for decision analysis. At the time this is being written, it is the only program available that incorporates both decision trees and influence diagrams. Decisions can be structured in both influence-diagram and decision-tree form. In addition, DPL sports a powerful analytical engine and numerous features for analyzing decision models in a variety of ways. Several of these techniques will be discussed later in the book. Here we examine the decision-structuring environment that DPL provides.

Figure 3.37 shows DPL's influence-diagram representation of the toxic-chemicals decision (Figure 3.20), and Figure 3.38 displays the default decision tree that DPL creates from the influence diagram. Several features are apparent. For example, in both views of the decision, DPL provides a number of drawing tools for creating and structuring the decision as required; these tools are represented by the various icons in the toolbar. In addition, commands for adding, deleting, and modifying the structure reside within the menus. Second, you can see in the decision-tree representation that, in creating a default decision tree in schematic form, DPL assumes that the influence diagram represents a *symmetric* decision; that is, the same sequence of decisions and chance nodes occur regardless of the path taken.

In DPL, data regarding the specific alternatives and outcomes are entered via the influence diagram. In the decision-tree view, though, the decision maker specifies the structure of the decision more fully, specifying the sequence of the decisions and chance events. In particular, although sequence arcs can be included in the influence diagram, DPL uses these only to order the nodes in the default decision tree; it is up

Figure 3.37
*Toxic-chemical
 influence diagram in
 DPL*

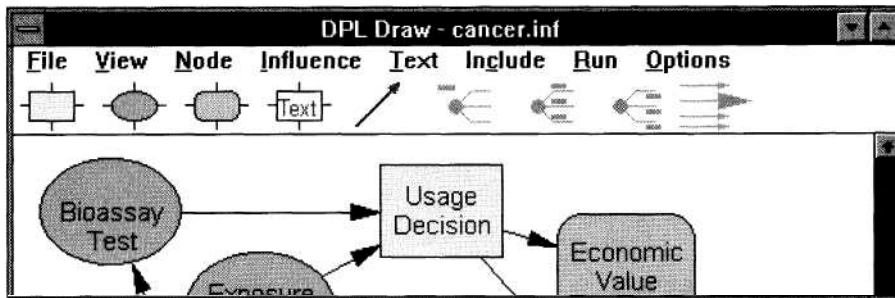
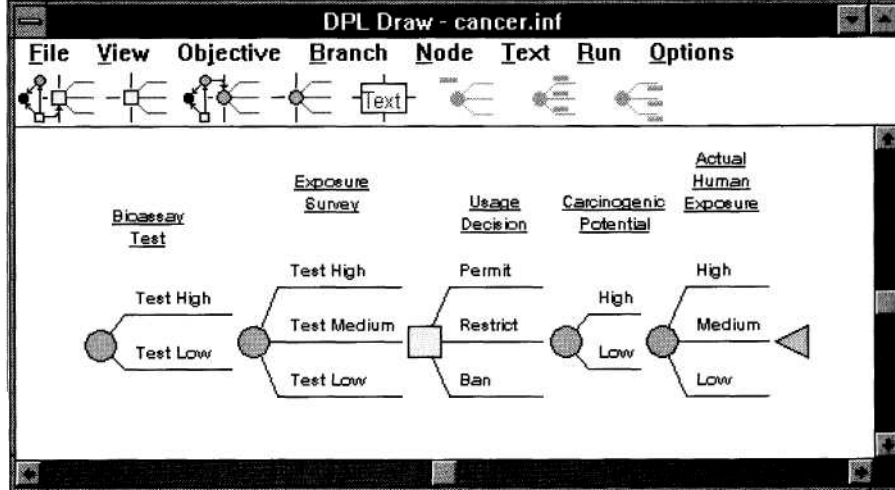


Figure 3.38
*Toxic-chemical
 decision tree in DPL.*



to the user to verify that the default sequence is correct or to modify the sequence as needed. In addition, the user can modify the default tree to create an asymmetric tree. Modeling asymmetry explicitly in the decision tree is straightforward with DPL's tools and commands. Moreover, DPL has the ability to represent decision trees in a compact way, including schematic representation where appropriate or including the same subtree at several different places in the decision structure. For example, Figure 3.39 shows the influence diagram for the hurricane-evacuation decision (Figure 3.12), and Figure 3.40 displays the corresponding decision tree as modified by the user to reflect the asymmetry in the decision.

Including cash flows and probabilities as described above is straightforward. DPL calls cash flows "gets" and "pays" because the decision maker may either get an incremental amount (of money or whatever the relevant attribute is) or pay some amount. In the analysis, the "gets" and "pays" are collapsed into net values at the ends of the branches.

DPL's treatment of arcs in an influence diagram is somewhat more elaborate than our discussion has been, but the two approaches are not inconsistent. DPL distinguishes between four different kinds of arcs, reflecting different kinds of relevance or dependence on the predecessor node. For example, the values (cash flows, or "gets" and "pays") associated with alternatives or chance outcomes can depend on the outcome of a predecessor node. Likewise, the probabilities associated with a chance node's outcomes can depend on a predecessor's outcome. The four different

Figure 3.39
DPL's influence diagram for the hurricane problem.

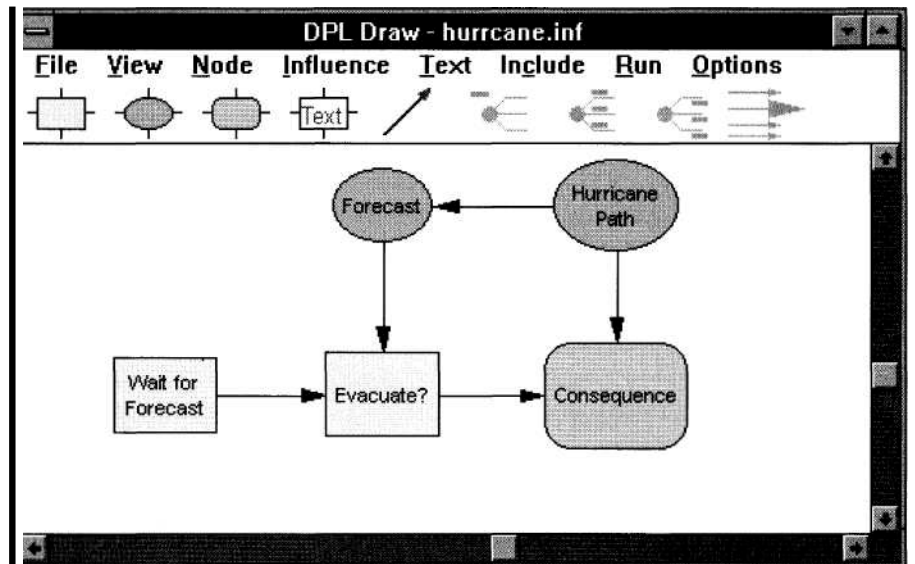
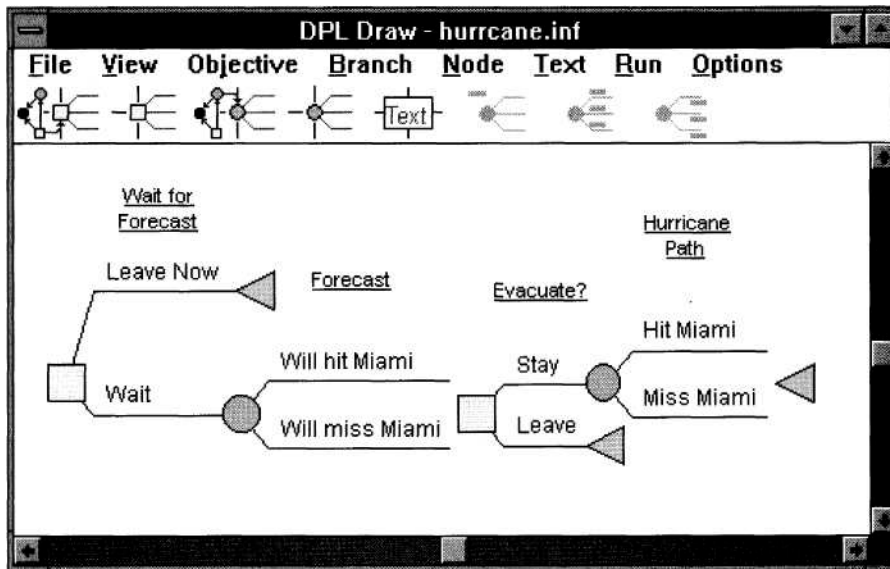


Figure 3.40
DPL's decision tree
for the hurricane
problem.



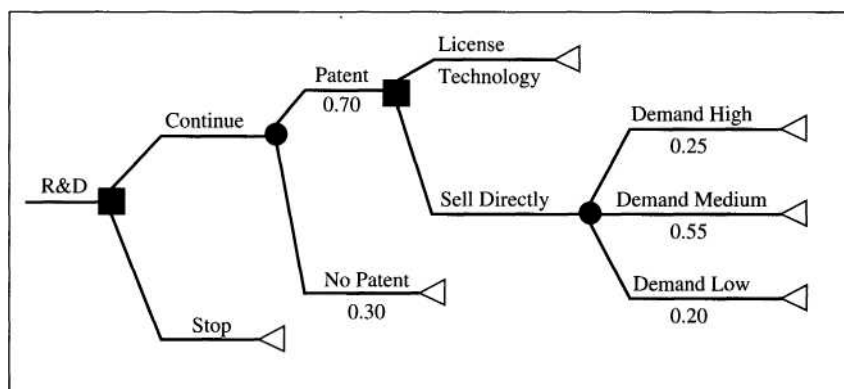
kinds of arcs represent relevance of predecessors on values, probabilities, both, or neither. The sequence arcs into decision nodes that we have defined, for example, can be represented by a DPL arc that represents neither value nor probability dependence. On the other hand, the modeler can use a different arc to specify that cash flows in a decision node differ depending on the predecessor's outcome; in this case, the arc into the decision node is something of a combination of sequence and relevance arcs. Having this flexibility for specifying different kinds of dependence is one of the features that makes DPL a powerful and efficient structuring tool.

This quick introduction to DPL barely scratches the surface of its structuring capabilities. In fact, the influence-diagram and decision-tree screens constitute the "front end" of a powerful language that DPL uses to represent decisions in a formal way. The language itself has considerable flexibility in how decisions can be represented. Other features include the ability to model multiple objectives and to interface with electronic spreadsheets.

DATA

This program provides an environment for constructing and analyzing decision trees. The graphical representation is fully consistent with our development in this textbook, with the exception that the tree includes triangles at the right-hand end of each branch. [You may have noticed that DPL does the same to represent the consequence. The use of the triangle as a termination goes back to Howard Raiffa's original *Decision Analysis* textbook (1968).] Figure 3.41 shows DATA'S version of the R&D problem (Figure 3.32). DATA makes it easy to create and edit a decision tree,

Figure 3.41
Research-and-development decision in DATA.



including adding and deleting nodes and branches and modifying information. In Figure 3.41, you can see probabilities attached to the branches. Cash flows are easily defined for variables that correspond to the various nodes but are not displayed in the tree itself.

All of the programs that we have mentioned here, as well as many others, are useful for modeling decisions. Although we have emphasized the structuring aspects of these programs, all of them have analytical capabilities. We will look at these analytical features in Chapters 4 and 5.

SUMMARY

This chapter has discussed the general process of structuring decision problems. It is impossible to overemphasize the importance of the structuring step, because it is here that one really understands the problem and all of its different aspects. We began with the process of structuring values, emphasizing the importance of identifying underlying fundamental objectives and separating those from means objectives. Fundamental objectives, structured in a hierarchy, are those things that the decision maker wants to accomplish, and means objectives, structured in a network, describe ways to accomplish the fundamental objectives.

With objectives specified, we can begin to structure a decision's specific elements. A decision maker may use both influence diagrams and decision trees as tools in the process of modeling decisions. Influence diagrams provide compact representations of decision problems while suppressing many of the details, and thus they are ideal for obtaining overviews, especially for complex problems. Influence diagrams are especially appropriate for communicating decision structure because they are easily understood by individuals with little technical background. On the other hand, decision trees display all of the minute details. Being able to see the details can be an advantage, but in complex decisions trees may be too large and "bushy" to be of much use in communicating with others.

The clarity test is used to ensure that the problem is defined well enough so that everyone can agree on the definitions of the basic decision elements, and we also discussed the specification of probabilities and cash flows at different points in the

problem. We also discussed the notion of attribute scales for measuring the extent to which fundamental objectives are accomplished, and we showed how scales can be constructed to measure achievement of those objectives that do not have natural measures. Finally, we introduced three decision-analysis programs for personal computers that are useful for structuring decisions.

EXERCISES

- 3.1 Describe in your own words the difference between a means objective and a fundamental objective. Why do we focus on coming up with attribute scales that measure accomplishment of fundamental objectives, but not means objectives? What good does it do to know what your means objectives are?
- 3.2 What are your fundamental objectives in the context of renting an apartment while attending college? What are your means objectives? Create a fundamental-objectives hierarchy and a means-objectives network.
- 3.3 In the context of renting an apartment (Exercise 3.2), some of the objectives may have natural attribute scales. Examples are minimizing rent (\$) or minimizing the distance to campus (kilometers or city blocks). But other attributes, such as ambiance, amount of light, or neighbors, have no natural scales. Construct an attribute scale with at least five different levels, ranked from best to worst, for some aspect of an apartment that is important to you but has no natural scale.
- 3.4 Before making an unsecured loan to an individual a bank orders a report on the applicant's credit history. To justify making the loan, the bank must find the applicant's credit record to be satisfactory. Describe the bank's decision. What are the bank's objectives? What risk does the bank face? What role does the credit report play? Draw an influence diagram of this situation. (*Hint:* Your influence diagram should include chance nodes for a credit report and for eventual default.) Finally, be sure to specify everything (decisions, chance events, objectives) in your model clearly enough to pass the clarity test.
- 3.5 When a movie producer decides whether to produce a major motion picture, the main question is how much revenue the movie will generate. Draw a decision tree of this situation, assuming that there is only one fundamental objective, to maximize revenue. What must be included in revenue to be sure that the clarity test is passed?
- 3.6 You have met an acquaintance for lunch, and he is worried about an upcoming meeting with his boss and some executives from his firm's headquarters. He has to outline the costs and benefits of some alternative investment strategies. He knows about both decision trees and influence diagrams but cannot decide which presentation to use. In your own words, explain to him the advantages and disadvantages of each.
- 3.7 Reframe your answer to Exercise 3.6 in terms of objectives and alternatives. That is, what are appropriate fundamental objectives to consider in the context of choosing how to present the investment information? How do decision trees and influence diagrams compare in terms of these objectives?
- 3.8 Draw the politician's decision in Figure 3.25 as an influence diagram. Include the tables showing decision alternatives, chance-event outcomes, and consequences.

- 3.9** A dapper young decision maker has just purchased a new suit for \$200. On the way out the door, the decision maker considers taking an umbrella. With the umbrella on hand, the suit will be protected in the event of rain. Without the umbrella, the suit will be ruined if it rains. On the other hand, if it does not rain, carrying the umbrella is an unnecessary inconvenience.
- Draw a decision tree of this situation.
 - Draw an influence diagram of the situation.
 - Before deciding, the decision maker considers listening to the weather forecast on the radio. Draw an influence diagram that takes into account the weather forecast.
- 3.10** When patients suffered from hemorrhagic fever, M*A*S*H doctors replaced lost sodium by administering a saline solution intravenously. However, headquarters (HQ) sent a treatment change disallowing the saline solution. With a patient in shock and near death from a disastrously low sodium level, B. J. Hunnicut wanted to administer a low-sodium-concentration saline solution as a last-ditch attempt to save the patient. Colonel Potter looked at B. J. and Hawkeye and summed up the situation. "O.K., let's get this straight. If we go by the new directive from HQ and don't administer saline to replace the sodium, our boy will die for sure. If we try B. J.'s idea, then he may survive, and we'll know how to treat the next two patients who are getting worse. If we try it and he doesn't make it, we're in trouble with HQ and may get court-martialed. I say we have no choice. Let's try it." (*Source*: "Mr. and Mrs. Who." Written by Ronny Graham, directed by Burt Metcalfe, 1980.)
- Structure the doctors' decision. What are their objectives? What risks do they face? Draw a decision tree for their decision.
- 3.11** Here is an example that provides a comparison between influence diagrams and decision trees.
- Suppose you are planning a party, and your objective is to have an enjoyable party for all the guests. An outdoor barbecue would be the best, but only if the sun shines; rain would make the barbecue terrible. On the other hand, you could plan an indoor party. This would be a good party, not as nice as an outdoor barbecue in the sunshine but better than a barbecue in the rain. Of course, it is always possible to forego the party altogether! Construct an influence diagram and a decision tree for this problem.
 - You will, naturally, consult the weather forecast, which will tell you that the weather will be either "sunny" or "rainy." The forecast is not perfect, however. If the forecast is "sunny," then sunshine is more likely than rain, but there still is a small chance that it will rain. A forecast of "rainy" implies that rain is likely, but the sun may still shine. Now draw an influence diagram for the decision, including the weather forecast. (There should be four nodes in your diagram, including one for the forecast, which will be available at the time you decide what kind of party to have, and one for the actual weather. Which direction should the arrow point between these two nodes? Why?) Now draw a decision tree for this problem. Recall that the events and decisions in a decision tree should be in chronological order.
- 3.12** The clarity test is an important issue in Exercise 3.11. The weather obviously can be somewhere between full sunshine and rain. Should you include an outcome like "cloudy"? Would it affect your satisfaction with an outdoor barbecue? How will you define rain? The National Weather Service uses the following definition: Rain has occurred if "measurable precipitation" (more than 0.004 inch) has occurred at the official rain gauge. Would this definition be suitable for your purposes? Define a set of weather outcomes that is appropriate relative to your objective of having an enjoyable party.
- 3.13** Draw the machine-replacement decision (Figure 3.10) as a decision tree.

QUESTIONS AND PROBLEMS

3.14 Modify the influence diagram in Figure 3.11 (the hurricane-forecast example) so that it contains nodes for each of the two objectives (maximize safety and minimize cost). Cost has a natural attribute scale, but how can you define safety? Construct an attribute scale that you could use to measure the degree of danger you might encounter during a hurricane.

3.15 Decision analysis can be used on itself! What do you want to accomplish in studying decision analysis? Why is decision analysis important to you? In short, what are your fundamental objectives in studying decision analysis? What are appropriate means objectives? Is your course designed in a way that is consistent with your objectives? If not, how could the course be modified to achieve your objectives?

3.16 In the spring of 1987 Gary Hart, the leading Democratic presidential candidate, told the news media that he was more than willing to have his private life scrutinized carefully. A few weeks later, the Miami Herald reported that a woman, Donna Rice, had been seen entering his Washington townhouse on a Friday evening but not leaving until Saturday evening. The result was a typical political scandal, with Hart contending that Rice had left Friday evening by a back door that the reporter on the scene was not watching. The result was that Hart's credibility as a candidate was severely damaged, thus reducing his chance of winning both the Democratic nomination and the election. The decision he had to make was whether to continue the campaign or to drop out. Compounding the issue was a heavy debt burden that was left over from his unsuccessful 1984 presidential bid.

Using both an influence diagram and a decision tree, structure Hart's decision. What is the main source of uncertainty that he faces? Are there conflicting objectives, and if so, what are they? What do you think he should have done? (He decided to drop out of the race. However, he eventually reentered, only to drop out again because of poor showings in the primary elections.)

3.17 When an amateur astronomer considers purchasing or building a telescope to view deep-sky objects (galaxies, clusters, nebulae, etc.), the three primary considerations are minimizing cost, having a stable mounting device, and maximizing the *aperture* (diameter of the main lens or mirror). The aperture is crucial because a larger aperture gathers more light. With more light, more detail can be seen in the image, and what the astronomer wants to do is to see the image as clearly as possible. As an example, many small telescopes have lens or mirrors up to 8 inches in diameter. Larger amateur telescopes use concave mirrors ranging from 10 to 16 inches in diameter. Some amateurs grind their own mirrors as large as 40 inches.

Saving money is important, of course, because the less spent on the telescope, the more can be spent on accessories (eyepieces, star charts, computer-based astronomy programs, warm clothing, flashlights, and so on) to make viewing as easy and comfortable as possible. Money might also be spent on an observatory to house a large telescope or on trips away from the city (to avoid the light pollution of city skies and thus to see images more clearly).

Finally, a third issue is the way the telescope is mounted. First, the mount should be very stable, keeping the telescope perfectly still. Any vibrations will show up dramatically in the highly magnified image, thus reducing the quality of the image and the detail that can be seen. The mount should also allow for easy and smooth movement of the telescope to view any part of the sky. Finally, if the astronomer wants to use the telescope to take pho-

tographs of the sky (astrophotos), it is important that the mount includes some sort of tracking device to keep the telescope pointing at the same point in the sky as the earth rotates beneath it.

Based on this description, what are the amateur astronomer's fundamental objectives in choosing a telescope? What are the means objectives? Structure these objectives into a fundamental-objectives hierarchy and a means-objectives network. (*Hint:* If you feel the need for more information, look in your library for recent issues of *Astronomy* magazine or *Sky and Telescope*, two publications for amateur astronomers.)

- 3.18** Consider the following situations that involve multiple objectives:
- a** Suppose you want to go out for dinner. What are your fundamental objectives? Create a fundamental-objectives hierarchy.
 - b** Suppose you are trying to decide where to go for a trip over spring break. What are your fundamental objectives? What are your means objectives?
 - c** You are about to become a parent (surprise!), and you have to choose a name for your child. What are important objectives to consider in choosing a name?
 - d** Think of any other situation where choices involve multiple objectives. Create a fundamental-objectives hierarchy and a means-objectives network.
- 3.19** Thinking about fundamental objectives and means objectives is relatively easy when the decision context is narrow (buying a telescope, renting an apartment, choosing a restaurant for dinner). But when you start thinking about your strategic objectives—objectives in the context of what you choose to do with your life or your career—the process becomes more difficult. Spend some time thinking about your fundamental strategic objectives. What do you want to accomplish in your life or in your career? Why are these objectives important? Try to create a fundamental-objectives hierarchy and a means-objectives network for yourself.
- If you succeed in this problem, you will have achieved a deeper level of self-knowledge than most people have, regardless of whether they use decision analysis. That knowledge can be of great help to you in making important decisions, but you should revisit your fundamental objectives from time to time; they might change!
- 3.20** Occasionally a decision is sensitive to the way it is structured. The following problem shows that leaving out an important part of the problem can affect the way we view the situation.
- a** Imagine that a close friend has been diagnosed with heart disease. The physician recommends bypass surgery. The surgery should solve the problem. When asked about the risks, the physician replies that a few individuals die during the operation, but most recover and the surgery is a complete success. Thus, your friend can (most likely) anticipate a longer and healthier life after the surgery. Without surgery, your friend will have a shorter and gradually deteriorating life. Assuming that your friend's objective is to maximize the quality of her life, diagram this decision with both an influence diagram and a decision tree.
 - b** Suppose now that your friend obtains a second opinion. The second physician suggests that there is a third possible outcome: Complications from surgery can develop which will require long and painful treatment. If this happens, the eventual outcome can be either a full recovery, partial recovery (restricted to a wheelchair until death), or death within a few months. How does this change the decision tree and influence diagram that you created in part a? Draw the decision tree and influence diagram that

represent the situation after hearing from both physicians. Given this new structure, does surgery look more or less positive than it did in part a? [For more discussion of this problem, see von Winterfeldt and Edwards (1986, pp. 8-14).]

- c Construct an attribute scale for the patient's quality of life. Be sure to include levels that relate to all of the possible outcomes from surgery.
- 3.21** Create an influence diagram and a decision tree for the difficult decision problem that you described in Problem 1.9. What are your objectives? Construct attribute scales if necessary. Be sure that all aspects of your decision model pass the clarity test.
- 3.22** To be, or not to be, that is the question:
 Whether 'tis nobler in the mind to suffer
 The slings and arrows of outrageous fortune
 Or to take arms against a sea of troubles,
 And by opposing end them. To die — to sleep —
 No more; and by a sleep to say we end
 The heartache, and the thousand natural shocks
 That flesh is heir to. 'Tis a consummation
 Devoutly to be wished. To die — to sleep.
 To sleep — perchance to dream: ay, there's the rub!
 For in that sleep of death what dreams may come
 When we have shuffled off this mortal coil,
 Must give us pause. There's the respect
 That makes calamity of so long life.
 For who would bear the whips and scorns of time,
 the oppressor's wrong, the proud man's contumely,
 The pangs of despised love, the law's delay,
 The insolence of office, and the spurns
 That patient merit of the unworthy takes,
 When he himself might his quietus make
 With a bare bodkin? Who would these fardels bear,
 To grunt and sweat under a weary life,
 But that the dread of something after death —
 The undiscovered country, from whose bourn
 No traveller returns — puzzles the will,
 And makes us rather bear those ills we have
 Than fly to others that we know not of?

— *Hamlet*, Act III, Scene 1

Describe Hamlet's decision. What are his choices? What risk does he perceive? Construct a decision tree for Hamlet.

- 3.23** On July 3, 1988, the USS *Vincennes* was engaged in combat in the Persian Gulf. On the radar screen a blip appeared that signified an incoming aircraft. After repeatedly asking

the aircraft to identify itself with no success, it appeared that the aircraft might be a hostile Iranian F-14 attacking the *Vincennes*. Captain Will Rogers had little time to make his decision. Should he issue the command to launch a missile and destroy the plane? Or should he wait for positive identification? If he waited too long and the plane was indeed hostile, then it might be impossible to avert the attack and danger to his crew.

Captain Rogers issued the command, and the aircraft was destroyed. It was reported to be an Iranian Airbus airliner carrying 290 people. There were no survivors.

What are Captain Rogers's fundamental objectives? What risks does he face? Draw a decision tree representing his decision.

- 3.24** Reconsider the research-and-development decision in Figure 3.32. If you decide to continue the project, you will have to come up with the \$2 million this year (Year 1). Then there will be a year of waiting (Year 2) before you know if the patent is granted. If you decide to license the technology, you would receive the \$25 million distributed as \$5 million per year beginning in Year 3. On the other hand, if you decide to sell the product directly, you will have to invest \$5 million in each of Years 3 and 4 (to make up the total investment of \$10 million). Your net proceeds from selling the product, then, would be evenly distributed over Years 5 through 9.

Assuming an interest rate of 15%, calculate the NPV at the end of each branch of the decision tree.

- 3.25** When you purchase a car, you may consider buying a brand-new car or a used one. A fundamental trade-off in this case is whether you pay repair bills (uncertain at the time you buy the car) or make loan payments that are certain.

Consider two cars, a new one that costs \$15,000 and a used one with 75,000 miles for \$5500. Let us assume that your current car's value and available cash amount to \$5500, so you could purchase the used car outright or make a down payment of \$5500 on the new car. Your credit union is willing to give you a five-year, 10% loan on the \$9500 difference if you buy the new car; this loan will require monthly payments of \$201.85 per month for five years. Maintenance costs are expected to be \$100 for the first year and \$300 per year for the second and third years.

After taking the used car to your mechanic for an evaluation, you learn the following. First, the car needs some minor repairs within the next few months, including a new battery, work on the suspension and steering mechanism, and replacement of the belt that drives the water pump. Your mechanic has estimated that these repairs will cost \$150.00. Considering the amount you drive, the tires will last another year but will have to be replaced next year for about \$200. Beyond that, the mechanic warns you that the cooling system (radiator and hoses) may need to be repaired or replaced this year or next and that the brake system may need work. These and other repairs that an older car may require could lead you to pay anywhere from \$500 to \$2500 in each of the next three years. If you are lucky, the repair bills will be low or will come later. But you could end up paying a lot of money when you least expect it.

Draw a decision tree for this problem. To simplify it, look at the situation on a yearly basis for three years. If you buy the new car, you can anticipate cash outflows of $12 \times \$201.85 = \2422.20 plus maintenance costs. For the used car, some of the repair costs are known (immediate repairs this year, tires next year), but we must model the uncertainty associated with the rest. In addition to the known repairs, assume that in each year there is a 20% chance that these uncertain repairs will be \$500, a 20% chance they will be \$2500, and a 60% chance they will be \$1500. (*Hint:* You need 3 chance nodes: one for each year!)

To even the comparison of the two cars, we must also consider their values after three years. If you buy the new car, it will be worth approximately \$8000, and you will still owe \$4374. Thus, its net salvage value will be \$3626. On the other hand, you would own the used car free and clear (assuming you can keep up with the repair bills!), and it would be worth approximately \$2000.

Include all of the probabilities and cash flows (outflows until the last branch, then an inflow to represent the car's salvage value) in your decision tree. Calculate the net values at the ends of the branches.

C A S E S T U D I E S

COLD FUSION

On March 23, 1989, Stanley Pons and Martin Fleischmann announced in a press conference at the University of Utah that they had succeeded in creating a small-scale nuclear fusion reaction in a simple apparatus at room temperature. They called the process "cold fusion." Although many details were missing from their description of the experiment, their claim inspired thoughts of a cheap and limitless energy supply, the raw material for which would be ocean water. The entire structure of the world economy potentially would change.

For a variety of reasons, Pons and Fleischmann were reluctant to reveal all of the details of their experiment. If their process really were producing energy from a fusion reaction, and any commercial potential existed, then they could become quite wealthy. The state of Utah also considered the economic possibilities and even went so far as to approve \$5 million to support cold-fusion research. Congressman Wayne Owens from Utah introduced a bill in the U.S. House of Representatives requesting \$100 million to develop a national cold-fusion research center at the University of Utah campus.

But were the results correct? Experimentalists around the world attempted to replicate Pons and Fleischmann's results. Some reported success, while many others did not. A team at Texas A&M claimed to have detected neutrons, the telltale sign of fusion. Other teams detected excess heat as had Pons and Fleischmann. Many experiments failed to confirm a fusion reaction, however, and several physicists claimed that the Utah pair simply had made mistakes in their measurements.

Questions

- 1 Consider the problem that a member of the U.S. Congress would have in deciding whether to vote for Congressman Owens's bill. What alternatives are available? What are the key uncertainties? What objectives might the Congress member consider? Structure the decision problem using an influence diagram and a decision tree.
- 2 A key part of the experimental apparatus was a core of palladium, a rare metal. Consider a speculator who is thinking of investing in palladium in response to the

announcement. Structure the investor's decision. How does it compare to the decision in Question 1 ?

Sources: "Fusion in a Bottle: Miracle or Mistake," *Business Week*, May 8, 1989, pp. 100-110; "The Race for Fusion," *Newsweek*, May 8, 1989, pp. 49-54.

PRESCRIBED FIRE

Using fire in forest management sounds contradictory. Prescribed fire, however, is an important tool for foresters, and a recent article describes how decision analysis is used to decide when, where, and what to burn. In one example, a number of areas in the Tahoe National Forest in California had been logged and were being prepared for replanting. Preparation included prescribed burning, and two possible treatments were available: burning the slash as it lay on the ground, or "yarding of unmerchantable material" (YUM) prior to burning. The latter treatment involves using heavy equipment to pile the slash. YUM reduces the difficulty of controlling the burn but costs an additional \$100 per acre. In deciding between the two treatments, two uncertainties were considered critical. The first was how the fire would behave under each scenario. For example, the fire could be fully successful, problems could arise which could be controlled eventually, or the fire could escape, entailing considerable losses. Second, if problems developed, they could result in high, low, or medium costs.

Questions

- 1 What do you think the U.S. Forest Service's objectives should be in this decision? In the article, only one objective was considered, minimizing cost (including costs associated with an escaped fire and the damage it might do). Do you think this is a reasonable criterion for the Forest Service to use? Why or why not?
- 2 Develop an influence diagram and a decision tree for this situation. What roles do the two diagrams play in helping to understand and communicate the structure of this decision?

Source: D. Cohan, S. Haas, D. Radloff, and R. Yancik (1984) "Using Fire in Forest Management: Decision Making under Uncertainty." *Interfaces*, 14, 8-19.

THE SS KUNIANG

In the early 1980s, New England Electric System (NEES), was deciding how much to bid for the salvage rights to a grounded ship, the SS *Kuniang*. If the bid

were successful, the ship could be repaired and fitted out to haul coal for its power-generation stations. The value of doing so, however, depended on the outcome of a Coast Guard judgment about the salvage value of the ship. The Coast Guard's judgment involved an obscure law regarding domestic shipping in coastal waters. If the judgment indicated a low salvage value, then NEES would be able to use the ship for its shipping needs. If the judgment were high, the ship would be considered ineligible for domestic shipping use unless a considerable amount of money was spent in fitting her with fancy equipment. The Coast Guard's judgment would not be known until after the winning bid was chosen, so there was considerable risk associated with actually buying the ship as a result of submitting the winning bid. If the bid failed, the alternatives included purchasing a new ship for \$18 million or a tug barge combination for \$15 million. One of the major issues was that the higher the bid, the more likely that NEES would win. NEES judged that a bid of \$3 million would definitely not win, whereas a bid of \$10 million definitely would win. Any bid in between was possible.

Questions

- 1 Draw an influence diagram and a decision tree for NEES's decision.
- 2 What roles do the two diagrams play in helping to understand and communicate the structure of this decision? Do you think one representation is more appropriate than the other? Why?

Source: David E. Bell, (1984) "Bidding for the SS *Kuniang*." *Interfaces*, 14, 17-23.

REFERENCES

Decision structuring as a topic of discussion and research is relatively new. Traditionally the focus has been on modeling uncertainty and preferences and solution procedures for specific kinds of problems. Recent discussions of structuring include von Winterfeldt and Edwards (1986, Chapter 2), Humphreys and Wisudha (1987), and Keller and Ho (1989).

The process of identifying and structuring one's objectives comes from Keeney's (1992) *Value-Focused Thinking*. Although the idea of specifying one's objectives clearly as part of the decision process has been accepted for years, Keeney has made this part of decision structuring very explicit. Value-focused thinking captures the ultimate in common sense; if you know what you want to accomplish, you will be able to make choices that help you accomplish those things. Thus, Keeney advocates focusing on values and objectives first, before considering your alternatives. For a more compact description of value-focused thinking, see Keeney (1994).

Relatively speaking, influence diagrams are brand-new on the decision-analysis circuit. Developed by Strategic Decisions Group as a consulting aid in the late seventies, they first appeared in the decision-analysis literature in Howard and Matheson (1984). Bodily (1985) presents an overview of influence diagrams. For more technical details, consult Shachter (1986, 1988) and Oliver and Smith (1989).

The idea of representing a decision with a network has spawned a variety of different approaches beyond influence diagrams. Two in particular are valuation networks (Shenoy, 1992) and sequential decision diagrams (Covaliu and Oliver, 1995). A recent

overview of influence diagrams and related network representations of decisions can be found in Matzkevich and Abramson (1995).

Decision trees, on the other hand, have been part of the decision-analysis tool kit since the discipline's inception. The textbooks by Holloway (1979) and Raiffa (1968) provide extensive modeling using decision trees. This chapter's discussion of basic decision trees draws heavily from Behn and Vaupel's (1982) typology of decisions.

The clarity test is another consulting aid invented by Ron Howard and his associates. It is discussed in Howard (1988).

- Behn, R. D., and J. D. Vaupel (1982) *Quick Analysis for Busy Decision Makers*. New York: Basic Books.
- Bodily, S. E. (1985) *Modern Decision Making*, New York: McGraw-Hill.
- Covaliu, Z., and R. Oliver (1995) "Representation and Solution of Decision Problems Using Sequential Decision Diagrams." *Management Science*, 41, in press.
- Holloway, C. A. (1979) *Decision Making under Uncertainty: Models and Choices*. Englewood Cliffs, NJ: Prentice-Hall.
- Howard, R. A. (1988) "Decision Analysis: Practice and Promise." *Management Science*, 34, 679-695.
- Howard, R. A., and J. E. Matheson (1984) "Influence Diagrams." In R. Howard and J. Matheson (eds.) *The Principles and Applications of Decision Analysis, Vol. II*, pp. 719-762. Palo Alto, CA: Strategic Decisions Group.
- Humphreys, P., and A. Wisudha (1987) "Methods and Tools for Structuring and Analyzing Decision Problems: A Catalogue and Review." Technical Report 87-1. London: Decision Analysis Unit, London School of Economics and Political Science.
- Keeney, R. L. (1992) *Value-Focused Thinking*. Cambridge, MA: Harvard University Press.
- Keeney, R. L. (1994) "Creativity in Decision Making with Value-Focused Thinking." *Sloan Management Review*, Summer, 33—41.
- Keller, L. R., and J. L. Ho (1989) "Decision Problem Structuring." In A. P. Sage (ed.) *Concise Encyclopedia of Information Processing in Systems and Organizations*. Oxford, England: Pergamon Press.
- Matzkevich, L, and B. Abramson (1995) "Decision-Analytic Networks in Artificial Intelligence." *Management Science*, 41, 1-22.
- Oliver, R. M., and J. Q. Smith (1989) *Influence Diagrams, Belief Nets and Decision Analysis* (Proceedings of an International Conference 1988, Berkeley). New York: Wiley.
- Raiffa, H. (1968) *Decision Analysis*. Reading, MA: Addison-Wesley.
- Shachter, R. (1986) "Evaluating Influence Diagrams." *Operations Research*, 34, 871-882.
- Shachter, R. (1988) "Probabilistic Inference and Influence Diagrams." *Operations Research*, 36, 389-604.
- Shenoy, P. (1992) "Valuation-Based Systems for Bayesian Decision Analysis." *Operations Research*, 40, 463-484.
- Ulvila, J., and R. B. Brown (1982) "Decision Analysis Comes of Age." *Harvard Business Review*, Sept-Oct, 130-141.
- von Winterfeldt, D., and W. Edwards (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press.

EPILOGUE

Toxic Chemicals The trade-off between economic value and cancer cost can be very complicated and lead to difficult decisions, especially when a widely used substance is found to be carcinogenic. Imposing an immediate ban can have extensive economic consequences. Asbestos is an excellent example of the problem. This material has been in use since Roman times and was used extensively after World War II. However, pioneering research by Dr. Irving Selikoff of the Mt. Sinai School of Medicine showed that breathing asbestos particles can cause lung cancer. This caused the EPA to list it as a hazardous air pollutant in 1972. In 1978, the EPA imposed further restrictions and banned spray-on asbestos insulation. Finally, in the summer of 1989 the EPA announced a plan that would result in an almost total ban of the substance by the year 1996. (*Sources:* "U.S. Orders Virtual Ban on Asbestos." *Los Angeles Times*, July 7, 1989; "Asbestos Widely Used Until Researcher's Warning," *The Associated Press*, July 7, 1989.)

Cold Fusion At a conference in Santa Fe, New Mexico, at the end of May 1989, Pons and Fleischmann's results were discussed by scientists from around the world. After many careful attempts by the best experimentalists in the world, no consensus was reached. Many researchers reported observing excess heat, while others observed neutrons. Many had observed nothing. With no agreement, research continued.

Over the next year, many labs attempted to replicate Pons and Fleischmann's experiments. The most thorough attempts were made at Caltech and MIT, and both failed to find evidence for a fusion reaction. In what appeared to be the death blow, exactly one year later the journal *Nature* published an article reporting work by Michael Salamon of the University of Utah. Using the electrolytic cells of Pons and Fleischmann, and observing them for several weeks, still no evidence of fusion was observed. In the same issue, *Nature* editor David Lindley wrote an editorial that essentially was an epitaph for cold fusion. In addition, two books by scientist John Huizenga (*Cold Fusion: The Scientific Fiasco of the Century*. Rochester, NY: University of Rochester Press, 1992) and journalist Gary Taubes (*Bad Science: The Short Life and Weird Times of Cold Fusion*. New York: Random House, 1993) have attempted to close the door definitively on cold fusion.

Surprisingly, though, the controversy continues. Although the top-level scientific journals no longer publish their articles, cold-fusion experimenters from around the world continue to hold conferences to report their results, and evidence is growing that some unusual and poorly understood phenomenon is occurring and can be reproduced in carefully controlled laboratory conditions. EPRI (the Electric Power Research Institute) has provided funding for cold-fusion research for several years. In its May/June 1994 cover story, *Technology Review* summarized the collected evidence relating to cold fusion and possible explanations — none consistent with conventional physical theory — of the phenomenon. Undoubtedly, research will continue for some time. Eventually the experimental effects will be confirmed and explained, or the entire enterprise will be debunked for good! (*Sources:* David Lindley (1990) "The Embarrassment of Cold Fusion." *Nature*, 344, 375-376; Robert Pool (1989) "Cold Fusion: End of Act I." *Science*, 244; 1039-1040; Edmund Storms (1994) "Warming Up to Cold Fusion." *Technology Review*, May/June, 20-29.)

CHAPTER 4

Making Choices

In this chapter, we will learn how to use the details in a structured problem to find a preferred alternative. "Using the details" typically means analysis: making calculations, creating graphs, and examining the results so as to gain insight into the decision. We will see that the kinds of calculations we make are essentially the same in solving decision trees and influence diagrams. We also introduce risk profiles and dominance considerations, ways to make decisions without doing all those calculations.

We begin by studying the analysis of decision models that involve only one objective or attribute. Although most of the examples we give use money as the attribute, it could be anything that can be measured as discussed in Chapter 3. After discussing calculation of expected values and the use of risk profiles for single-attribute decisions, we turn to decisions with multiple attributes and present some simple analytical approaches. The chapter concludes with a discussion of software for doing decision-analysis calculations on personal computers.

Our main example for this chapter is from the famous Texaco-Pennzoil court case.

TEXACO VERSUS PENNZOIL

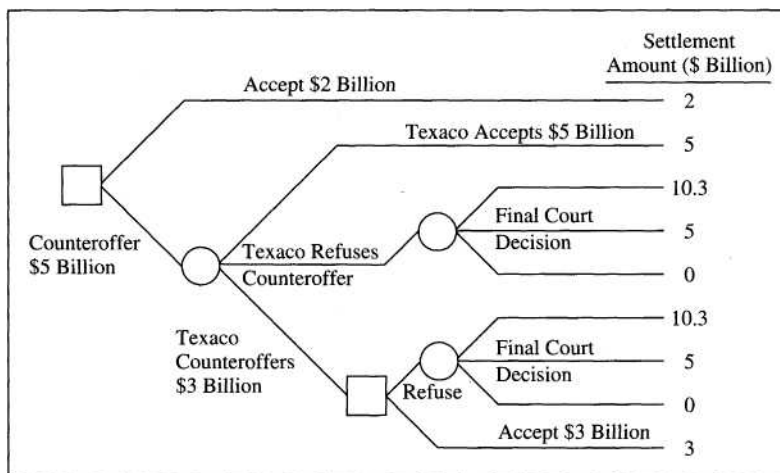
In early 1984, Pennzoil and Getty Oil agreed to the terms of a merger. But before any formal documents could be signed, Texaco offered Getty a substantially better price, and Gordon Getty, who controlled most of the Getty stock, reneged on the Pennzoil deal and sold to Texaco. Naturally, Pennzoil felt as if it had been dealt with unfairly and immediately filed a lawsuit against Texaco alleging that Texaco

had interfered illegally in the Pennzoil-Getty negotiations. Pennzoil won the case; in late 1985, it was awarded \$11.1 billion, the largest judgment ever in the United States at that time. A Texas appeals court reduced the judgment by \$2 billion, but interest and penalties drove the total back up to \$10.3 billion. James Kinnear, Texaco's chief executive officer, had said that Texaco would file for bankruptcy if Pennzoil obtained court permission to secure the judgment by filing liens against Texaco's assets. Furthermore, Kinnear had promised to fight the case all the way to the U.S. Supreme Court if necessary, arguing in part that Pennzoil had not followed Security and Exchange Commission regulations in its negotiations with Getty. In April 1987, just before Pennzoil began to file the liens, Texaco offered to pay Pennzoil \$2 billion to settle the entire case. Hugh Liedtke, chairman of Pennzoil, indicated that his advisors were telling him that a settlement between \$3 and \$5 billion would be fair.

What do you think Liedtke (pronounced "lid-key") should do? Should he accept the offer of \$2 billion, or should he refuse and make a firm counteroffer? If he refuses the sure \$2 billion, then he faces a risky situation. Texaco might agree to pay \$5 billion, a reasonable amount in Liedtke's mind. If he counteroffered \$5 billion as a settlement amount, perhaps Texaco would counter with \$3 billion or simply pursue further appeals. Figure 4.1 is a decision tree that shows a simplified version of Liedtke's problem.

The decision tree in Figure 4.1 is simplified in a number of ways. First, we assume that Liedtke has only one fundamental objective: maximizing the amount of the settlement. No other objectives need be considered. Also, Liedtke has a more varied set of decision alternatives than those shown. He could counteroffer a variety of possible values in the initial decision, and in the second decision, he could counteroffer some amount between \$3 and \$5 billion. Likewise, Texaco's counteroffer, if it makes one, need not be exactly \$3 billion. The outcome of the final court decision

Figure 4.1
Hugh Liedtke's
decision in the Texaco-
Pennzoil affair.



could be anything between zero and the current judgment of \$10.3 billion. Finally, we have not included in our model of the decision anything regarding Texaco's option of filing for bankruptcy.

Why all of the simplifications? A straightforward answer (which just happens to have some validity) is that for our purposes in this chapter we need a relatively simple decision tree to work with. But this is just a pedagogical reason. If we were to try to analyze Liedtke's problem in all of its glory, how much detail should be included? As you now realize, all of the relevant information should be included, and the model should be constructed in a way that makes it easy to analyze. Does our representation accomplish this? Let us consider the following points.

- 1 *Liedtke's objective.* Certainly maximizing the amount of the settlement is a valid objective. The question is whether other objectives, such as minimizing attorney fees or improving Pennzoil's public image, might also be important. Although Liedtke may have other objectives, the fact that the settlement can range all the way from zero to \$10.3 billion suggests that this objective will swamp any other concerns.
- 2 *Liedtke's initial counteroffer.* The counteroffer of \$5 billion could be replaced by an offer for another amount, and then the decision tree reanalyzed. Different amounts may change the chance of Texaco accepting the counteroffer. At any rate, other possible counteroffers are easily dealt with.
- 3 *Liedtke's second counteroffer.* Other possible offers could be built into the tree, leading to a Texaco decision to accept, reject, or counter. The reason for leaving these out reflects an impression from the media accounts (especially *Fortune*, May 11, 1987, pp. 50-58) that Kinnear and Liedtke were extremely tough negotiators and that further negotiations were highly unlikely.
- 4 *Texaco's counteroffer.* The \$3 billion counteroffer could be replaced by a fan representing a range of possible counteroffers. It would be necessary to find a "break-even" point, above which Liedtke would accept the offer and below which he would refuse. Another approach would be to replace the \$3 billion value with other values, recomputing the tree each time. Thus, we have a variety of ways to deal with this issue.
- 5 *The final court decision.* We could include more branches, representing additional possible outcomes, or we could replace the three branches with a fan representing a range of possible outcomes. For a first-cut approximation, the possible outcomes we have chosen do a reasonably good job of capturing the uncertainty inherent in the court outcome.
- 6 *Texaco's bankruptcy option.* A detail left out of the case is that Texaco's net worth is much more than the \$10.3 billion judgment. Thus, even if Texaco does file for bankruptcy, Pennzoil probably would still be able to collect. In reality, negotiations can continue even if Texaco has filed for bankruptcy; the purpose of filing is to protect the company from creditors seizing assets while the company proposes a financial reorganization plan. In fact, this is exactly what Texaco needs to do in

order to figure out a way to deal with Pennzoil's claims. In terms of Liedtke's options, however, whether Texaco files for bankruptcy appears to have no impact.

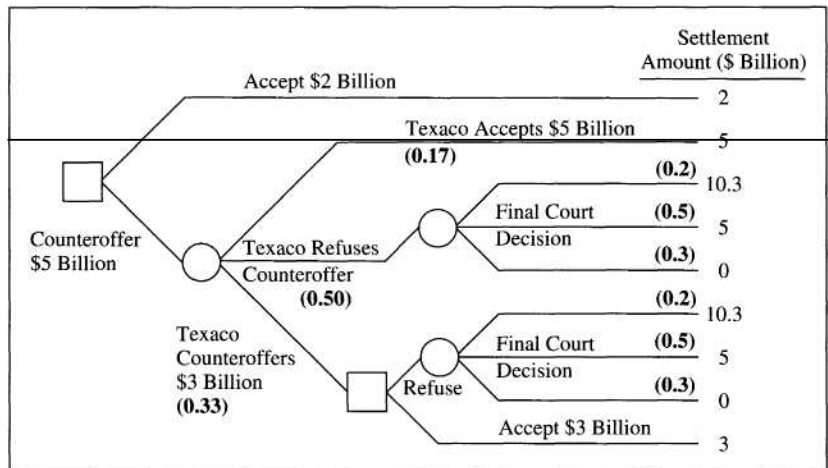
The purpose of this digression has been to explore the extent to which our structure for Liedtke's problem is requisite in the sense of Chapter 1. The points above suggest that the main issues in the problem have been represented in the problem. While it may be necessary to rework the analysis with slightly different numbers or structure later, the structure in Figure 4.1 should be adequate for a first analysis. The objective is to develop a representation of the problem that captures the essential features of the problem so that the ensuing analysis will provide the decision maker with insight and understanding.

One small detail remains before we can solve the decision tree. We need to specify the chances associated with Texaco's possible reactions to the \$5 billion counteroffer, and we also need to assess the chances of the various court awards. The probabilities that we assign to the outcome branches in the tree should reflect Liedtke's beliefs about the uncertain events that he faces. For this reason, any numbers that we include to represent these beliefs should be based on what Liedtke has to say about the matter or on information from individuals whose judgments in this matter he would trust. For our purposes, imagine overhearing a conversation between Liedtke and his advisors. Here are some of the issues they might raise:

- Given the tough negotiating stance of the two executives, it could be an even chance (50%) that Texaco will refuse to negotiate further. If Texaco does not refuse, then what? What are the chances that Texaco would accept a \$5 billion counteroffer? How likely is this outcome compared to the \$3 billion counteroffer from Texaco? Liedtke and his advisors might figure that a counteroffer of \$3 billion from Texaco is about twice as likely as Texaco accepting the \$5 billion. Thus, because there is already a 50% chance of refusal, there must be a 33% chance of a Texaco counteroffer and a 17% chance of Texaco accepting \$5 billion.
- What are the probabilities associated with the final court decision? In the *Fortune* article cited above, Liedtke is said to admit that Texaco could win its case, leaving Pennzoil with nothing but lawyer bills. Thus, there is a significant possibility that the outcome would be zero. Given the strength of Pennzoil's case so far, there is also a good chance that the court will uphold the judgment as it stands. Finally, the possibility exists that the judgment could be reduced somewhat (to \$5 billion in our model). Let us assume that Liedtke and his advisors agree that there is a 20% chance that the court will award the entire \$10.3 billion and a slightly larger, or 30%, chance that the award will be zero. Thus, there must be a 50% chance of an award of \$5 billion.

Figure 4.2 shows the decision tree with these chances included. The chances have been written in terms of probabilities rather than percentages.

Figure 4.2
 Hugh Liedtke's
 decision tree with
 chances (probabilities)
 included.



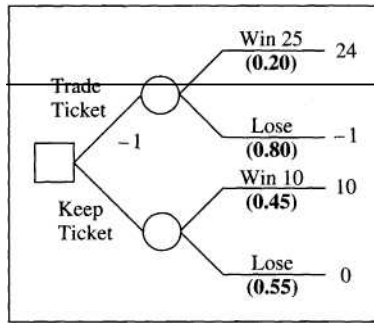
Decision Trees and Expected Monetary Value

One way to choose among risky alternatives is to pick the alternative with the highest *expected value* (EV). When the decision's consequences involve only money, we can calculate the *expected monetary value* (EMV). Finding EMVs when using decision trees is called "folding back the tree" for reasons that will become obvious. (The procedure is called "rolling back" in some texts.) We start at the endpoints of the branches on the far right-hand side and move to the left, (1) calculating expected values (to be defined momentarily) when we encounter a chance node, or (2) choosing the branch with the highest value or expected value when we encounter a decision node. These instructions sound rather cryptic. It is much easier to understand the procedure through a few examples. We will start with a simple example, the double-risk dilemma shown in Figure 4.3.

Recall that a double-risk dilemma is a matter of choosing between two risky alternatives. The situation is one in which you have a ticket that will let you participate in a game of chance (a lottery) that will pay off \$10 with a 45% chance, and nothing with a 55% chance. Your friend has a ticket to a different lottery that has a 20% chance of paying \$25 and an 80% chance of paying nothing. Your friend has offered to let you have his ticket if you will give him your ticket plus one dollar. Should you agree to the trade and play to win \$25, or should you keep your ticket and have a better chance of winning \$10?

Figure 4.3 displays your decision situation. In particular, notice that the dollar consequences at the ends of the branches are the net values as discussed in Chapter 3. Thus, if you trade tickets and win, you will have gained a net amount of \$24, having paid one dollar to your friend.

Figure 4.3
A double-risk dilemma.



To solve the decision tree using EMV, begin by calculating the expected value of keeping the ticket and playing for \$10. This expected value is simply the weighted average of the possible outcomes of the lottery, the weights being the chances with which the outcomes occur. The calculations are

$$\begin{aligned} \text{EMV(Keep Ticket)} &= 0.45(10) + 0.55(0) \\ &= \$4.5 \end{aligned}$$


One interpretation of this EMV is that playing this lottery many times would yield on average of approximately \$4.50 per game. Calculating EMV for trading tickets gives

$$\begin{aligned} \text{EMV(Trade Ticket)} &= 0.20(24) + 0.80(-1) \\ &= \$4 \end{aligned}$$

Now we can replace the chance nodes in the decision tree with their expected values, as shown in Figure 4.4. Finally, choosing between trading and keeping the ticket amounts to choosing the branch with the highest expected value. The double slash through the "Trade Ticket" branch indicates that this branch would not be chosen.

This simple example is only a warm-up exercise. Now let us see how the solution procedure works when we have a more complicated decision problem. Consider Hugh Liedtke's situation as diagrammed in Figure 4.2. Our strategy, as indicated, will be to work from the right-hand side of the tree. First, we will calculate the expected value of the final court decision. The second step will be to decide whether it is better for Liedtke to accept a \$3 billion counteroffer from Texaco or to refuse and take a chance on the final court decision. We will do this by comparing the expected value of the judgment with the sure \$3 billion. The third step will be to calculate the

Figure 4.4
Replacing chance nodes with EMVs.



expected value of making the \$5 billion counteroffer, and finally we will compare this expected value with the sure \$2 billion that Texaco is offering now.

The expected value of the court decision is the weighted average of the possible outcomes:

$$\begin{aligned} \text{EMV}(\text{Court Decision}) &= [P(\text{Award} = 10.3) \times 10.3] + [P(\text{Award} = 5) \times 5] \\ &\quad + [P(\text{Award} = 0) \times 0] = [0.2 \times 10.3] \\ &\quad + [0.5 \times 5] + [0.3 \times 0] = 4.56 \end{aligned}$$

We replace both uncertainty nodes representing the court decision with this expected value, as in Figure 4.5. Now, comparing the two alternatives of accepting and refusing Texaco's \$3 billion counteroffer, it is obvious that the expected value of \$4.56 billion is greater than the certain value of \$3 billion, and hence the slash through the "Accept \$3 Billion" branch.

To continue folding back the decision tree, we replace the decision node with the preferred alternative. The decision tree as it stands after this replacement is shown in Figure 4.6. The third step is to calculate the expected value of the alternative "Counteroffer \$5 Billion." This expected value is

$$\begin{aligned} \text{EMV}(\text{Counteroffer } \$5 \text{ Billion}) &= [P(\text{Texaco Accepts}) \times 5] \\ &\quad + [P(\text{Texaco Refuses}) \times 4.56] \\ &\quad + [P(\text{Texaco Counteroffers}) \times 4.56] \\ &= [0.17 \times 5] + [0.50 \times 4.56] + [0.33 \times 4.56] \\ &= 4.63 \end{aligned}$$

Replacing the chance node with its expected value results in the decision tree shown in Figure 4.7. Comparing the values of the two branches, it is clear that the expected value of \$4.63 billion is preferred to the \$2 billion offer from Texaco. According to

Figure 4.5
Hugh Liedtke's decision tree after calculating expected value of court decision.

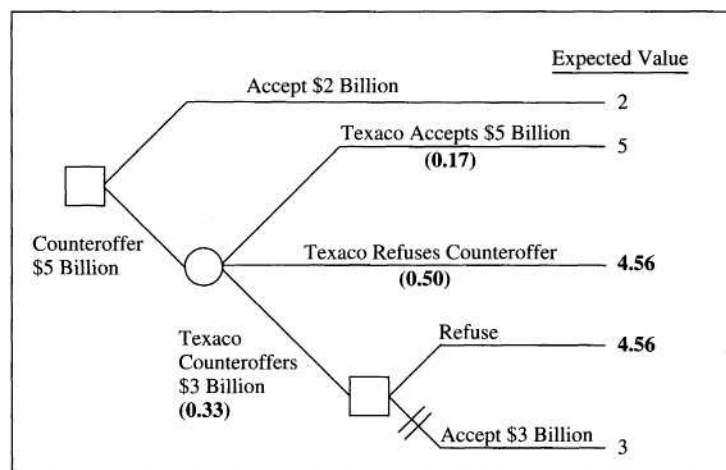


Figure 4.6
 Hugh Liedtke's decision tree after decision node replaced with expected value of preferred alternative.





Figure 4.7
 Hugh Liedtke's decision tree after original tree completely folded back.

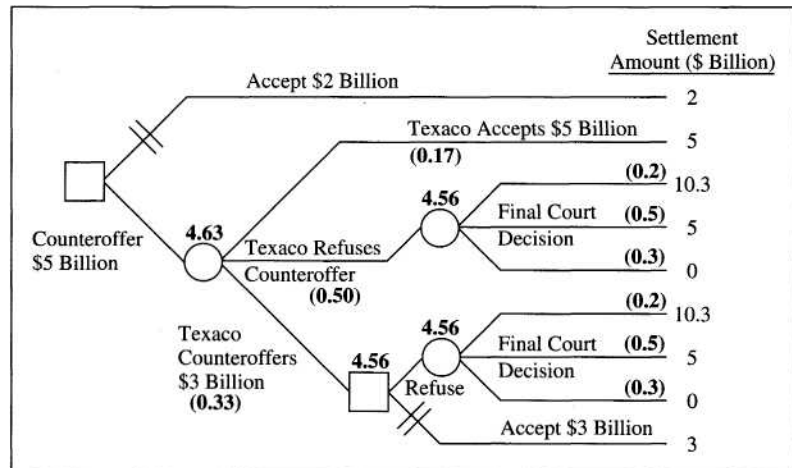


this solution, which implies that decisions should be made by comparing expected values, Liedtke should turn down Texaco's offer but counteroffer a settlement of \$5 billion. If Texaco turns down the \$5 billion and makes another counteroffer of \$3 billion, Liedtke should refuse the \$3 billion and take his chances in court.

We went through this decision in gory detail so that you could see clearly the steps involved. In fact, in solving a decision tree, we usually do not redraw the tree at each step, but simply indicate on the original tree what the expected values are at each of the chance nodes and which alternative is preferred at each decision node. The solved decision tree for Liedtke would look like the tree shown in Figure 4.8, which shows all of the details of the solution. Expected values for the chance nodes are placed above the nodes. The 4.56 above the decision node indicates that if Liedtke gets to this decision point, he should refuse Texaco's offer and take his chances in court for an expected value of \$4.56 billion. The decision tree also shows that his best current choice is to make the \$5 billion counteroffer with an expected payoff of \$4.63 billion.

The decision tree shows clearly what Liedtke should do if Texaco counteroffers \$3 billion: He should refuse. This is the idea of a contingent strategy. If a particular course of events occurs (Texaco's counteroffer), then there is a specific course of action to take (refuse the counteroffer). Moreover, in deciding whether to accept Texaco's current \$2 billion offer, Liedtke must know what he will do in the event that Texaco returns with a counteroffer of \$3 billion. This is why the decision tree is solved backward. In order to make a good decision at the current time, we have to know what the appropriate contingent strategies are in the future.

Figure 4.8
Hugh Liedtke's solved
decision tree.



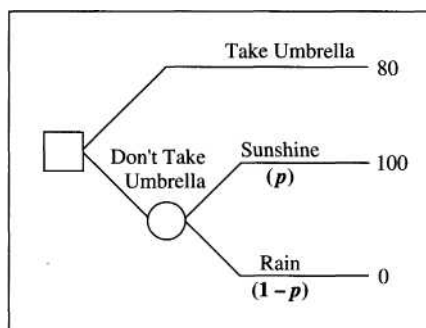
Solving Influence Diagrams: Overview

Solving decision trees is straightforward, and EMVs for small trees can be calculated by hand relatively easily. The procedure for solving an influence diagram, though, is somewhat more complicated. Fortunately, computer programs are available to do the calculations. In this short section we give an overview of the issues involved in solving an influence diagram. For interested readers, the following optional section goes through a complete solution of the influence diagram of the Texaco-Pennzoil decision.

While influence diagrams appear on the surface to be rather simple, much of the complexity is hidden. Our first step is to take a close look at how an influence diagram translates information into an internal representation. An influence diagram "thinks" about a decision in terms of a symmetric expansion of the decision tree from one node to the next.

For example, suppose we have the basic decision tree shown in Figure 4.9, which represents the "umbrella problem" (see Exercise 3.9). The issue is whether or

Figure 4.9
Umbrella problem.



not to take your umbrella. If you do not take the umbrella, and it rains, your good clothes (and probably your day) are ruined, and the consequence is zero (units of satisfaction). However, if you do not take the umbrella and the sun shines, this is the best of all possible consequences with a value of 100. If you decide to take your umbrella, your clothes will not get spoiled. However, it is a bit of a nuisance to carry the umbrella around all day. Your consequence is 80, between the other two values.

If we were to represent this problem with an influence diagram, it would look like the diagram in Figure 4.10. Note that it does not matter whether the sun shines or not if you take the umbrella. If we were to reconstruct exactly how the influence diagram "thinks" about the umbrella problem in terms of a decision tree, the representation would be that shown in Figure 4.11. Note that the uncertainty node on the "Take Umbrella" branch is an unnecessary node. The payoff is the same regardless of the weather. In a decision-tree model, we can take advantage of this fact by not even drawing the unnecessary node. Influence diagrams, however, use the symmetric decision tree, even though this may require unnecessary nodes (and hence unnecessary calculations).

With an understanding of the influence diagram's internal representation, we can talk about how to solve an influence diagram. The procedure essentially solves the symmetric decision tree, although the terminology is somewhat different. Nodes are *reduced*; reduction amounts to calculating expected values for chance nodes and choosing the largest expected value at decision nodes, just as we did with the decision tree. Moreover, also parallel with the decision-tree procedure, as nodes are reduced, they are removed from the diagram. Thus, solving the influence diagram in Figure 4.10 would require first reducing the "Weather" node (calculating the expected values) and then reducing the "Take Umbrella?" node by choosing the largest expected value.

Figure 4.10
Influence diagram of
the umbrella problem.

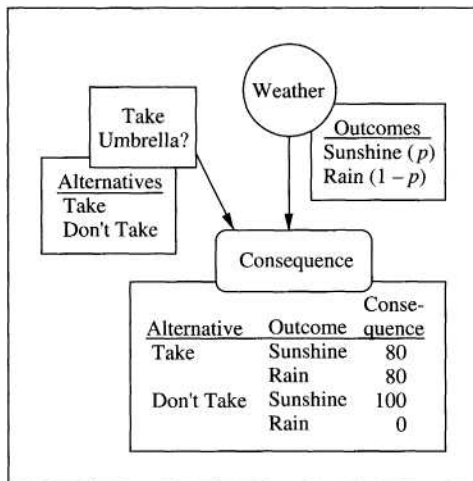



Figure 4.11
How the influence diagram "thinks" about the umbrella problem.



Solving Influence Diagrams: The Details (Optional)

Consider the Texaco-Pennzoil case in influence-diagram form, as shown in Figure 4.12. This diagram shows the tables of alternatives, outcomes (with probabilities), and consequences that are contained in the nodes. The consequence table in this case is too complicated to put into Figure 4.12. We will work with it later in great detail, but if you want to see it now, it is displayed in Table 4.1.

Figure 4.12 needs explanation. The initial decision is whether to accept Texaco's offer of \$2 billion. Within this decision node a table shows that the available alternatives are to accept the offer or make a counteroffer. Likewise, under the "Pennzoil Reaction" node is a table that lists "Accept 3" and "Refuse" as alternatives. The chance node "Texaco Reaction" contains a table showing the probabilities of Texaco accepting a counteroffer of \$5 billion, making an offer of \$3 billion, or refusing to

Figure 4.12
Influence diagram for Liedtke's decision.

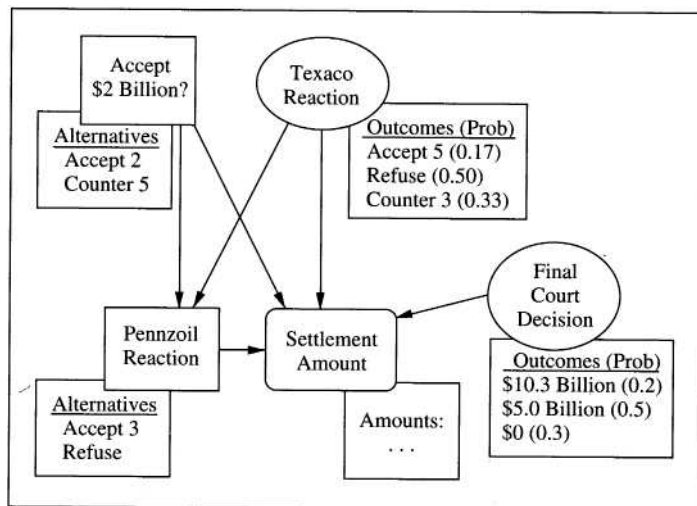


Table 4.1
Consequence table for the influence diagram of Liedtke's decision.

Accept \$2 Billion?	Texaco Reaction (\$ Billion)	Pennzoil Reaction (\$ Billion)	Final Court Decision (\$ Billion)	Settlement Amount (\$ Billion)
Accept 2	Accept 5	Accept 3	10.3	2.0
			5	2.0
			0	2.0
		Refuse	10.3	2.0
			5	2.0
			0	2.0
	Offer 3	Accept 3	10.3	2.0
			5	2.0
			0	2.0
		Refuse	10.3	2.0
			5	2.0
			0	2.0
Offer 5	Accept 5	Accept 3	10.3	5.0
			5	5.0
			0	5.0
		Refuse	10.3	5.0
			5	5.0
			0	5.0
	Offer 3	Accept 3	10.3	3.0
			5	3.0
			0	3.0
		Refuse	10.3	10.3
			5	5.0
			0	0.0
Refuse	Accept 3	10.3	10.3	
		5	5.0	
		0	0.0	
	Refuse	10.3	10.3	
		5	5.0	
		0	0.0	

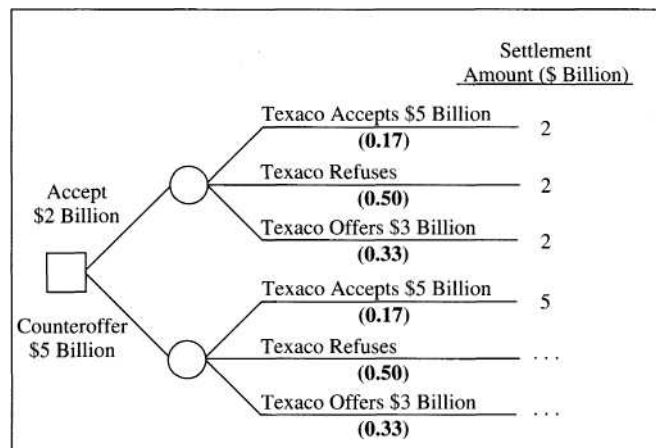
negotiate. Finally, the "Final Court Decision" node has a table with its outcomes and associated probabilities.

The thoughtful reader should have an immediate reaction to this. After all, whether Texaco reacts depends on whether Liedtke makes his \$5 billion counteroffer in the first place! Shouldn't there be an arrow from the decision node "Accept \$2 Billion" to the "Texaco Reaction" node? The answer is yes, there could be such an arrow, but it is unnecessary and would only complicate matters. The reason is that, as with the umbrella example above, the influence diagram "thinks" in terms of a symmetric expansion of the decision tree. Figure 4.13 shows a portion of the tree that deals with Liedtke's initial decision and Texaco's reaction. An arrow in Figure 4.12 from "Accept \$2 Billion" to "Texaco Reaction" would indicate that the decision made (accepting or rejecting the \$2 billion) would affect the chances associated with Texaco's reaction to a counteroffer. But the uncertainty about Texaco's response to a \$5 billion counteroffer does not depend on whether Liedtke accepts the \$2 billion. Essentially, the influence diagram is equivalent to a decision tree that is symmetric.

For similar reasons, there are no arrows between "Final Court Decision" and the other three nodes. If some combination of decisions comes to pass so that Pennzoil and Texaco agree to a settlement, it does not matter what the court decision would be. The influence diagram implicitly includes the "Final Court Decision" node with the agreed-upon settlement regardless of the "phantom" court outcome.

How is all of this finally resolved in the influence-diagram representation? Everything is handled in the consequence node. This node contains a table that gives Liedtke's settlement for every possible combination of decisions and outcomes. That table (Table 4.1) shows that the settlement is \$2 billion if Liedtke accepts the current offer, regardless of the other outcomes. It also shows that if Liedtke counteroffers \$5 billion and Texaco accepts, then the settlement is \$5 billion regardless of the court decision or Pennzoil's reaction (neither of which have any impact if Texaco accepts the \$5 billion). The table also shows the details of the

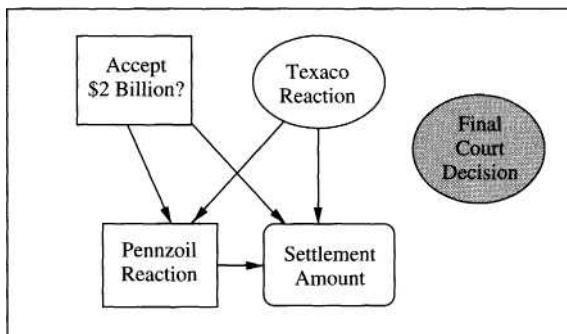
Figure 4.13
How the influence diagram "thinks" about the Texaco-Pennzoil case.



court outcomes if either Texaco refuses to negotiate after Liedtke's counteroffer or if Liedtke refuses a Texaco counteroffer. And so on. The table shows exactly what the payoff is to Pennzoil under all possible combinations. The column headings in Table 4.1 represent nodes that are predecessors of the value node. In this case, both decision nodes and both chance nodes are included because all are predecessors of the value node. We can now discuss how the algorithm for solving an influence diagram proceeds. Take the Texaco-Pennzoil diagram as drawn in Figure 4.12. As mentioned above, our strategy will be to *reduce nodes* one at a time. The order of reduction is reminiscent of our solution in the case of the decision tree. The first node reduced is "Final Court Decision," resulting in the diagram in Figure 4.14. In this first step, expected values are calculated using the "Final Court Decision" probabilities, which yields Table 4.2. All combinations of decisions and possible outcomes of Texaco's reaction are shown. For example, if Liedtke counteroffers \$5 billion and Texaco refuses to negotiate, the expected value of \$4.56 billion is listed regardless of the decision in the "Pennzoil Reaction" node (because that decision is meaningless if Texaco initially refuses to negotiate). If Liedtke accepts the \$2 billion offer, the expected value is listed as \$2 billion, regardless of other outcomes. (Of course, there is nothing uncertain about this outcome; the value that we know will happen is the expected value.) If Liedtke offers 5, Texaco offers 3, and finally Liedtke refuses to continue negotiating, then the expected value is given as 4.56. And so on.

The next step is to reduce the "Pennzoil Reaction" node. The resulting influence diagram is shown in Figure 4.15. Now the table in the consequence node (Table 4.3) reflects the decision that Liedtke should choose the alternative with the highest expected value (refuse to negotiate) if Texaco makes the counteroffer of \$3 billion. Thus, the table now says that, if Liedtke offers \$5 billion and Texaco either refuses to negotiate or counters with \$3 billion, the expected value is \$4.56 billion. If Texaco accepts the \$5 billion counteroffer, the expected value is \$5 billion, and if Liedtke accepts the current offer, the expected value is \$2 billion. (Again, there is nothing uncertain about these values; the expected value in these cases is just the value that we know will occur.)

Figure 4.14
First step in solving
the influence diagram.



The third step is to reduce the "Texaco Reaction" node, as shown in Figure 4.16. As with the first step, this involves taking the table of consequences (now expected values) within the "Settlement Amount" node and calculating expected values again. The resulting table has only two entries (Table 4.4). The expected value of Liedtke accepting \$2 billion is just \$2 billion, and the expected value of countering with \$5 billion is \$4.63 billion.

The fourth and final step is simply to figure out which decision is optimal in the "Accept \$2 Billion?" node and to record the result. This final step is shown in Figure 4.17. The table associated with the decision node indicates that Liedtke's optimal choice is to counteroffer \$5 billion. The payoff table now contains only one value, \$4.63 billion, the expected value of the optimal decision.

Reviewing the procedure, you should be able to see that it followed basically the same steps that we followed in folding back the decision tree.

Table 4.2
Table for Liedtke's decision after reducing the "Final Court Decision" node.

Accept \$2 Billion?	Texaco Reaction	Pennzoil Reaction	Expected Value (\$ Billion)
Accept 2	Accept 5	Accept 3	2
		Refuse	2
	Offer 3	Accept 3	2
		Refuse	2
		Accept 3	2
		Refuse	2
Offer 5	Accept 5	Accept 3	
		Refuse	5
	Offer 3	Accept 3	3
		Refuse	4.56
		Accept 3	4.56
		Refuse	4.56

Figure 4.15
Second step in solving the influence diagram.

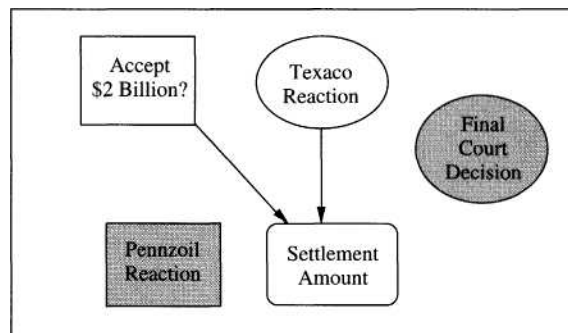


Table 4.3
Table for Liedtke's decision after reducing "Final Court Decision" and "Pennzoil Reaction" nodes.

Accept \$2 Billion?	Texaco Reaction	Expected Value (\$ Billion)
Accept 2	Accept 5	2.00
	Offer 3	2.00
	Refuse	2.00
Offer 5	Accept 5	5.00
	Offer 3	4.56
	Refuse	4.56

Table 4.4
Table for Liedtke's decision after reducing "Final Court Decision," "Pennzoil Reaction," and "Texaco Reaction" nodes.

Accept \$2 Billion?	Expected Value (\$ Billion)
Accept 2	2.00
Offer 5	4.63

Figure 4.16
Third step in solving the influence diagram.

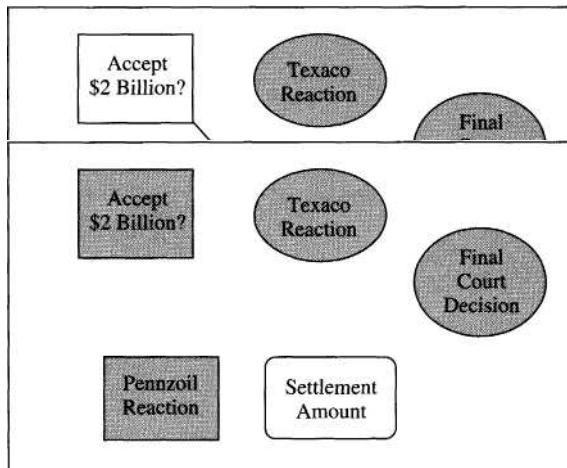
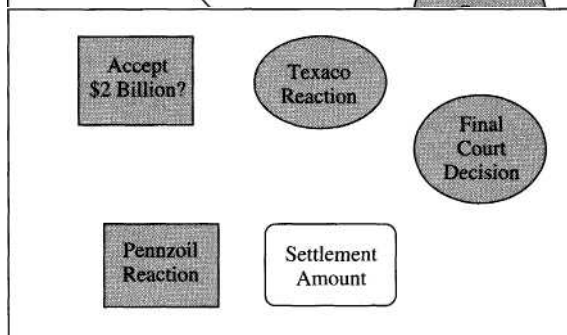


Figure 4.17
Final step in solving the influence diagram.



Solving Influence Diagrams: An Algorithm (Optional)

The example above should provide some insight into how influence diagrams are solved. Fortunately, you will not typically have to solve influence diagrams by hand; computer programs are available to accomplish this. It is worthwhile, however, to spend a few moments describing the procedure that is used to solve influence diagrams. A set procedure for solving a problem is called an *algorithm*. You have already learned the algorithm for solving a decision tree (the folding-back procedure). Now let us look at an algorithm for solving influence diagrams.

1 First, we simply clean up the influence diagram to make sure it is ready for solution. Check to make sure the influence diagram has only one consequence node (or a series of consequence nodes that feed into one "super" consequence node) and that there are no cycles. If your diagram does not pass this test, you must fix it before it can be solved. In addition, if any nodes other than the consequence node have arrows into them but not out of them, they can be eliminated. Such nodes are called *barren nodes* and have no effect on the decision that would be made. Replace any intermediate-calculation nodes with chance nodes. (This includes any consequence nodes that feed into a "super" consequence node representing a higher-level objective in the objectives hierarchy.) For each possible combination of the predecessor node outcomes, such a node has only one outcome that happens with probability 1.

2 Next, look for any chance nodes that (a) directly precede the consequence node and (b) do *not* directly precede any other node. Any such chance node found should be reduced by calculating expected values. The consequence node then inherits the predecessors of the reduced nodes. (That is, any arrows that went into the node you just reduced should be redrawn to go into the consequence node.)

This step is just like calculating expected values for chance nodes at the far right-hand side of a decision tree. You can see how this step was implemented in the Texaco-Pennzoil example. In the original diagram, Figure 4.12, the "Final Court Decision" node is the only chance node that directly precedes the consequence node *and* does not precede any decision node. Thus it is reduced by the expected-value procedure, resulting in Table 4.2. The consequence node does not inherit any new direct predecessors as a result of this step because "Final Court Decision" has no direct predecessors.

3 Next, look for a decision node that (a) directly precedes the consequence node and (b) has as predecessors all of the other direct predecessors of the consequence node. *If you do not find any such decision node, go directly to Step 5.* If you find such a decision node, you can reduce it by choosing the optimum value. When decision nodes are reduced, the consequence node does not inherit any new predecessors. This step may create some barren nodes, which can be eliminated from the diagram.

This step is like folding a decision tree back through a decision node at the far right-hand side of the tree. In the Texaco-Pennzoil problem, this step was

implemented when we reduced "Pennzoil Reaction." In Figure 4.14, this node satisfies the criteria for reduction because it directly precedes the consequence node, and the other nodes that directly precede the consequence node also precede "Pennzoil Reaction." In reducing this node, we choose the option for "Pennzoil Reaction" that gives the highest expected value, and as a result we obtain Table 4.3. No barren nodes are created in this step.

- 4 Return to Step 2 and continue until the influence diagram is completely solved (all nodes reduced). This is just like working through a decision tree until all of the nodes have been processed from right to left.
- 5 You arrived at this step after reducing all possible chance nodes (if any) and then not finding any decision nodes to reduce. How could this happen? Consider the influence diagram of the hurricane problem in Figure 3.12. None of the chance nodes satisfy the criteria for reduction, and the decision node also cannot be reduced. In this case, one of the arrows between chance nodes must be reversed. This is a procedure that requires probability manipulations through the use of Bayes' theorem (Chapter 7). We will not go into the details of the calculations here because most of the simple influence diagrams that you might be tempted to solve by hand will not require arrow reversals.

Finding an arrow to reverse is a delicate process. First, find the correct chance node. The criteria are that (a) it directly precedes the consequence node and (b) it does not directly precede any decision node. Call the selected node A. Now look at the arrows out of node A. Find an arrow from A to chance node B (call it $A \rightarrow B$) such that there is no other way to get from A to B by following arrows. The arrow $A \rightarrow B$ can be reversed using Bayes' theorem. Afterward, both nodes inherit each other's direct predecessors *and* keep their own direct predecessors.

After reversing an arrow, return to Step 2 and continue until the influence diagram is solved. (More arrows may need to be reversed before a node can be reduced, but that only means that you may come back to Step 5 one or more times in succession.)

This description of the influence-diagram solution algorithm is based on the complete (and highly technical) description given in Shachter (1986). The intent is not to present a "cookbook" for solving an influence diagram because, as indicated, virtually all but the simplest influence diagrams will be solved by computer. The description of the algorithm, however, is meant to show the parallels between the influence-diagram and decision-tree solution procedures.

Risk Profiles

The idea of expected value is appealing, and comparing two alternatives on the basis of their EMVs is straightforward. For example, Liedtke's expected values are \$2 billion and \$4.63 billion for his two immediate alternatives. But you might have noticed that these two numbers are not exactly perfect indicators of what might happen. In

particular, suppose that Liedtke decides to counteroffer \$5 billion: He might end up with \$10.3 billion, \$5 billion, or nothing, given our simplification of the situation. Moreover, the interpretation of EMV as the average amount that would be obtained by "playing the game" a large number of times is not appropriate here. The "game" in this case amounts to suing Texaco — not a game that Pennzoil will play many times!

That Liedtke could come away from his dealings with Texaco with nothing indicates that choosing to counteroffer is a somewhat risky alternative. In later chapters we will look at the idea of risk in more detail. For now, however, we can intuitively grasp the relative riskiness of alternatives by studying their *risk profiles*.

A risk profile is a graph that shows the chances associated with possible consequences. Each risk profile is associated with a *strategy*, a particular immediate alternative, as well as specific alternatives in future decisions. For example, the risk profile for the "Accept \$2 Billion" alternative is shown in Figure 4.18. There is a 100% chance that Liedtke will end up with \$2 billion. The risk profile for the strategy "Counteroffer \$5 Billion; Refuse Texaco Counteroffer" is somewhat more complicated and is shown in Figure 4.19. There is a 58.5% chance that the eventual settlement is \$5 billion, a 16.6% chance of \$10.3 billion, and a 24.9% chance of nothing. These numbers are easily calculated. For example, take the \$5 billion amount. This can happen in three different ways. There is a 17% chance that it happens because Texaco accepts. There is a 25% chance that it happens because Texaco refuses and the judge awards \$5 billion. (That is, there is a 50% chance that Texaco refuses times a 50% chance that the award is \$5 billion.) Finally, the chances are 16.5% that the settlement is \$5 billion because Texaco counteroffers \$3 billion, Liedtke refuses and goes to court, and the judge awards \$5 billion. That is, 16.5% equals 33% times 50%. Adding up, we get the chance of \$5 billion = 17% + 25% + 16.5% = 58.5%.

In constructing a risk profile, we collapse a decision tree by multiplying out the probabilities on sequential chance branches. At a decision node, only one branch is taken; in the case of "Counteroffer \$5 Billion; Refuse Texaco Counteroffer," we use only the indicated alternative for the second decision, and so this decision node need not be included in the collapsing process. You can think about the process as one in which nodes are gradually removed from the tree in much the same sense as we did

Figure 4.18
Risk profile for the
"Accept \$2 Billion"
alternative.





Figure 4.19

Risk profile for the "Counteroffer \$5 Billion; Refuse Texaco Counteroffer" strategy.



with the folding-back procedure, except that in this case we keep track of the possible outcomes and their probabilities. Figures 4.20, 4.21, and 4.22 show the progression of collapsing the decision tree in order to construct the risk profile for the "Counteroffer \$5 Billion; Refuse Texaco Counteroffer" strategy.

By looking at the risk profiles, the decision maker can tell a lot about the riskiness of the alternatives. In some cases a decision maker can choose among alternatives on the basis of their risk profiles. Comparing Figures 4.18 and 4.19, it is clear that the worst possible consequence for "Counteroffer \$5 Billion; Refuse Texaco Counteroffer" is less than the value for "Accept \$2 billion." On the other hand, the largest amount (\$10.3 billion) is much better than \$2 billion. Hugh Liedtke has to decide whether the risk of perhaps coming away empty-handed is worth the possibility of getting more than \$2 billion. This is clearly a case of a basic risky decision, as we can see from the collapsed decision tree in Figure 4.22.

Risk profiles can be calculated for strategies that might not have appeared as opti-

Figure 4.20

First step in collapsing the decision tree to make risk profile for "Counteroffer \$5 Billion; Refuse Texaco Counteroffer" strategy. The decision node has been removed to leave only the outcomes associated with the "Refuse" branch.

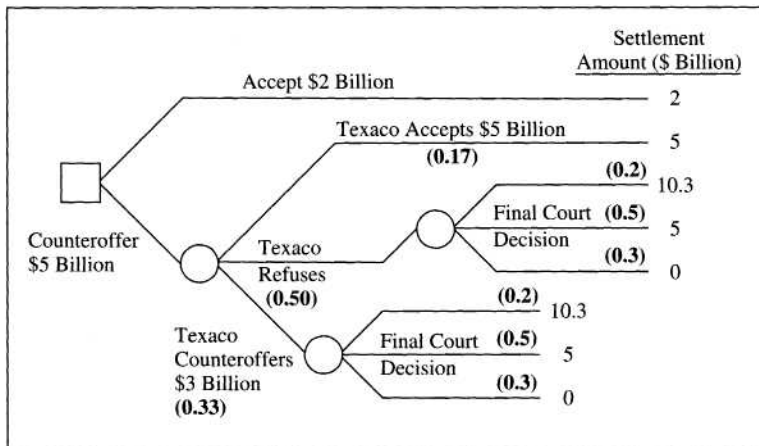


Figure 4.21

Second step in collapsing the decision tree to make a risk profile.

The three chance nodes have been collapsed into one chance node. The probabilities on the branches are the product of the probabilities from sequential branches in Figure 4.20.

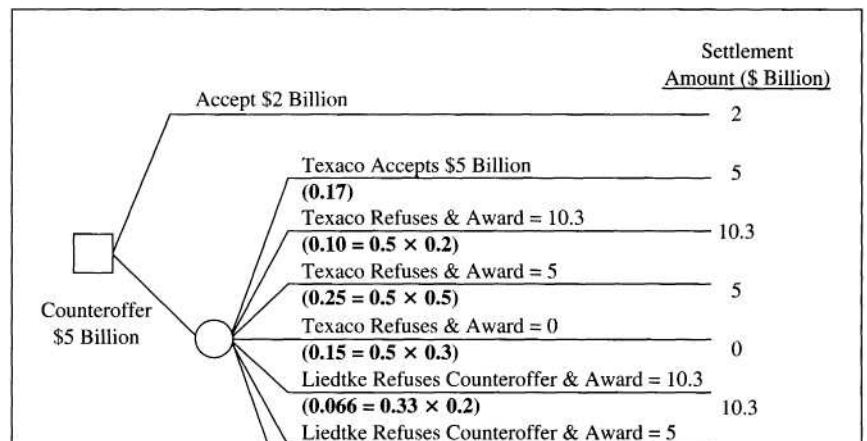
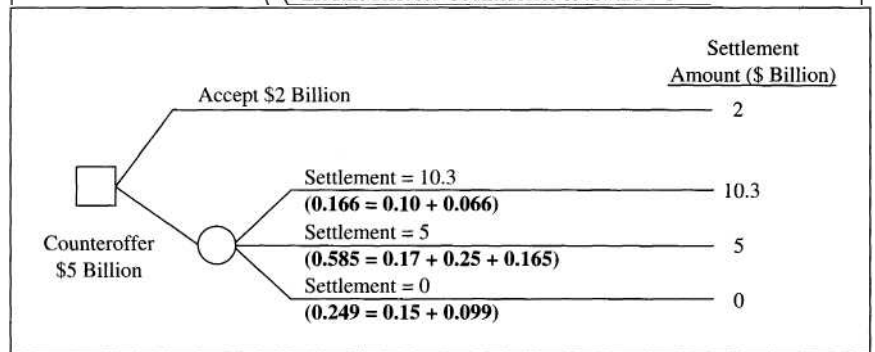


Figure 4.22

Third step in collapsing the decision tree to make a risk profile.


The seven branches from the chance node in Figure 4.21 have been combined into three branches.



mal in an expected-value analysis. For example, Figure 4.23 shows the risk profile for "Counteroffer \$5 Billion; Accept \$3 Billion," which we ruled out on the basis of EMV. Comparing Figures 4.23 and 4.19 indicates that this strategy yields a smaller chance of getting nothing, but also less chance of a \$10.3 billion judgment. Compensating for this is the greater chance of getting something in the middle: \$3 or \$5 billion.

Although risk profiles can in principle be used as an alternative to EMV to check every possible strategy, for complex decisions it can be tedious to study many risk profiles. Thus, a compromise is to look at strategies only for the first one or two decisions, on the assumption that future decisions would be made using a decision rule such as maximizing expected value, which is itself a kind of strategy. (This is the approach used by many decision-analysis computer programs, DPL included.) Thus, in

Figure 4.23
 Risk profile for the
 "Counteroffer \$5
 Billion; Accept
 \$3 Billion "
 strategy.




the Texaco-Pennzoil example, one might compare only the "Accept \$2 Billion" and "Counteroffer \$5 Billion; Refuse Texaco Counteroffer" strategies.

Cumulative Risk Profiles

We also can present the risk profile in cumulative form. Figure 4.24 shows the *cumulative risk profile* for "Counteroffer 5 Billion; Refuse Texaco Counteroffer." In this format, the vertical axis is the chance that the payoff is less than or equal to the corresponding value on the horizontal axis. This is only a matter of translating the information contained in the risk profile in Figure 4.19. There is no chance that the settlement will be less than zero. At zero, the chance jumps up to 24.9%, because there is a substantial chance that the court award will be zero. The graph continues at 24.9% across to \$5 billion. (For example, there is a 24.9% chance that the settlement is less than or equal to \$3.5 billion; that is, there is the 24.9% chance that the settlement is zero, and that is less than \$3.5 billion.) At \$5 billion, the line jumps up to 83.4% (which is 24.9% + 58.5%), because there is an 83.4% chance that the settlement is less than or equal to \$5 billion. Finally, at \$10.3 billion, the cumulative graph jumps up to 100%: The chance is 100% that the settlement is less than or equal to \$10.3 billion.

Figure 4.24
 Cumulative risk profile
 for "Counteroffer \$5
 Billion; Refuse Texaco
 Counteroffer."



Thus, you can see that creating a cumulative risk profile is just a matter of adding up, or accumulating, the chances of the individual payoffs. For any specific value along the horizontal axis, we can read off the chance that the payoff will be less than or equal to that specific value. Cumulative risk profiles will be very helpful in the next section in our discussion of dominance.

Dominance: An Alternative to EMV

Comparing expected values of different risky prospects is useful, but in many cases EMV inadequately captures the nature of the risks that must be compared. With risk profiles, however, we can make a more comprehensive comparison of the risks. But how can we choose one risk profile over another? Unfortunately, there is no clear answer that can be used in all situations. By using the idea of *dominance*, though, we can identify those profiles (and their associated strategies) that can be ignored. Such strategies are said to be *dominated*, because we can show logically, according to some rules relating to cumulative risk profiles, that there are better risks (strategies) available.

Suppose we modify Liedtke's decision as shown in Figure 4.2 so that \$2.5 billion is the minimum amount that he believes he could get in a court award. This decision is diagrammed in Figure 4.25. Now what should he do? It is rather obvious. Because he believes that he could do no worse than \$2.5 billion if he makes a counteroffer, he should clearly shun Texaco's offer of 2 billion. This kind of dominance is called *deterministic dominance*, signifying that the dominating alternative pays off at least as much as the one that is dominated.

We can show deterministic dominance in terms of the cumulative risk profiles displayed in Figure 4.26. The cumulative risk profile for "Accept \$2 Billion" goes from zero to 100% at \$2 billion, because the settlement for this alternative is bound to be \$2 billion. The risk profile for "Counteroffer \$5 Billion; Refuse Texaco

Figure 4.25
Hugh Liedtke's decision tree, assuming \$2.5 billion is minimum court award.

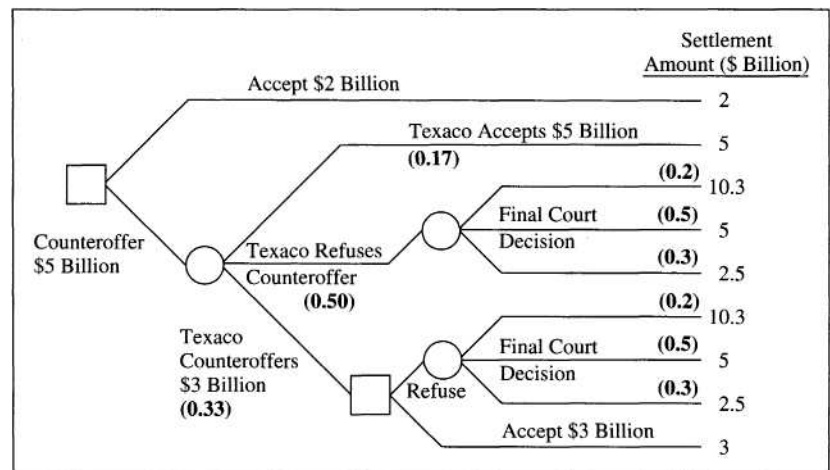



Figure 4.26
 Cumulative risk profiles for alternatives in Figure 4.25.



Counteroffer" starts at \$2.5 billion but does not reach 100% until \$10.3 billion. Deterministic dominance can be detected in the risk profiles by comparing the value where one cumulative risk profile reaches 100% with the value where another risk profile begins. If there is a value x such that the chance of the payoff being less than or equal to x is 100% in alternative B, and the chance of the payoff being *less than* x is 0% in Alternative A, then A deterministically dominates B. Graphically, continue the vertical line where alternative A first leaves 0% (the vertical line at \$2.5 billion for "Counteroffer 5 Billion"). If that vertical line corresponds to 100% for the other cumulative risk profile, then A dominates B. Thus, even if the minimum court award had been \$2 billion instead of \$2.5 billion, "Counteroffer \$5 Billion" still would have dominated "Accept \$2 Billion."

The following example shows a similar kind of dominance. Suppose that Liedtke is choosing between two different law firms to represent Pennzoil. He considers both law firms to be about the same in terms of their abilities to deal with the case, but one charges less in the event that the case goes to court. The full decision tree for this problem appears in Figure 4.27. Which choice is preferred? Again, it's rather obvious; the settlement amounts for choosing Firm A are the same as the corresponding amounts for choosing Firm B, except that Pennzoil gets more with Firm A if the case results in a damage award in the final court decision. Choosing Firm A is like choosing Firm B and possibly getting a bonus as well. Firm A is said to display *stochastic dominance* over Firm B. Many texts also use the term *probabilistic dominance* to indicate the same thing. (Strictly speaking, this is first-order stochastic dominance. Higher-order stochastic dominance comes into play when we consider preferences regarding risk.)

The cumulative risk profiles corresponding to Firms A and B (and assuming that Liedtke refuses a Texaco counteroffer) are displayed in Figure 4.28. The two cumulative risk profiles almost coincide; the only difference is that Firm A's profile is slightly to the right of Firm B's at \$5 and \$10 billion, which represents the possibility of Pennzoil having to pay less in fees. Stochastic dominance is represented in the cumulative risk profiles by the fact that the two profiles do not cross and that there is some space between them. That is, if two cumulative risk profiles are such that no part

Figure 4.27 A decision tree comparing two law firms. Firm A charges less than Firm B if Pennzoil is awarded damages in court.

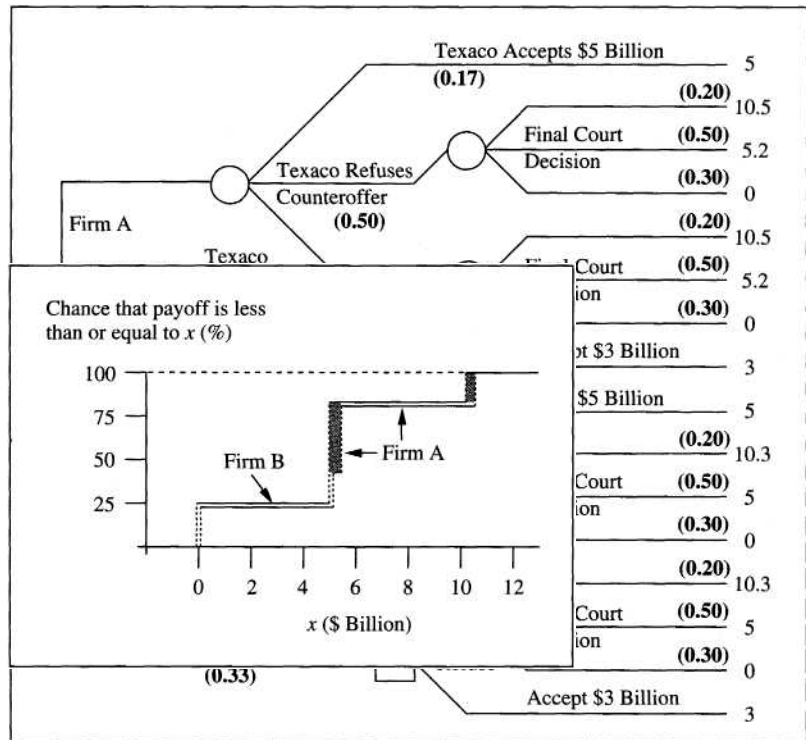


Figure 4.28 Cumulative risk profiles for two law firms in Figure 4.27. Firm A stochastically dominates Firm B.

of Profile A lies to the left of B, and at least some part of it lies to the right of B, then the strategy corresponding to Profile A stochastically dominates the strategy for Profile B.

The next example demonstrates stochastic dominance in a slightly different form. Instead of the consequences, the pattern of the probability numbers makes the preferred alternative apparent. Suppose Liedtke's choice is between two law firms

that he considers to be of different abilities. The decision tree is shown in Figure 4.29. Carefully examine the probabilities in the branches for the final court decision. Which law firm is preferred? This is a somewhat more subtle situation than the preceding. The essence of the problem is that for Firm C, the larger outcome values have higher probabilities. The settlement with Firm C is not bound to be at least as great or greater than that with Firm D, but with Firm C the settlement is more likely to be greater. Think of Firm C as being a better gamble if the situation comes down to a court decision. Situations like this are characterized by two alternatives that offer the same possible consequences, but the dominating alternative is more likely to bring a better consequence.

Figure 4.30 shows the cumulative risk profiles for the two law firms in this example. As in the last example, the two profiles nearly coincide, although space is found between the two profiles because of the different probabilities associated with the court award. Because Firm C either coincides with or lies to the right of Firm D, we can conclude that Firm C stochastically dominates Firm D.

Stochastic dominance can show up in a decision problem in several ways. One way is in terms of the consequences (as in Figure 4.27), and another is in terms of the probabilities (as in Figure 4.29). Sometimes stochastic dominance may emerge as a mixture of the two; both slightly better payoffs and slightly better probabilities may lead to one alternative dominating another.

Figure 4.29
 Decision tree comparing two law firms. Firm C has a better chance of winning a damage award in court than does Firm D.

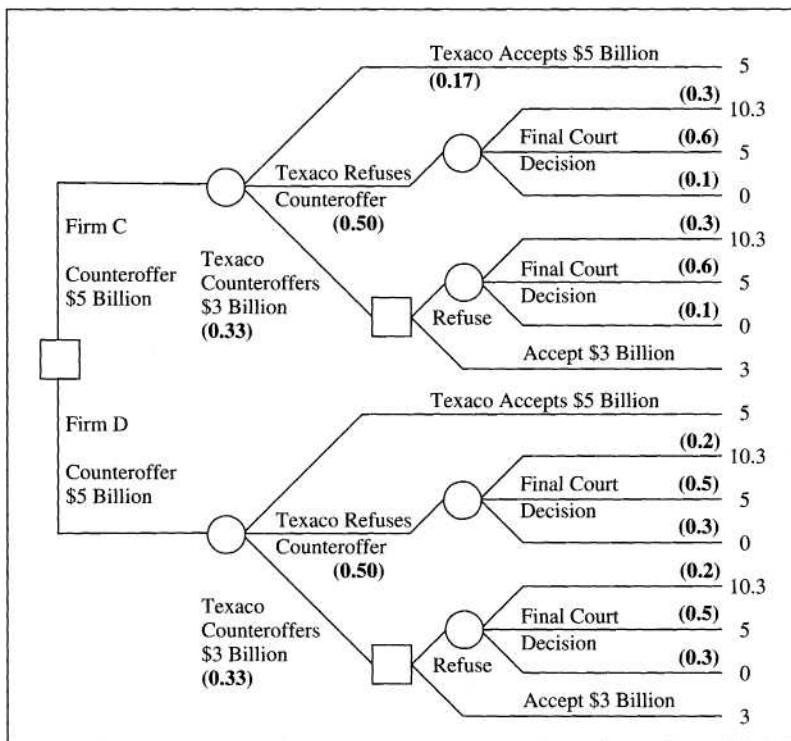



Figure 4.30
 Cumulative risk profiles for law firms in Figure 4.29. Firm C stochastically dominates Firm D.



What is the relationship between stochastic dominance and expected value? It turns out that if one alternative dominates another, then the dominating alternative must have the higher expected value. This is a property of dominant alternatives that can be proven mathematically. To get a feeling for why it is true, think about the cumulative risk profiles, and imagine the EMV for a dominated Alternative B. If Alternative A dominates B, then its cumulative risk profile must lie at least partly to the right of the profile for B. Because of this, the EMV for A must also lie to the right of, and hence be greater than, the EMV for B.

Although this discussion of dominance has been fairly brief, one should not conclude that dominance is not important. Indeed, screening alternatives on the basis of dominance begins implicitly in the structuring phase of decision analysis, and, as alternatives are considered, they usually are at least informally compared to other alternatives. Screening alternatives formally on the basis of dominance is an important decision-analysis tool. If an alternative can be eliminated early in the selection process on that basis, considerable cost can be saved in large-scale problems. For example, suppose that the decision is where to build a new electric power plant. Analysis of proposed alternatives can be exceedingly expensive. If a potential site can be eliminated in an early phase of the analysis on the grounds that another dominates it, then that site need not undergo full analysis.

Making Decisions with Multiple Objectives

So far we have learned how to analyze a single-objective decision; in the Texaco-Pennzoil example, we have focused on Liedtke's objective of maximizing the settlement amount. How would we deal with a decision that involves multiple objectives? In this section, we learn how to extend the concepts of expected value and risk profiles to multiple-objective situations. In contrast to the grandiose Texaco-Pennzoil example, consider the following down-to-earth example of a young person deciding which of two summer jobs to take.

THE SUMMER JOB

Sam Chu was in a quandary. With two job offers in hand, the choice he should make was far from obvious. The first alternative was a job as an assistant at a local small business; the job would pay minimum wage (\$5.25 per hour), require 25 to 35 hours per week, and the hours would be primarily during the week, leaving the weekends free. The job would last for three months, but the exact amount of work, and hence the amount Sam could earn, was uncertain. On the other hand, the free weekends could be spent with friends.

The second alternative was to work as a member of a trail-maintenance crew for a conservation organization. This job would require 10 weeks of hard work, 40 hours per week at \$6.50 per hour, in a national forest in a neighboring state. The job would involve extensive camping and backpacking. Members of the maintenance crew would come from a large geographic area and spend the entire 10 weeks together, including weekends. Although Sam had no doubt about the earnings this job would provide, the real uncertainty was what the staff and other members of the crew would be like. Would new friendships develop? The nature of the crew and the leaders could make for 10 weeks of a wonderful time, 10 weeks of misery, or anything in between.

From the description, it appears that Sam has two objectives in this context: earning money and having fun this summer. Both are reasonable, and the two jobs clearly differ in these two dimensions; they offer different possibilities for the amount of money earned and the quality of summer fun.

The amount of money to be earned has a natural scale (dollars), and like most of us Sam prefers more money to less. The objective of having fun has no natural scale, though. Thus, a first step is to create such a scale. After considering the possibilities, Sam has created the scale in Table 4.5 to represent different levels of summer fun in the context of choosing a summer job. Although living in town and living in a forest camp pose two very different scenarios, the scale has been constructed in such a way that it can be applied to either job (as well as to any other prospect that might arise). The levels are numbered so that the higher numbers are more preferred.

Table 4.5

A constructed scale for summer fun.

- | | |
|---|---|
| 5 | (Best) A large, congenial group. Many new friendships made. Work is enjoyable, and time passes quickly. |
| 4 | A small but congenial group of friends. The work is interesting, and time off work is spent with a few friends in enjoyable pursuits. |
| 3 | No new friends are made. Leisure hours are spent with a few friends doing typical activities. Pay is viewed as fair for the work done. |
| 2 | Work is difficult. Coworkers complain about the low pay and poor conditions. On some weekends it is possible to spend time with a few friends, but other weekends are boring. |
| 1 | (Worst) Work is extremely difficult, and working conditions are poor. Time off work is generally boring because outside activities are limited or no friends are available. |

With the constructed scale for summer fun, we can represent Sam's decision with the influence diagram and decision tree shown in Figures 4.31 and 4.32, respectively. The influence diagram shows the uncertainty about fun and amount of work, and that these have an impact on their corresponding consequences. The tree reflects Sam's belief that summer fun with the in-town job will amount to Level 3 in the constructed scale, but there is considerable uncertainty about how much fun the forest job will be. This uncertainty has been translated into probabilities based on Sam's uncertainty; how to make such judgments is the topic of Chapter 8. Likewise, the decision tree reflects uncertainty about the amount of work available at the in-town job.

Figure 4.31
Influence diagram for summer-job example.


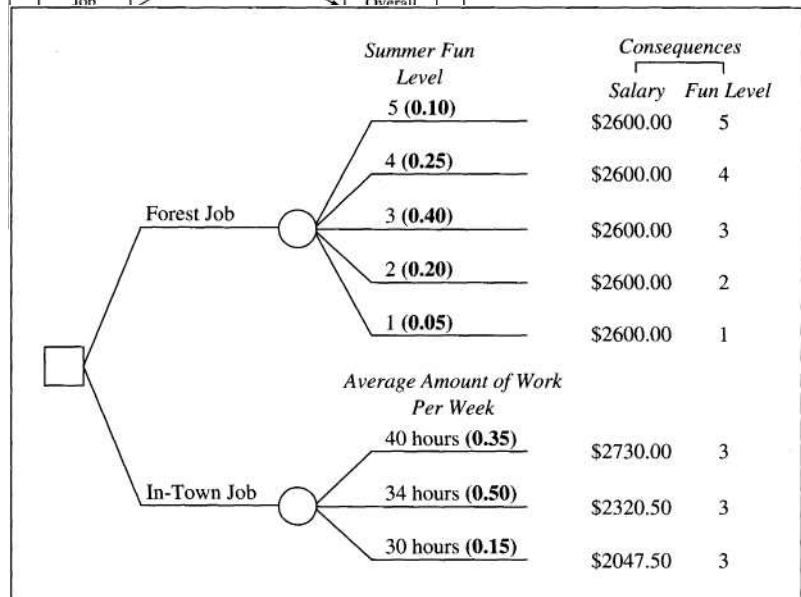


Figure 4.32
Decision tree for summer-job example.




Analysis: One Objective at a Time

One way to approach the analysis of a multiple-objective decision is to calculate the expected value or create the risk profile for each individual objective. In the summer-job example, it is easy enough to do these things for salary. For the forest job, in which there is no uncertainty about salary, the expected value is \$2600, and the risk profile is a single bar at \$2600, as in Figure 4.33. For the in-town job, the expected salary is

$$E(\text{Salary}) = 0.35(\$2730.00) + 0.50(\$2320.50) + 0.15(\$2047.50) = \\ \$2422.88$$

The risk profile for salary at the in-town job is also shown in Figure 4.33.


Figure 4.33
Risk profiles for salary
in the summer-job
example.



Subjective Ratings for Constructed Attribute Scales

For the summer-fun constructed attribute scale, risk profiles can be created and compared (Figure 4.34), but expected-value calculations are not meaningful because no meaningful numerical measurements are attached to the specific levels in the scale. The levels are indeed ordered, but the ordering is limited in what it means. The labels do not mean, for example, that going from Level 2 to Level 3 would give Sam the same increase in satisfaction as going from Level 4 to Level 5. Thus, before we can do any meaningful analysis, Sam must *rate* the different levels in the scale, indicating how much each level is worth (to Sam) relative to the other levels. This is a subjective judgment on Sam's part. Different people with different preferences would be expected to give different ratings for the possible levels of summer fun.

Figure 4.34
*Risk profiles for
 summer fun in the
 summer-job example.*



To make the necessary ratings, we begin by setting the endpoints of the scale. Let the best possible level (Level 5 in the summer-job example) have a value of 100 and the worst possible level (Level 1) a value of 0. Now all Sam must do is indicate how the intermediate levels rate on this scale from 0 to 100 points. For example, Level 4 might be worth 90 points, Level 3, 60 points, and Level 2, 25 points. Sam's assessments indicate that going from Level 3 to Level 4, with an increase of 30 points, is three times as good as going from Level 4 to Level 5 with an increase of only 10 points. Note that there is no inherent reason for the values of the levels to be evenly spaced; in fact, it might be surprising to find perfectly even spacing.

This same procedure can be used to create meaningful measurements for any constructed scale. The best level is assigned 100 points, the worst 0 points, and the decision maker must then assign rating points between 0 and 100 to the intermediate levels. A scale like this assigns more points to the preferred consequences, and the rating points for intermediate levels should reflect the decision maker's relative preferences for those levels.

With Sam's assessments, we can now calculate and compare the expected values for the amount of fun in the two jobs. For the in-town job, this is trivial because there is no uncertainty; the expected value is 60 points. For the forest job, the expected value is

$$\begin{aligned} E(\text{Fun Points}) &= 0.10(100) + 0.25(90) + 0.40(60) + 0.20(25) + 0.05(0) \\ &= 61.5 \end{aligned}$$

With individual expected values and risk profiles, alternatives can be compared. In doing so, we can hope for a clear winner, an alternative that dominates all other alternatives on all attributes. Unfortunately, comparing the forest and in-town jobs does not produce a clear winner. The forest job appears to be better on salary, having no risk and a higher expected value. Considering summer fun, the news is mixed. The in-town job has less risk but a lower expected value. It is obvious that going from one job to the other involves trading risks. Would Sam prefer a slightly higher salary for sure and take a risk on how much fun the summer will be? Or would the in-town job be better, playing it safe with the amount of fun and taking a risk on how much money will be earned?

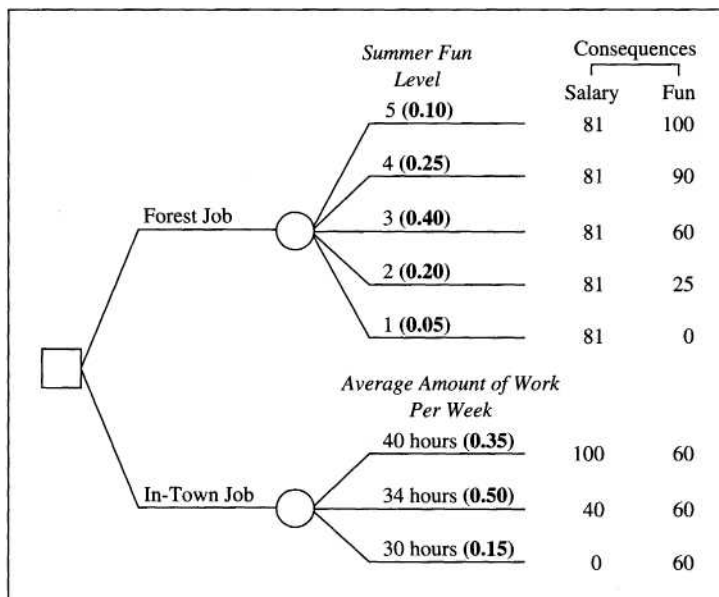
Assessing Trade-Off Weights

The summer-job decision requires Sam to make an explicit trade-off between the objectives of maximizing fun and maximizing salary. How can Sam make this trade-off? Although this seems like a formidable task, a simple thought experiment is possible that will help Sam to understand the relative value of salary and fun.

In order to make the comparison between salary and fun, it is helpful to measure these two on similar scales, and the most convenient arrangement is to put salary on the same 0 to 100 scale that we used for summer fun. As before, the best (\$2730) and worst (2047.50) take values of 100 and 0, respectively. To get the values for the intermediate salaries (\$2320.50 and \$2600), a simple approach is to calculate them proportionately. Thus, we find that \$2320.50 is 40% of the way from \$2047.50 to \$2730, and so it gets a value of 40 on the converted scale. (That is, $[\$2320.50 - \$2047.50] / [\$2730 - \$2047.50] = 0.40$.) Likewise, \$2600 is 81% of the way from \$2047.50 to \$2730, and so it gets a value of 81. (In Chapter 15, we will call this approach *proportional scoring*.) With the ratings for salary and summer fun, we now can create a new consequence matrix, giving the decision tree in Figure 4.35.

Now the trade-off question can be addressed in a straightforward way. The question is how Sam would trade points on the salary scale for points on the fun scale. To do this we introduce the idea of *weights*. What we want to do is assign weights to salary and fun to reflect their relative importance to Sam. Call the weights k_s and k_f , where the subscripts s and f stand for salary and fun, respectively. We will use the weights to calculate a weighted average of the two ratings for any given consequence in order to get an overall score. For example, suppose that $k_s = 0.70$ and

Figure 4.35
Decision tree with ratings for consequences.



$k_f = 0.30$, reflecting a judgment that salary is a little more than twice as important as fun. The overall score (U) for the forest job with fun at Level 3 would be

$$\begin{aligned} U(\text{Salary: } 81, \text{ Fun: } 60) &= 0.70(81) + 0.30(60) \\ &= 74.7 \end{aligned}$$

It is up to Sam to make an appropriate judgment about the relative importance of the two attributes. Although details on making this judgment are in Chapter 15, one important issue in making this judgment bears discussion here. Sam *must* take into consideration the ranges of the two attributes. Strictly speaking, the two weights should reflect the relative value of going from best to worst on each scale. That is, if Sam thinks that improving salary from \$2047.50 to \$2730 is three times as important as improving fun from Level 1 to Level 5, this judgment would imply weights $k_s = 0.75$ and $k_f = 0.25$.

Paying attention to the ranges of the attributes in assigning weights is crucial. Too often we are tempted to assign weights on the basis of vague claims that Attribute A (or its underlying objective) is worth three times as much as Attribute B. Suppose you are buying a car, though. If you are looking at cars that all cost about the same amount but their features differ widely, why should price play a role in your decision? It should have a low weight in the overall score. In the Texaco-Pennzoil case, we argued that we could legitimately consider only the objective of maximizing the settlement amount because its range was so wide; any other objectives would be overwhelmed by the importance of moving from worst to best on this one. In an overall score, the weight for settlement amount would be near 1, and the weight for any other attribute would be near zero.

Suppose that, after carefully considering the possible salary and summer-fun outcomes, Sam has come up with weights of 0.6 for salary and 0.4 for fun, reflecting a judgment that the range of possible salaries is 1.5 times as important as the range of possible summer-fun ratings. With these weights, we can collapse the consequence matrix in Figure 4.35 to get Figure 4.36. For example, if Sam chooses the forest job and the level of fun turns out to be Level 4, the overall score is $0.6(81) + 0.4(90) = 84.6$. The other endpoint values in Figure 4.36 can be found in the same way.

In these last two sections we have discussed some straightforward ways to make subjective ratings and trade-off assessments. These topics are treated more completely in Chapters 13, 15, and 16. For now you can rest assured that the techniques described here are fully compatible with those described in later chapters.

Analysis: Expected Values and Risk Profiles for Two Objectives

The decision tree in Figure 4.36 is now ready for analysis. The first thing we can do is fold back the tree to calculate expected values. Using the overall scores from Figure 4.36, the expected values are:

Figure 4.36

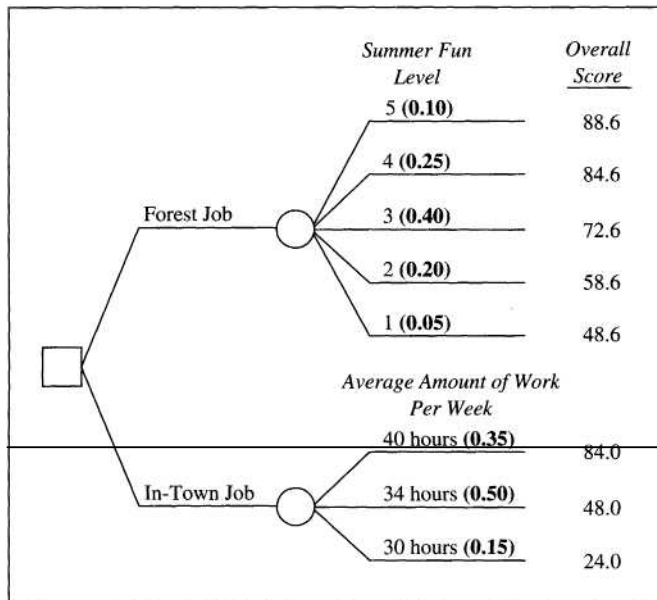
Decision tree with overall scores for summer-job example.

Weights used are $k_s = 0.60$ and $k_f = 0.40$.

For example, consider the forest job that has an outcome of

Level 4 on the fun scale. The rating for salary is 81, and the rating for fun is 90.

Thus, the overall score is $0.60(81) + 0.40(90) = 84.6$.



$$\begin{aligned}
 E(\text{Score for Forest Job}) &= 0.10(88.6) + 0.25(84.6) \\
 &\quad + 0.40(72.6) + 0.20(58.6) + 0.05(48.6) = \\
 &= 73.2
 \end{aligned}$$

$$\begin{aligned}
 E(\text{Score for In-Town Job}) &= 0.35(84) + 0.50(48) + 0.15(24) \\
 &= 57
 \end{aligned}$$

Can we also create risk profiles for the two alternatives? We can; the risk profiles would represent the uncertainty associated with the overall weighted score Sam will get from either job. To the extent that this weighted score is meaningful to Sam as a measure of overall satisfaction, the risk profiles will represent the uncertainty associated with Sam's overall satisfaction. Figures 4.37 and 4.38 show the risk profiles and cumulative risk profiles for the two jobs. Figure 4.38 shows that, given the ratings and the trade-off between fun and salary, the forest job stochastically dominates the in-town job in terms of the overall score. Thus, the decision may be clear for Sam at this point; given Sam's assessed probabilities, ratings, and the trade-off, the forest job is a better risk. (Before making the commitment, though, Sam may want to do some *sensitivity analysis*, the topic of Chapter 5; small changes in some of those subjective judgments might result in a less clear choice between the two.)

Two final caveats are in order regarding the risk profiles of the overall score. First, it is important to understand that the overall score is something of an artificial outcome; it is an amalgamation in this case of two rating scales. As indicated above, Figures 4.37 and 4.38 only makes sense to the extent that Sam is willing to interpret them as representing the uncertainty in the overall satisfaction from the two jobs.

Second, the stochastic dominance displayed by the forest job in Figure 4.38 is a relatively weak result; it relies heavily on Sam's assessed trade-off between the two attributes. A stronger result — one in which Sam could have confidence that the forest job is preferred regardless of his trade-off — requires that the forest job stochastically dominate the in-town job on each individual attribute. (Technically, however, even individual stochastic dominance is not quite enough; the risk profiles for the attributes must be combined into a single two-dimensional risk profile, or *bivariate probability distribution*, for each attribute. Then these two-dimensional risk profiles must be compared in much the same way we did with the single-attribute risk profiles. The good news is that as long as amount of work and amount of fun are *independent* (no arrow between these two chance nodes in the influence diagram in Figure 4.31), then finding that the same job stochastically dominates the other on each attribute guarantees that the same relationship holds in terms of the technically correct two-dimensional risk profile. Independence and stochastic dominance for multiple attributes will be discussed in Chapter 7.)

Figure 4.37
Risk profiles for summer jobs.





Figure 4.38
Cumulative risk profiles for summer jobs. The forest job stochastically dominates the in-town job.



Computer Programs for Decision Analysis

In this section we continue the discussion from Chapter 3 of personal-computer software for decision analysis, demonstrating how DPL and DATA perform the analysis described in this chapter. These programs do more than calculate expected values and create risk profiles, of course. In later chapters we will see how they perform sensitivity analysis and incorporate utility functions.

DPL

Figures 4.39 and 4.40 show the DPL influence diagram and decision tree, respectively, for the Texaco-Pennzoil case. Figure 4.39 is essentially the same as Figure 4.12. Figure 4.40, though, differs slightly from Figure 4.1. Notice that the "Final Court Decision" node and branches only appear once, at the end of the "Refuse" branch of the "Texaco Reaction" node. The "Final Court Decision" node is marked with an "a." A chance node with no branches but also marked with an "a" appears at the end of the "Refuse" branch of the "Pennzoil Reaction" node, indicating that from this point the structure of the tree is the same — identical nodes, branches, cash flows, and probabilities — as from the first "a." This shorthand approach for displaying similar subtrees can greatly simplify the presentation of complex decisions.

Figure 4.39
DPL's influence
diagram for Texaco-
Pennzoil decision.

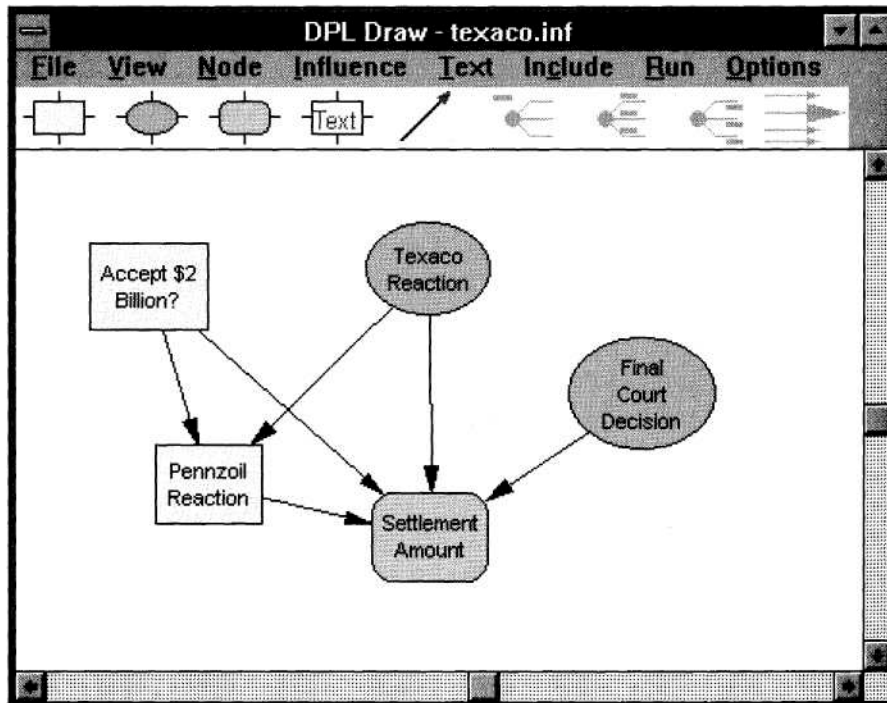
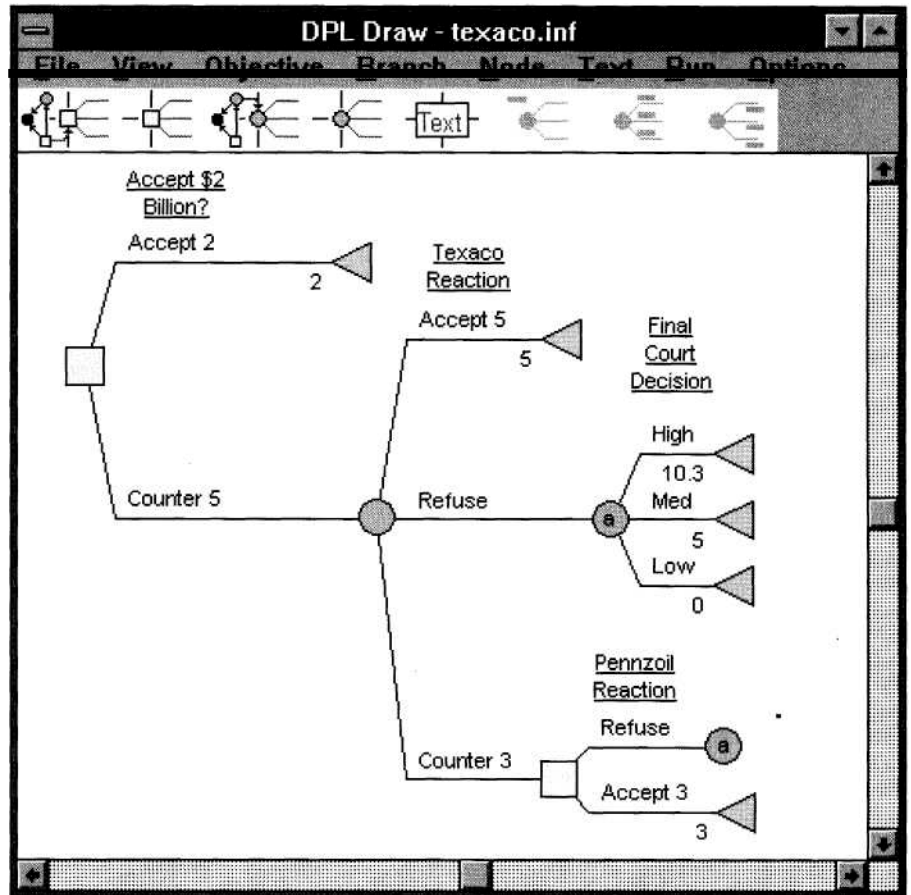


Figure 4.40
 DPL's decision tree for
 Texaco-Pennzoil.



With the problem fully structured, choosing the "Decision Analysis" command from the "Run" menu generates two new windows. First, the DPL Decision Policy window (Figure 4.41) displays the folded-back decision tree with expected values calculated and the optimal strategy shown with heavy lines. Second, DPL creates the optimal strategy's risk profile and displays it in the DPL Distributions window (Figure 4.42). An optional setting permits display of the risk profile in both regular and cumulative formats.

DATA

The "Roll Back" command on the "Analysis" menu in DATA does exactly what it suggests: It solves the decision tree. Figure 4.43 shows the folded-back decision tree for the summer-job decision as displayed by DATA. As in DPL, the optimal policy is highlighted, in this case taking the forest job. This analysis has taken into account the two attributes, summer fun and salary; DATA has a facility for identifying up to four different "payoffs" (consequences), any one of which can be calculated according to

Figure 4.41
DPL's solved decision tree for Texaco-Pennzoil decision.





Figure 4.42
Risk profile for optimal Texaco-Pennzoil strategy in DPL



a formula using defined variables whose values depend on the path taken. In this rendition of the problem, variables "Fun_Score" and "Salary_Score" have been defined and used to calculate the weighted "Overall" consequence. (DPL also allows the decision maker to specify and use variables to calculate probabilities and payoffs. The ability to define and use variables in representing a decision problem makes a program flexible and easy to use in complex decision situations.)

Figure 4.44 shows the cumulative risk profile for the forest job. This graph is created by choosing the "Probability Distribution" command from DATA's "Analysis" menu. As with DPL, the risk profile can be displayed in cumulative or regular format.

Figure 4.43
*Folded-back summer-
 job decision tree in
 DATA.*

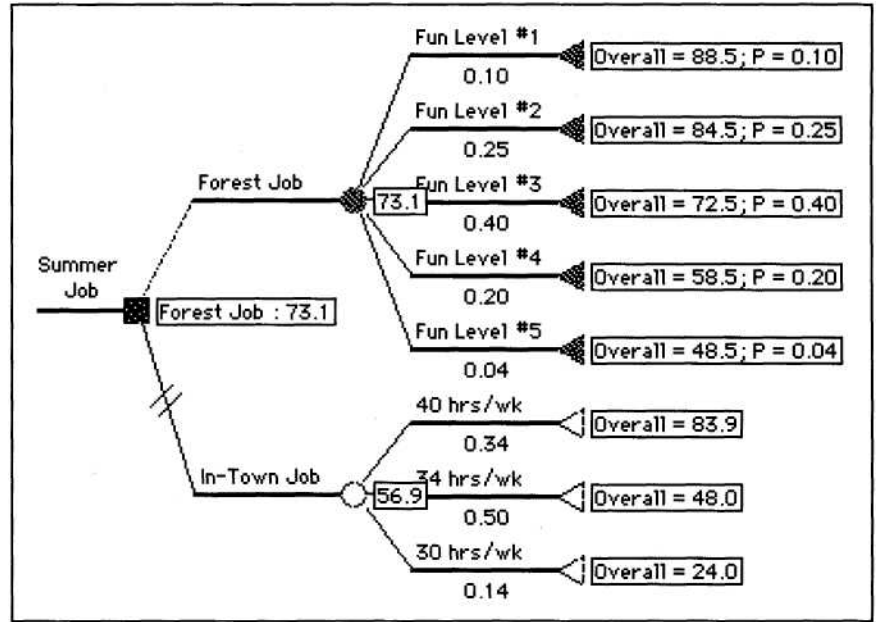



Figure 4.44
*Risk profile for forest
 job in DATA.*



Electronic Spreadsheets

No discussion of decision-analysis software would be complete without a discussion of electronic spreadsheets such as Microsoft Excel. Many decision problems revolve around accounting information. A spreadsheet environment is ideal for working with accounting data and building financial models for the possible outcomes and consequences in a decision situation. In this way, an analyst may use a spreadsheet to create subsidiary models to provide some of the inputs to a decision tree or influence diagram. In fact, both DPL and DATA have provisions for linking with Excel. Thus, it is possible to create financial models in Excel and access appropriate figures on the worksheet directly from DPL or DATA. Moreover, the links can be dynamic; changes made in the spreadsheet can be immediately reflected in the decision-tree or influence-diagram model.

Linking a spreadsheet with a specialized decision-analysis program is one of many ways to use electronic spreadsheets in decision analysis, and in future chapters we will see other uses. In addition to these specific uses, though, it is important to emphasize the usefulness and flexibility of electronic spreadsheets. Constructing decision models, calculating cash flows and expected values, collapsing decision trees to create risk profiles, performing sensitivity analysis, and more can all be done in a spreadsheet. Spreadsheet expertise pays off in many areas other than decision analysis; these programs can be used to do virtually anything that requires calculations. Of course, gaining the necessary expertise takes time and patience. Virtually all managers now have personal computers on their desks with the ability to run sophisticated spreadsheet programs, which suggests that aspiring managers would be well advised to become proficient in the use of this flexible tool.

SUMMARY

This chapter has demonstrated a variety of ways to use quantitative tools to make choices in uncertain situations. We first looked at the solution process for decision trees using expected value [or expected monetary value (EMV) when consequences are dollars]. This is the most straight-forward way to analyze a decision model; the algorithm for solving a decision tree is easy to apply, and expected values are easy to calculate.

We also explored the process of solving influence diagrams using expected values. To understand the solution process for influence diagrams, we had to look at their internal structures. In a sense, we had to fill in certain gaps left from Chapter 3 about how influence diagrams work. The solution procedure works out easily once we understand how the problem's numerical details are represented internally. The procedure for reducing nodes involves calculating expected values in a way that parallels the solution of a decision tree.

Risk profiles can be used to compare the riskiness of strategies and give comprehensive views of risks faced by a decision maker. Thus, risk profiles provide additional information to the decision maker trying to gain insight into the decision situation and the available alternatives. We also showed how cumulative risk profiles can be used to identify dominated alternatives.


The chapter ended with a discussion of how current decision-analysis computer programs perform the analysis described in the chapter. We also discussed the elec-

tronic spreadsheet as a flexible modeling tool for decision makers. The use of these tools can greatly enhance a decision maker's ability to model and analyze a wide variety of decision situations.

EXERCISES


- 4.1 Is it possible to solve a decision-tree version of a problem and an equivalent influence-diagram version and come up with different answers? If so, explain. If not, why not?
- 4.2 Explain in your own words what it means when one alternative stochastically dominates another.
- 4.3 The analysis of the Texaco-Pennzoil example shows that the EMV of counteroffering with \$5 billion far exceeds \$2 billion. Why might Liedtke want to accept the \$2 billion anyway? If you were Liedtke, what is the smallest offer from Texaco that you would accept?
- 4.4 Solve the decision tree in Figure 4.45.

Figure 4.45
Generic decision tree
for Exercise 4.4.




- 4.5 Draw and solve the influence diagram that corresponds to the decision tree in Figure 4.45.
- 4.6 Solve the decision tree in Figure 4.46. What principle discussed in Chapter 4 is illustrated by this decision tree?

Figure 4.46
Generic decision tree
for Exercise 4.6.




4.7 Which alternative is preferred in Figure 4.47? Do you have to do any calculations? Explain.



4.8 Solve the decision tree in Figure 4.48.

Figure 4.48
Generic decision tree
for Exercise 4.8.



- 4.9 Create risk profiles and cumulative risk profiles for all possible strategies in Figure 4.48. Is one strategy stochastically dominant? Explain.
- 4.10 Draw and solve the influence diagram that corresponds to the decision tree in Figure 4.48.
- 4.11 Explain why deterministic dominance is a special case of stochastic dominance.
- 4.12 Explain in your own words why it is important to consider the ranges of the consequences in determining a trade-off weight.
- 4.13 Solve the influence diagram for the umbrella problem shown in Figure 4.10.

QUESTIONS AND PROBLEMS

4.14 A real-estate investor has the opportunity to purchase an apartment complex. The apartment complex costs \$400,000 and is expected to generate net revenue (net after all operating and finance costs) of \$6000 per month. Of course, the revenue could vary because

the occupancy rate is uncertain. Considering the uncertainty, the revenue could vary from a low of -\$1000 to a high of \$10,000 per month. Assume that the investor's objective is to maximize the value of the investment at the end of 10 years.

- a Do you think the investor should buy the apartment complex or invest the \$400,000 in a 10-year certificate of deposit earning 9.5%? Why?
- b The city council is currently considering an application to rezone a nearby empty parcel of land. The owner of that land wants to build a small electronics-assembly plant. The proposed plant does not really conflict with the city's overall land use plan, but it may have a substantial long-term negative effect on the value of the nearby residential district in which the apartment complex is located. Because the city council currently is divided on the issue and will not make a decision until next month, the real-estate investor is thinking about waiting until the city council makes its decision.

If the investor waits, what could happen? What are the trade-offs that the investor has to make in deciding whether to wait or to purchase the complex now?

- c Suppose the investor could pay the seller \$1000 in earnest money now, specifying in the purchase agreement that if the council's decision is to approve the rezoning, the investor can forfeit the \$1000 and forego the purchase. Draw and solve a decision tree showing the investor's three options. Examine the alternatives for dominance. If you were the investor, which alternative would you choose? Why?

4.15 A stock market investor has \$500 to spend and is considering purchasing an option contract on 1000 shares of Apricot Computer. The shares themselves are currently selling for \$28.50 per share. Apricot is involved in a lawsuit, the outcome of which will be known within a month. If the outcome is in Apricot's favor, analysts expect Apricot's stock price to increase by \$5 per share. If the outcome is unfavorable, then the price is expected to drop by \$2.75 per share. The option costs \$500, and owning the option would allow the investor to purchase 1000 shares of Apricot stock for \$30 per share. Thus, if the investor buys the option and Apricot prevails in the lawsuit, the investor would make an immediate profit. Aside from purchasing the option, the investor could (1) do nothing and earn about 8% on his money, or (2) purchase \$500 worth of Apricot shares.

- a Construct cumulative risk profiles for the three alternatives, assuming Apricot has a 25% chance of winning the lawsuit. Can you draw any conclusions?
- b If the investor believes that Apricot stands a 25% chance of winning the lawsuit, should he purchase the option? What if he believes the chance is only 10%? How large does the probability have to be for the option to be worthwhile?

4.16 Johnson Marketing is interested in producing and selling an innovative new food processor. The decision they face is the typical "make or buy" decision often faced by manufacturers. On one hand, Johnson could produce the processor itself, subcontracting different sub-assemblies, such as the motor or the housing. Cost estimates in this case are as follows:

Alternative: Make Food Processor	
Cost Per Unit (\$)	Chance (%)
35.00	25
42.50	25
45.00	37
49.00	13

The company also could have the entire machine made by a subcontractor. The subcontractor, however, faces similar uncertainties regarding the costs and has provided Johnson Marketing with the following schedule of costs and chances:

Alternative: Buy Food Processor

Cost Per Unit (\$)	Chance (%)
37.00	10
43.00	40
46.00	30
50.00	20

If Johnson Marketing wants to minimize its expected cost of production in this case, should it make or buy? Construct cumulative risk profiles to support your recommendation. (*Hint*: Use care when interpreting the graph!)

4.17 Analyze the difficult decision situation that you identified in Problem 1.9. and structured in Problem 3.21. Be sure to examine alternatives for dominance. Does your analysis suggest any new alternatives?

4.18 Stacy Ennis eats lunch at a local restaurant two or three times a week. In selecting a restaurant on a typical workday, Stacy uses three criteria. First is to minimize the amount of travel time, which means that close-by restaurants are preferred on this attribute. The next objective is to minimize cost, and Stacy can make a judgment of the average lunch cost at most of the restaurants that would be considered. Finally, variety comes into play. On any given day, Stacy would like to go someplace different from where she has been in the past week.

Today is Monday, her first day back from a two-week vacation, and Stacy is considering the following six restaurants:

	Distance (Walking Time)	Average Price (\$)
Sam's Pizza	10	3.50
Sy's Sandwiches	9	2.85
Bubba's Italian Barbecue	7	6.50
Blue China Cafe	2	5.00
The Eating Place	2	7.50
The Excel-Soaring Restaurant	5	9.00

a If Stacy considers distance, price, and variety to be equally important (given the range of alternatives available), where should she go today for lunch? (*Hints*: Don't forget to convert both distance and price to similar scales, such as a scale from 0 to 100. Also, recall that Stacy has just returned from vacation; what does this imply for how the restaurants compare on the variety objective?)

b Given your answer to part a, where should Stacy go on Thursday?

CASE STUDIES

GPC'S NEW PRODUCT DECISION

The executives of the General Products Company (GPC) have to decide which of three products to introduce, A, B, or C. Product C is essentially a risk-free proposition, from which the company will obtain a net profit of \$1 million. Product B is considerably more risky. Sales may be high, with resulting net profit of \$8 million, medium with net profit of \$4 million, or low, in which case the company just breaks even. The probabilities for these outcomes are

$$P(\text{Sales High for B}) = 0.38$$

$$P(\text{Sales Medium for B}) = 0.12$$

$$P(\text{Sales Low for B}) = 0.50$$

Product A poses something of a difficulty; a problem with the production system has not yet been solved. The engineering division has indicated its confidence in solving the problem, but there is a slight (5%) chance that devising a workable solution may take a long time. In this event, there will be a delay in introducing the product and that delay will result in lower sales and profits. Another issue is the price for Product A. The options are to introduce it at either high or low price; the price would not be set until just before the product is to be introduced. Both of these issues have an impact on the ultimate net profit.

Finally, once the product is introduced, sales can be either high or low. If the company decides to set a low price, then low sales are just as likely as high sales. If the company sets a high price, the likelihood of low sales depends on whether the product was delayed by the production problem. If there was no delay and the company sets a high price, the probability is 0.4 that sales will be high. However, if there is a delay and the price is set high, the probability is only 0.3 that sales will be high. The following table shows the possible net profit figures (in millions) for Product A:

	Price	High Sales (\$ Million)	Low Sales (\$ Million)
Time delay	High	5.0	(0.5)
	Low	3.5	1.0
No delay	High	8.0	0.0
	Low	4.5	1.5

Questions

- 1 Draw an influence diagram for GPC's problem. Specify the possible outcomes and the probability distributions for each chance node. Specify the possible alternatives

for each decision node. Write out the complete table for the consequence node. (If possible, use a computer program for creating influence diagrams.)

- 2 Draw a complete decision tree for GPC. Solve the decision tree. What should GPC do? (If possible, do this problem using a computer program for creating and solving decision trees.)
- 3 Create cumulative risk profiles for each of the three products. Plot all three profiles on one graph. Can you draw any conclusions?
- 4 One of the executives of GPC is considerably less optimistic about Product B and assesses the probability of medium sales as 0.3 and the probability of low sales as 0.4. Based on expected value, what decision would this executive make? Should this executive argue about the probabilities? Why or why not? (*Hint*: Don't forget that probabilities have to add up to 1!)
- 5 Comment on the specification of chance outcomes and decision alternatives. Would this specification pass the clarity test? If not, what changes in the problem must be made in order to pass the clarity test?

SOUTHERN ELECTRONICS, PART I

Steve Sheffler is president, CEO, and majority stockholder of Southern Electronics, a small firm in the town of Silicon Mountain. Steve faces a major decision: Two firms, Big Red Business Machines and Banana Computer, are bidding for Southern Electronics.

Steve founded Southern 15 years ago, and the company has been extremely successful in developing progressive computer components. Steve is ready to sell the company (as long as the price is right!) so that he can pursue other interests. Last month, Big Red offered Steve \$5 million and 100,000 shares of Big Red stock (currently trading at \$50 per share and not expected to change substantially in the future). Until yesterday, Big Red's offer sounded good to Steve, and he had planned on accepting it this week. But a lawyer from Banana Computer called last week and indicated that Banana was interested in acquiring Southern Electronics. In discussions this past week, Steve has learned that Banana is developing a new computer, code-named EYF, that, if successful, will revolutionize the industry. Southern Electronics could play an important role in the development of the machine.

In their discussions, several important points have surfaced. First, Banana has said that it believes the probability that the EYF will succeed is 0.6, and that if it does, the value of Banana's stock will increase from the current value of \$30 per share. Although the future price is uncertain, Banana judges that, conditional on the EYF's success, the expected price of the stock is \$50 per share. If the EYF is not successful, the price will probably decrease slightly. Banana judges that if the EYF fails, Banana's share price will be between \$20 and \$30, with an expected price of \$25.

Yesterday Steve discussed this information with his financial analyst, who is an expert regarding the electronics industry and whose counsel Steve trusts completely. The analyst pointed out that Banana has an incentive to be very optimistic about the

EYF project. "Being realistic, though," said the analyst, "the probability that the EYF succeeds is only 0.4, and if it does succeed, the expected price of the stock would be only \$40 per share. On the other hand, I agree with Banana's assessment for the share price if the EYF fails."

Negotiations today have proceeded to the point where Banana has made a final offer to Steve of \$5 million and 150,000 shares of Banana stock. The company's representative has stated quite clearly that Banana cannot pay any more than this in a straight transaction. Furthermore, the representative claims, it is not clear why Steve will not accept the offer because it appears to them to be more valuable than the Big Red offer.

Questions

- 1 In terms of expected value, what is the least that Steve should accept from Banana? (This amount is called his *reservation price*.)
- 2 Steve obviously has two choices, to accept the Big Red offer or to accept the Banana offer. Draw an influence diagram representing Steve's decision. (If possible, do this problem using a computer program for structuring influence diagrams.)
- 3 Draw and solve a complete decision tree representing Steve's decision. (If possible, do this problem using a computer program for creating and solving decision trees.)
- 4 Why is it that Steve cannot accept the Banana offer as it stands?

SOUTHERN ELECTRONICS, PART II

Steve is well aware of the difference between his probabilities and Banana's, and he realizes that because of this difference, it may be possible to design a contract that benefits both parties. In particular, he is thinking about put options for the stock. A put option gives the owner of the option the right to sell an asset at a specific price. (For example, if you own a put option on 100 shares of General Motors (GM) with an exercise price of \$75, you could sell 100 shares of GM for \$75 per share before the expiration date of the option. This would be useful if the stock price fell below \$75.) Steve reasons that if he could get Banana to include a put option on the stock with an exercise price of \$30, then he would be protected if the EYF failed.

Steve proposes the following deal: He will sell Southern Electronics to Banana for \$530,000 plus 280,000 shares of Banana stock and a put option that will allow him to sell the 280,000 shares back to Banana for \$30 per share any time within the next year (during which time it will become known whether the EYF succeeds or fails).

Questions

- 1 Calculate Steve's expected value for this deal. Ignore tax effects and the time value of money.

- 2 The cost to Banana of their original offer was simply

$$\$5,000,000 + 150,000(\$30) = \$9,500,000$$

Show that the expected cost to Banana of Steve's proposed deal is less than \$9.5 million, and hence in Banana's favor. Again, ignore tax effects and the time value of money.

STRENLAR

Fred Wallace scratched his head. By this time tomorrow he had to have an answer for Joan Sharkey, his former boss at Plastics International (PI). The decision was difficult to make. It involved how he would spend the next 10 years of his life.

Four years ago, when Fred was working at PI, he had come up with an idea for a revolutionary new polymer. A little study — combined with intuition, hunches, and educated guesses — had convinced him that the new material would be extremely strong for its weight. Although it would undoubtedly cost more than conventional materials, Fred discovered that a variety of potential uses existed in the aerospace, automobile manufacturing, robotics, and sporting goods industries.

When he explained his idea to his supervisors at PI, they had patiently told him that they were not interested in pursuing risky new projects. His appeared to be even riskier than most because, at the time, many of the details had not been fully worked out. Furthermore, they pointed out that efficient production would require the development of a new manufacturing process. Sure, if that process proved successful, the new polymer could be a big hit. But without that process the company simply could not provide the resources Fred would need to develop his idea into a marketable product.

Fred did not give up. He began to work at home on his idea, consuming most of his evenings and weekends. His intuition and guesses had proven correct, and after some time he had worked out a small-scale manufacturing process. With this process, he had been able to turn out small batches of his miracle polymer, which he dubbed Strenlar. At this point he quietly began to assemble some capital. He invested \$100,000 of his own, managed to borrow another \$200,000, and quit his job at PI to devote his time to Strenlar.

That was 15 months ago. In the intervening time he had made substantial progress. The product was refined, and several customers eagerly awaited the first production run. A few problems remained to be solved in the manufacturing process, but Fred was 80% sure that these bugs could be worked out satisfactorily. He was eager to start making profits himself; his capital was running dangerously low. When he became anxious, he tried to soothe his fears by recalling his estimate of the project's potential. His best guess was that sales would be approximately \$35 million over 10 years, and that he would net some \$8 million after costs.

Two weeks ago, Joan Sharkey at PI had surprised him with a telephone call and had offered to take Fred to lunch. With some apprehension, Fred accepted the offer. He had always regretted having to leave PI, and was eager to hear how his friends were doing. After some pleasantries, Joan came to the point.

"Fred, we're all impressed with your ability to develop Strenlar on your own. I guess we made a mistake in turning down your offer to develop it at PI. But we're interested in helping you out now, and we can certainly make it worth your while. If you will grant PI exclusive rights to Strenlar, we'll hire you back at, say \$40,000 a year, and we'll give you a 2.5 percent royalty on Strenlar sales. What do you say?"

Fred didn't know whether to laugh or become angry. "Joan, my immediate reaction is to throw my glass of water in your face! I went out on a limb to develop the product, and now you want to capitalize on my work. There's no way I'm going to sell out to PI at this point!"

The meal proceeded, with Joan sweetening the offer gradually, and Fred obstinately refusing. After he got back to his office, Fred felt confused. It would be nice to work at PI again, he thought. At least the future would be secure. But there would never be the potential for the high income that was possible with Strenlar. Of course, he thought grimly, there was still the chance that the Strenlar project could fail altogether.

At the end of the week, Joan called him again. PI was willing to go either of two ways. The company could hire him for \$50,000 plus a 6% royalty on Strenlar gross sales. Alternatively, PI could pay him a lump sum of \$500,000 now plus options to purchase up to 70,000 shares of PI stock at the current price of \$40 any time within the next three years. No matter which offer Fred accepted, PI would pay off Fred's creditors and take over the project immediately. After completing development of the manufacturing process, PI would have exclusive rights to Strenlar. Furthermore, it turned out that PI was deadly serious about this game. If Fred refused both of these offers, PI would file a lawsuit claiming rights to Strenlar on the grounds that Fred had improperly used PI's resources in the development of the product.

Consultation with his attorney just made him feel worse. After reviewing Fred's old contract with PI, the attorney told him that there was a 60% chance that he would win the case. If he won the case, PI would have to pay his court costs. If he lost, his legal fees would amount to about \$20,000.

Fred's accountant helped him estimate the value of the stock options. First, the exercise date seemed to pose no problem; unless the remaining bugs could not be worked out, Strenlar should be on the market within 18 months. If PI were to acquire the Strenlar project and the project succeeded, PI's stock would go up to approximately \$52. On the other hand, if the project failed, the stock price probably would fall slightly to \$39.

As Fred thought about all of the problems he faced, he was quite disturbed. On one hand, he yearned for the comradery he had enjoyed at PI four years ago. He also realized that he might not be cut out to be an entrepreneur. He reacted unpleasantly to the risk he currently faced. His physician had warned him that he may be developing hypertension and had tried to persuade him to relax more. Fred knew that his health was important to him, but he had to believe that he would be able to weather the tension of getting Strenlar onto the market. He could always relax later, right? He

sighed as he picked up a pencil and pad of paper to see if he could figure out what he should tell Joan Sharkey.

Question

- 1 Do a complete analysis of Fred's decision. Your analysis should include at least structuring the problem with an influence diagram, drawing and solving a decision tree, creating risk profiles, and checking for stochastic dominance. What do you think Fred should do? Why? (*Hint:* This case will require you to make certain assumptions in order to do a complete analysis. State clearly any assumptions you make, and be careful that the assumptions you make are both reasonable and consistent with the information given in the case. You may want to analyze your decision model under different sets of assumptions. Do not forget to consider issues such as the time value of money, riskiness of the alternatives, and so on.)

JOB OFFERS

Robin Pinelli is considering three job offers. In trying to decide which to accept, Robin has concluded that three objectives are important in this decision. First, of course, is to maximize disposable income — the amount left after paying for housing, utilities, taxes, and other necessities. Second, Robin likes cold weather and enjoys winter sports. The third objective relates to the quality of the community. Being single, Robin would like to live in a city with a lot of activities and a large population of single professionals.

Developing attributes for these three objectives turns out to be relatively straightforward. Disposable income can be measured directly by calculating monthly take-home pay minus average monthly rent (being careful to include utilities) for an appropriate apartment. The second attribute is annual snowfall. For the third attribute, Robin has located a magazine survey of large cities that scores those cities as places for single professionals to live. Although the survey is not perfect from Robin's point of view, it does capture the main elements of her concern about the quality of the singles community and available activities. Also, all three of the cities under consideration are included in the survey.

Here are descriptions of the three job offers:

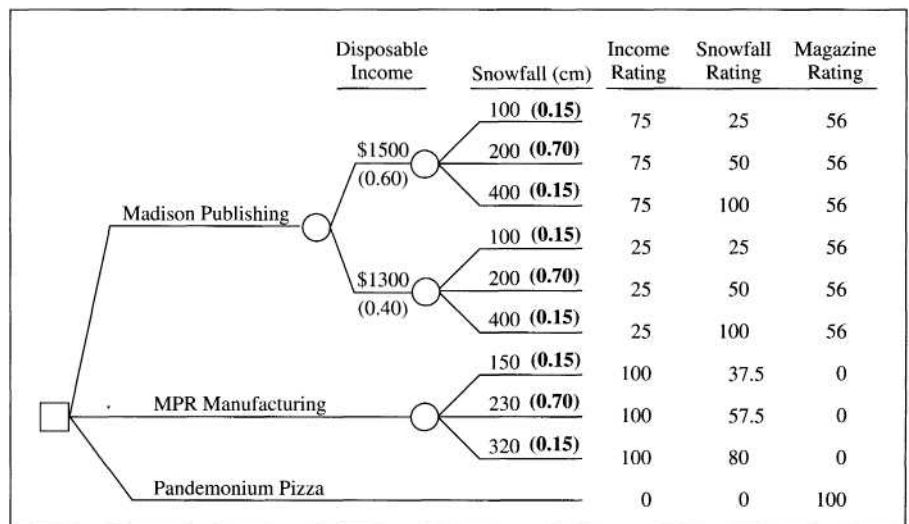
- 1 MPR Manufacturing in Flagstaff, Arizona. Disposable income estimate: \$1600 per month. Snowfall range: 125 to 320 cm per year. Magazine score: 50 (out of 100).
- 2 Madison Publishing in St. Paul, Minnesota. Disposable income estimate: \$1300 to \$1500 per month. (This uncertainty here is because Robin knows there is a wide variety in apartment rental prices and will not know what is appropriate and available until spending some time in the city.) Snowfall range: 100 to 400 cm per year. Magazine score: 75.
- 3 Pandemonium Pizza in San Francisco, California. Disposable income estimate: \$1200 per month. Snowfall range: negligible. Magazine score: 95.

Robin has created the decision tree in Figure 4.49 to represent the situation. The uncertainty about snowfall and disposable income are represented by the chance nodes as Robin has included them in the tree. The ratings in the consequence matrix are such that the worst consequence has a rating of zero points and the best has 100.

Questions

- 1 Verify that the ratings in the consequence matrix are proportional scores (that is, that they were calculated the same way we calculated the ratings for salary in the summer-fun example in the chapter).
- 2 Comment on Robin's choice of annual snowfall as a measure for the cold-weather-winter-sports attribute. Is this a good measure? Why or why not?
- 3 After considering the situation, Robin concludes that the quality of the city is most important, the amount of snowfall is next, and the third is income. (Income is important, but the variation between \$1200 and \$1600 is not enough to make much difference to Robin.) Furthermore, Robin concludes that the weight for the magazine rating in the consequence matrix should be 1.5 times the weight for the snowfall rating and three times as much as the weight for the income rating. Use this information to calculate the weights for the three attributes and to calculate overall scores for all of the end branches in the decision tree.
- 4 Analyze the decision tree using expected values. Calculate expected values for the three measures as well as for the overall score.
- 5 Do a risk-profile analysis of the three cities. Create risk profiles for each of the three attributes as well as for the overall score. Does any additional insight arise from this analysis?
- 6 What do you think Robin should do? Why?

Figure 4.49
Robin Pinelli's decision tree.



SS KUNIANG, PART II

This case asks you to find the optimal amount for NEES to bid for the SS *Kuniang* (page 97). Before doing so, though, you need additional details. Regarding the Coast Guard's (CG) salvage judgment, NEES believes that the following probabilities are an appropriate representation of its uncertainty about the salvage-value judgment:

$$P(\text{CG judgment} = \$9 \text{ million}) = 0.185$$

$$P(\text{CG judgment} = \$4 \text{ million}) = 0.630$$

$$P(\text{CG judgment} = \$1.5 \text{ million}) = 0.185$$

The obscure-but-relevant law required that NEES pay an amount (including both the winning bid and refitting cost) at least 1.5 times the salvage value for the ship in order to use it for domestic shipping. For example, if NEES bid \$3.5 million and won, followed by a CG judgment of \$4 million, then NEES would have to invest at least \$2.5 million more: $\$3.5 + \$2.5 = \$6 = \4×1.5 . Thus, assuming NEES submits the winning bid, the total investment amount required is either the bid or 1.5 times the CG judgment, whichever is greater.

As for the probability of submitting the highest bid, recall that winning is a function of the size of the bid; a bid of \$3 million is sure to lose, and a bid of \$10 million is sure to win. For this problem, we can model the probability of winning (P) as a linear function of the bid: $P = (\text{Bid} - \$3 \text{ million}) / (\$7 \text{ million})$.

Finally, NEES's values of \$18 million for the new ship and \$15 million for the tug-barge alternatives are adjusted to reflect differences in age, maintenance, operating costs, and so on. The two alternatives provide equivalent hauling capacity. Thus, at \$15 million, the tug-barge combination appears to be the better choice.

Questions

- 1 Reasonable bids may fall anywhere between \$3 and \$10 million. Some bids, though, have greater expected values and some less. Describe a strategy you can use to find the optimal bid, assuming that NEES's objective is to minimize the cost of acquiring additional shipping capacity. (*Hint:* This question just asks you to describe an approach to finding the optimal bid.)
- 2 Use your structure of the problem (or one supplied by the instructor), along with the details supplied above, to find the optimal bid.

REFERENCES

The solution of decision trees as presented in this chapter is commonly found in textbooks on decision analysis, management science, and statistics. The decision-analysis

texts listed at the end of Chapter 1 can provide more guidance in the solution of decision trees if needed. In contrast, the material presented here on the solution of influence diagrams is relatively new. For additional basic instruction in the construction and analysis of decisions using influence diagrams, the user's manual for DPL and other influence-diagram programs can be helpful.

The solution algorithm presented here is based on Shachter (1986). The fact that this algorithm deals with a decision problem in a way which corresponds to solving a symmetric decision tree means that the practical upper limit for the size of an influence diagram that can be solved using the algorithm is relatively small. Recent work has explored a variety of ways to exploit asymmetry in decision models and to solve influence diagrams and related representations more efficiently (Call and Miller 1990; Covaliu and Oliver (1995); Smith et al. 1993; Shenoy 1993).

An early and quite readable article on risk profiles is that by Hertz (1964). We have developed them as a way to examine the riskiness of alternatives in a heuristic way and also as a basis for examining alternatives in terms of deterministic and stochastic dominance. Stochastic dominance itself is an important topic in probability. Bunn (1984) gives a good introduction to stochastic dominance. Whitmore and Findlay (1978) and Levy (1992) provide thorough reviews of stochastic dominance.

Our discussion of assigning rating points and trade-off rates is necessarily brief in Chapter 4. These topics are covered in depth in Chapters 13 to 16. In the meantime, interested readers can get more information from Keeney (1992) and Keeney and Raiffa (1976).

Bodily, S. E. (1985) *Modern Decision Making*. New York: McGraw-Hill.

Bunn, D. (1984) *Applied Decision Analysis*. New York: McGraw-Hill.

Call, H., and W. Miller (1990) "A Comparison of Approaches and Implementations for Automating Decision Analysis." *Reliability Engineering and System Safety*, 30, 115-162.

Covaliu, Z., and R. Oliver (1995) "Representation and Solution of Decision Problems Using Sequential Decision Diagrams." *Management Science*, 41.

Hertz, D. B. (1964) "Risk Analysis in Capital Investment." *Harvard Business Review*. Reprinted in *Harvard Business Review*, September-October, 1979, 169-181.

Keeney, R. L. (1992) *Value-Focused Thinking*. Cambridge, MA: Harvard University Press.

Keeney, R., and H. Raiffa (1976) *Decisions with Multiple Objectives*. New York: Wiley.

Levy, H. (1992) "Stochastic Dominance and Expected Utility: Survey and Analysis." *Management Science*, 38, 555-593.

Shachter, R. (1986) "Evaluating Influence Diagrams." *Operations Research*, 34, 871-882.

Shenoy, P. (1993) "Valuation Network Representation and Solution of Asymmetric Decision Problems." Working paper, School of Business, University of Kansas, Lawrence.

Smith, J., S. Holtzman, and J. Matheson (1993) "Structuring Conditional Relationships in Influence Diagrams." *Operations Research*, 41, 280-297.

Whitmore, G. A., and M. C. Findlay (1978) *Stochastic Dominance*. Lexington, MA: Heath.

EPILOGUE

What happened with Texaco and Pennzoil? You may recall that in April of 1987 Texaco offered a \$2 billion settlement. Hugh Liedtke turned down the offer. Within days of that decision, and only one day before Pennzoil began to file liens on Texaco's assets, Texaco filed for protection from creditors under Chapter 11 of the federal bankruptcy code, fulfilling its earlier promise. In the summer of 1987, Pennzoil submitted a financial reorganization plan on Texaco's behalf. Under their proposal, Pennzoil would receive approximately \$4.1 billion, and the Texaco shareholders would be able to vote on the plan. Finally, just before Christmas 1987, the two companies agreed on a \$3 billion settlement as part of Texaco's financial reorganization.

CHAPTER 5

Sensitivity Analysis

The idea of sensitivity analysis is central to the structuring and solving of decision models using decision-analysis techniques. In this chapter we will discuss sensitivity-analysis issues, think about how sensitivity analysis relates to the overall decision-modeling strategy, and introduce a variety of graphical sensitivity-analysis techniques.

The main example for this chapter is a hypothetical one in which the owner of a small airline considers expanding his fleet.

EAGLE AIRLINES

Dick Carothers, president of Eagle Airlines, had been considering expanding his operation, and now the opportunity was available. An acquaintance had put him in contact with the president of a small airline in the Midwest that was selling an airplane. Many aspects of the situation needed to be thought about, however, and Carothers was having a hard time sorting them out.

Eagle Airlines owned and operated three twin-engine aircraft. With this equipment, Eagle provided both charter flights and scheduled commuter service among several communities in the eastern United States. Scheduled flights constituted approximately 50% of Eagle's flights, averaging only 90 minutes of flying time and a distance of some 300 miles. The remaining 50% of flights were chartered. The mixture of charter flights and short scheduled flights had proved profitable, and

Carothers felt that he had found a niche for his company. He was aching to increase the level of service, especially in the area of charter flights, but this was impossible without more aircraft.

A Piper Seneca was for sale at a price of \$95,000, and Carothers figured that he could buy it for between \$85,000 and \$90,000. This twin-engine airplane had been maintained according to FAA regulations. In particular, the engines were almost new, with only 150 hours of operation since a major overhaul. Furthermore, having been used by another small commercial charter service, the Seneca contained all of the navigation and communication equipment that Eagle required. There were seats for five passengers and the pilot, plus room for baggage. Typical airspeed was approximately 175 nautical miles per hour (knots), or 200 statute miles per hour (mph). Operating cost was approximately \$245 per hour, including fuel, maintenance, and pilot salary. Annual fixed costs included insurance (\$20,000) and finance charges. Carothers figured that he would have to borrow some 40% of the money required, and he knew that the interest rate would be two percentage points above the prime rate (currently 9.5% but subject to change). Based on his experience at Eagle, Carothers knew that he could arrange charters for \$300 to \$350 per hour or charge a rate of approximately \$100 per person per hour on scheduled flights. He could expect on average that the scheduled flights would be half full. He hoped to be able to fly the plane for up to 1000 hours per year, but realized that 800 might be more realistic. In the past his business had been approximately 50% charter flights, but he wanted to increase that percentage if possible.

The owner of the Seneca has told Carothers that he would either sell the airplane outright or sell Carothers an option to purchase it within a year at a specified price. (The current owner would continue to operate the plane during the year.) Although the two had not agreed on a price for the option, the discussions had led Carothers to believe that the option would cost between \$2500 and \$4000. Of course, he could always invest his cash in the money market and expect to earn about 8%.

As Carothers pondered this information, he realized that many of the numbers he was using were estimates. Furthermore, some were within his control (for example, the amount financed and prices charged) while others, such as the cost of insurance or the operating cost, were not. How much difference did these numbers make? What about the option? Was it worth considering? Last, but not least, did he really want to expand the fleet? Or was there something else that he should consider?

Sensitivity Analysis: A Modeling Approach

Sensitivity analysis answers the question, "What makes a difference in this decision?" Returning to the idea of requisite decision models discussed in Chapter 1, you may recall that such a model is one whose form and content are just sufficient to solve a particular problem. That is, the issues that are addressed in a requisite decision model are the ones that matter, and those issues left out are the ones that do not

matter. Determining what matters and what does not requires incorporating sensitivity analysis throughout the modeling process.

No "optimal" sensitivity-analysis procedure exists for decision analysis. To a great extent, model building is an art. Because sensitivity analysis is an integral part of the modeling process, its use as part of the process also is an art. Thus, in this chapter we will discuss the philosophy of model building and how sensitivity analysis helps with model development. Several sensitivity-analysis tools are available, and we will see how they work in the context of the Eagle Airlines example.

Problem Identification and Structure

The flowchart of the decision-analysis process in Figure 1.1 shows that sensitivity analysis can lead the decision maker to reconsider the very nature of the problem. The question that we ask in performing sensitivity analysis at this level is, "Are we solving the right problem?" The answer does not require quantitative analysis, but it does demand careful thought and introspection about the appropriate decision context. Why is this an important sensitivity-analysis concern? The answer is quite simple: Answering a different question, addressing a different problem, or satisfying different objectives can lead to a very different decision.

Solving the wrong problem sometimes is called an "error of the third kind." The terminology contrasts this kind of a mistake with Type I and Type II errors in statistics, where incorrect conclusions are drawn regarding a particular question. An error of the third kind, or Type III error, implies that the wrong question was asked; in terms of decision analysis, the implication is that an inappropriate decision context was used, and hence the wrong problem was solved.

Examples of Type III errors abound; we all can think of times when a symptom was treated instead of a cause. Consider lung disease. Researchers and physicians have developed expensive medical treatments for lung disease, the objective apparently being to reduce the suffering of lung-disease patients. If the fundamental objective is to reduce suffering from lung disease in general, however, these treatments are not as effective as antismoking campaigns. We can, in fact, broaden the context further. Is the objective really to reduce patient suffering? Or is it to reduce discomfort in general, including patient suffering as well as the discomfort of nonsmokers exposed to second-hand smoke? Considering the broader problem suggests an entirely different range of options.

For another example, think about a farmer who considers using expensive sprays in the context of deciding how to control pests and disease in an orchard. To a great extent, the presence of pests and disease in orchards result from the practice of monoculture — that is, growing a lot of one crop rather than a little each of many crops. A monoculture does not promote a balanced ecological system in which diseases and pests are kept under control naturally. Viewed from this broader perspective, the farmer might want to consider new agricultural practices rather than relying exclusively on sprays. Admittedly a long-term project, this requires the development

of efficient methods for growing, harvesting, and distributing crops that are grown on a smaller scale.

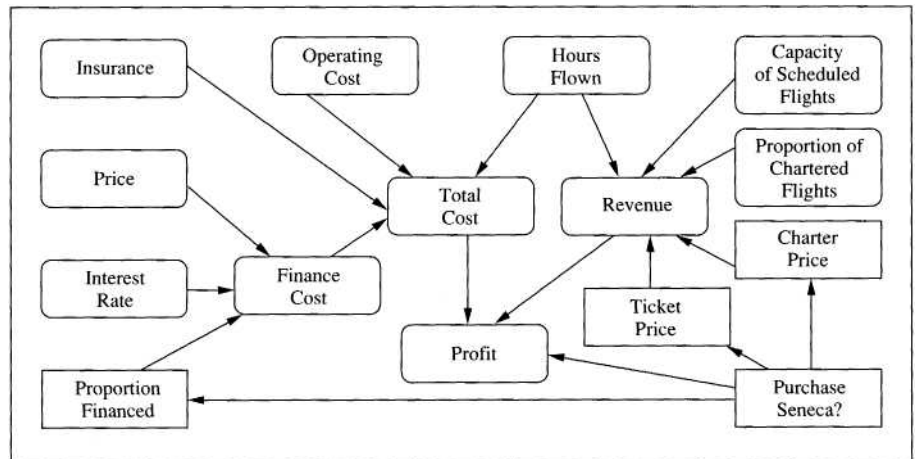
How can one avoid a Type III error? The best solution is simply to keep asking whether the problem on the surface is the real problem. Is the decision context properly specified? What exactly is the "unscratched itch" that the decision maker feels? In the case of Eagle Airlines, Carothers appears to be eager to expand operations by acquiring more aircraft. Could he "scratch his itch" by expanding in a different direction? In particular, even though he, like many pilots, may be dedicated to the idea of flying for a living, it might be wise to consider the possibility of helping his customers to communicate more effectively at long distance. To some extent, efficient communication channels such as those provided by computer links and facsimile service, coupled with an air cargo network, can greatly reduce the need for travel. Pursuing ideas such as these might satisfy Carothers's urge to expand while providing a more diversified base of operations. So the real question may be how to satisfy Carothers's desires for expansion rather than simply how to acquire more airplanes.

We also can talk about sensitivity analysis in the context of problem structuring. Problem 3.12 gave an example in a medical context in which a decision might be sensitive to the structure. In this situation, the issue is the inclusion of a more complete description of outcomes; coronary bypass surgery can lead to complications that require long and painful treatment. Inclusion of this outcome in a decision tree might make surgery appear considerably less appealing. Von Winterfeldt and Edwards (1986) describe a problem involving the setting of standards for pollution from oil wells in the North Sea. This could have been structured as a standard regulatory problem: Different possible standards and enforcement policies made up the alternatives, and the objective was to minimize pollution while maintaining efficient oil production. The problem, however, was perhaps more appropriately structured as a competitive situation in which the players were the regulatory agency, the industry, and the potential victims of pollution. This is an example of how a decision situation might be represented in a variety of different ways. Sensitivity analysis can aid the resolution of the problem of multiple representations by helping to identify the appropriate perspective on the problem as well as by identifying the specific issues that matter to the decision maker.

Is problem structuring an issue in the Eagle Airlines case? In this case, the alternatives are to purchase the airplane, the option, or neither. Although Carothers might consider a variety of fundamental objectives, such as company growth or increased influence in the community, in the context of deciding whether to purchase the Seneca, it seems reasonable for him to focus on one objective: maximize profit. Carothers could assess the probabilities associated with the various unknown quantities such as operating costs, amount of business, and so on. Thus, it appears that a straightforward decision tree or influence diagram may do the trick.

Figure 5.1 shows an initial influence diagram for Eagle Airlines. Note that the diagram consists entirely of decision nodes and rounded rectangles. "Profit" is obviously the consequence node, and "Finance Cost," "Total Cost," and "Revenue" are intermediate-calculation nodes. All of the other rounded rectangles ("Interest Rate," "Price," "Insurance," "Operating Cost," "Hours Flown," "Capacity of Scheduled Flights," "Proportion of Chartered Flights") represent inputs to the calculations, and for now we represent these inputs as being constant. (Thus, in Figure 5.1 you can see the three dif-

Figure 5.1
Influence diagram
representing the Eagle
Airlines decision.



ferent roles — constants, intermediate calculations, and consequences — that rounded rectangles can play. Although these different roles may seem confusing, the basic idea is the same in each case; for any variable represented by a rounded rectangle, as soon as you know what its inputs are, you can calculate the value of the variable. In the case of the constants, there are no inputs, and so there is no calculation to do!)

Table 5.1 provides a description of the input and decision variables. This table also includes estimates (base values) and reasonable upper and lower bounds. The upper and lower bounds represent Carothers's ideas about how high and how low each of these variables might be. He might specify upper and lower bounds as absolute extremes, beyond which he is absolutely sure that the variable cannot fall. Another approach would be to specify the bounds such that he would be "very surprised" (a 1-in-10 chance, say) that the variable would fall outside the bounds.

The "Base Value" column in Table 5.1 indicates Carothers's initial guess regarding the 10 input variables. We can use these to make an estimate of annual profit (ignoring taxes for simplicity). The annual profit would be the total annual revenue minus the total annual cost:

$$\begin{aligned}
 \text{Total Revenue} &= \text{Revenue from Charters} + \text{Revenue from Scheduled Flights} \\
 &= (\text{Charter Proportion} \times \text{Hours Flown} \times \text{Charter Price}) + [(1 - \\
 &\quad \text{Charter Proportion}) \times \text{Hours Flown} \times \text{Ticket Price} \times \text{Number of} \\
 &\quad \text{Passenger Seats} \times \text{Capacity of Scheduled Flights}] \\
 &= (0.5 \times 800 \times \$325) + (0.5 \times 800 \times \$100 \times 5 \times 0.5) \\
 &= \$230,000 \\
 \text{Total Cost} &= (\text{Hours Flown} \times \text{Operating Cost}) + \\
 &\quad \text{Insurance} + \text{Finance Cost} \\
 &= (\text{Hours Flown} \times \text{Operating Cost}) + \text{Insurance} + \\
 &\quad (\text{Price} \times \text{Proportion Financed} \times \text{Interest Rate}) \\
 &= (800 \times \$245) + \$20,000 + (\$87,500 \times 0.4 \times 11.5\%) \\
 &= \$220,025
 \end{aligned}$$

Table 5.1

Input variables and ranges of possible values for Eagle Airlines aircraft-purchase decision.

Variable	Base Value	Lower Bound	Upper Bound
Hours Flown	800	500	1000
Charter Price/Hour	\$325	\$300	\$350
Ticket Price/Hour	\$100	\$95	\$108
Capacity of Scheduled Flights	50%	40%	60%
Proportion of Chartered Flights	0.50	0.45	0.70
Operating Cost/Hour	\$245	\$230	\$260
Insurance	\$20,000	\$18,000	\$25,000
Proportion Financed	0.40	0.30	0.50
Interest Rate	11.5%	10.5%	13%
Purchase Price	\$87,500	\$85,000	\$90,000


Thus, using the base values, Carothers's annual profit is estimated to be $\$230,000 - \$220,025 = \$9975$. This represents a return of approximately 19% on his investment of $\$52,500$ (60% of the purchase price).

One-Way Sensitivity Analysis

The sensitivity-analysis question in the Eagle airlines case is, what variables really make a difference in terms of the decision at hand? For example, do different possible interest rates really matter? Does it matter that we can set the ticket price? If Hours Flown changes by some amount, how much impact is there on Profit? We can begin to address questions like these with one-way sensitivity analysis.

Let us consider Hours Flown. From Table 5.1, we see that Carothers is not at all sure what Hours Flown might turn out to be, and that it can vary from 500 to 1000 hours. What does this imply for Profit? The simplest way to answer this question is with a one-way sensitivity graph as in Figure 5.2. The upward-sloping line in Figure 5.2 shows profit as Hours Flown varies from 500 to 1000; to create this line, we have substituted different values for Hours Flown into the calculations detailed above. The horizontal line represents the amount of money (\$4200) that Carothers could earn from the money market. The point where these lines cross is the threshold at which the two alternatives each yield the same profit (\$4200), which occurs when Hours Flown equals 664. The heavy line indicates the maximum profit Carothers could obtain at different values of Hours Flown, and the different segments of this line is associated with different strategies (buy the Seneca versus invest in the money market). The fact that Carothers believes that Hours Flown could be above or below 664 suggests that this is a crucial variable and that he may need to think more carefully about the uncertainty associated with it.

Figure 5.2
One-way sensitivity
analysis of hours
flown.



Tornado Diagrams

A *tornado diagram* allows us to compare one-way sensitivity analysis for many input variables at once. Suppose we take each input variable in Table 5.1 and "wiggle" that variable between its high and low values to determine how much change is induced in Profit. Figure 5.3 graphically shows how annual profit varies as the input variables are independently wiggled between the high and low values. For instance, with everything else held at the base value, setting Capacity of Scheduled Flights at 0.4 instead of 0.5 implies a loss of \$10,025. That is, plug all the base values into the revenue equation above, except use 0.4 for Capacity of Scheduled Flights:


$$\begin{aligned} \text{Total Revenue} &= \text{Revenue from Charters} + \text{Revenue from Scheduled Flights} = \\ &= (\text{Charter Proportion} \times \text{Hours Flown} \times \text{Charter Price}) + [(1 - \\ &\quad \text{Charter Proportion}) \times \text{Hours Flown} \times \text{Ticket Price} \times \text{Number of} \\ &\quad \text{Passenger Seats} \times \text{Capacity on Scheduled Flights}] = (0.5 \times 800 \times \\ &\quad \$325) + (0.5 \times 800 \times \$100 \times 5 \times 0.4) = \$210,000 \end{aligned}$$

Nothing in the cost equation changes, and so cost still is estimated as \$220,025. The estimated loss is just the difference between cost and revenue: $\$210,000 - \$220,025 = -\$10,025$. This is plotted on the graph as the left end of the bar labeled Capacity of Scheduled Flights. On the other hand, setting Capacity of Scheduled Flights at the high end of its range, 0.6, leads to a profit of \$29,975. (Again, plug all of the same values into the revenue equation, but use 0.6 for capacity.) Thus, the right end of the capacity bar is at \$29,975.

We follow this same procedure for each input variable. The length of the bar for any given variable represents the extent to which annual profit is sensitive to this

Figure 5.3

Tornado diagram for the Eagle Airlines case. The bars represent the range for the annual profit when the specified quantity is varied from one end of its range to the other, keeping all other variables at their base values.



variable. The graph is laid out so that the most sensitive variable — the one with the longest bar — is at the top, and the least sensitive is at the bottom. With the bars arranged in this order, it is easy to see why the graph is called a tornado diagram.

The vertical line at \$4,200 represents what Carothers could make on his investment if he left his \$52,500 in the money market account earning 8%. If he does not think he can earn more than \$4,200, he should not purchase the Seneca.

Interesting insights can be gleaned from Figure 5.3. For example, Carothers's uncertainty regarding Capacity of Scheduled Flights is extremely important. On the other hand, the annual profit is very insensitive to Aircraft Price. What can we do with information like this? The tornado diagram tells us which variables we need to consider more closely and which ones we can leave at their base values. In this case, annual profit is insensitive to Proportion Financed, Interest Rate, and Aircraft Price, so in further analyzing this decision we simply can leave these variables at their base values. And yet Capacity of Scheduled Flights, Operating Cost, Hours Flown, and Charter Price all have substantial effects on the annual profit; the bars for these four variables cross the critical \$4,200 line. Proportion of Chartered Flights, Ticket Price, and Insurance each have a substantial effect on the profit, but the bars for all of these variables lie entirely above the \$4,200 line. In a first pass, these variables might be left at their base values, and the analyst might perform another sensitivity analysis at a later stage.

Dominance Considerations

In our discussion of making decisions in Chapter 4, we learned that alternatives can be screened on the basis of deterministic and stochastic dominance, and inferior alternatives can be eliminated. Identifying dominant alternatives can be viewed as a version of sensitivity analysis for use early in an analysis. In sensitivity-analysis terms, analyzing alternatives for dominance amounts to asking whether there is any

way that one alternative could end up being better than a second. If not, then the first alternative is dominated by the second and can be ignored.

In the case of Eagle Airlines, an immediate question is whether purchasing the option is a dominated alternative. Why would Carothers want to buy the option? There are two possibilities. First, it would allow him to lock in a favorable price on a suitable aircraft while he tried to gather more information. Having constructed a tornado diagram for the problem, we can explore the potential value of purchasing the option by considering the amount of information that we might obtain and the potential impact of this information. A second typical motivation for purchasing an option is to wait and see whether the economic climate for the venture becomes more favorable. In this case, if the commuter/charter air-travel market deteriorates, then Carothers has only lost the cost of the option. (Some individuals also purchase options to lock in a price while they raise the required funds. Carothers, however, appears to have the required capital and credit.)

It is conceivable that Carothers could obtain more accurate estimates of certain input variables. Considering the tornado diagram, he would most like to obtain information about the more critical variables. Some information regarding market variables (Capacity of Scheduled Flights, Hours Flown, and Charter Ratio) might be obtainable through consumer-intentions surveys, but it would be far from perfect as well as costly. The best way to obtain such information would be to purchase or lease an aircraft for a year and try it — but then he might as well buy the Seneca!

What about Operating Cost and Insurance? The main source of uncertainty for Operating Cost is fuel cost, and this is tied to the price of oil, which can fluctuate dramatically. Increases in Insurance are tied to changes in risk as viewed by the insurance companies. Rates have risen dramatically over the years, and stability is not expected. The upshot of this discussion is that good information regarding many of the input variables probably is not available. As a result, if Carothers is interested in acquiring the option in order to have the chance to gather information, he might discover that he is unable to find what he needs.

What about the second motivation, waiting to see whether the climate improves? The question here is whether any uncertainty will be resolved during the term of the option, and whether or not the result would be favorable to Eagle Airlines. In general, considerable uncertainty regarding all of the market variables will remain regardless of how long Carothers waits. Market conditions can fluctuate, oil prices can jump around, and insurance rates can change. On the other hand, if some event is anticipated, such as settlement of a major lawsuit or the creation of new regulations, then the option could protect Carothers until this uncertainty is resolved. (Notice that, even in this case, the option provides Carothers with an opportunity to collect information — all he must do is wait until the uncertain situation is resolved.) But Carothers does not appear to be awaiting the resolution of some major uncertainty. Thus, if his motivation for purchasing the option is to wait to see whether the climate improves, it is not clear whether he would be less uncertain about the economic climate when the option expires.

What are the implications of this discussion? It is fairly clear that, unless an inexpensive information-gathering strategy presents itself, purchasing the option prob-

ably is a dominated alternative. For the purposes of the following analysis, we will assume that Carothers has concluded that no such information-gathering strategy exists, and that purchasing the option is unattractive. Thus, we can reduce his alternatives to (1) buying the airplane outright and (2) investing in the money market.

Two-Way Sensitivity Analysis


The tornado-diagram analysis provides considerable insights, although these are limited to what happens when only one variable changes at a time. Suppose we wanted to explore the impact of several variables at one time? This is a difficult problem, but a graphical technique is available for studying the interaction of two variables.

Suppose, for example, that we want to consider the joint impact of changes in the two most critical variables, Operating Cost and Capacity of Scheduled Flights. Imagine a rectangular space (Figure 5.4) that represents all of the possible values that these two variables could take. Now, let us find those values of Operating Cost and Capacity for which the annual profit would be less than \$4200. If this is to be the case, then we must have total revenues minus total costs less than \$4200 or total revenues less than total costs plus \$4200:

$$\begin{aligned}
 & (\text{Charter Proportion} \times \text{Hours Flown} \times \text{Charter Price}) + [(1 - \text{Charter Proportion}) \\
 & \quad \times \text{Hours Flown} \times \text{Ticket Price} \times \text{Number of Seats} \\
 & \quad \times \text{Capacity of Scheduled Flights}] < (\text{Hours Flown} \\
 & \quad \times \text{Operating Cost}) + \text{Insurance} \\
 & \quad + (\text{Price} \times \text{Percent Financed} \times \text{Interest Rate}) + 4200
 \end{aligned}$$

Inserting the base values for all but the two variables of interest, we obtain

Figure 5.4
Two-way sensitivity graph for Eagle Airlines. The Line AB represents the points for which profit would be \$4200.



$$(0.5 \times 800 \times 325) + [0.5 \times 800 \times 100 \times 5 \times \text{Capacity}] < (800 \times \text{Operating Cost}) + 20,000 + (87,500 \times 0.4 \times 0.115) + 4200$$

which reduces to

$$130,000 + (200,000 \times \text{Capacity}) < (800 \times \text{Operating Cost}) + 28,225$$

Now solve this inequality for Capacity in terms of Operating Costs to get Capacity

$$< 0.004 \times \text{Operating Cost} - 0.509$$

This inequality defines the region in which purchasing the airplane would lead to a profit of less than \$4200. When the "<" sign is replaced with an equality, we have the points for which profit equals \$4200, or the line of points where the venture just breaks even relative to investing in the money market. To plot this line, notice that we only need two points. The simplest way to come up with these is to plug in the extreme values for Operating Cost and calculate the corresponding values for Capacity. Doing this gives the break-even points A (Capacity = 0.411 when Operating Cost = 230) and B (Capacity = 0.531 when Operating Cost = 260). These points define Line AB in Figure 5.4. The area below the line (Capacity < 0.004 × Operating Cost — 0.509) represents the region when the profit would be less than \$4200. The area above the line represents the region in which the profit would be greater than \$4200.

What insight can Carothers gain from Figure 5.4? The point labeled "Base Values" shows that when we plug in the base values for the capacity and operating-cost variables, we get an estimated profit that is greater than \$4200, and so the project looks promising. But Carothers might be wondering how likely it is that the two variables might work together to lead to a profit of less than \$4200. For example, suppose that Operating Cost was slightly more than the base value, say \$248, and that Capacity was just slightly less than the base value, say 48%. Taken individually, these two values do not seem to cause a problem. That is, substituting either one into the profit calculations, while keeping the other at its base value, leads to profit that is still greater than \$4200. When we consider these two values jointly (Point C in Figure 5.4), however, they lead to a situation in which it would have been better not to buy the airplane. If Carothers thinks that the two variables might be likely to fall in the "Profit < \$4200" region, then he may wish to forego the purchase. But such a situation would indicate that he really needs to model his uncertainty about these variables using probability methods. In the next section we will see how two-way sensitivity analysis can be used in conjunction with probabilities.

Sensitivity to Probabilities

The next step in our analysis will be to model the uncertainty surrounding the critical variables identified by our analysis of the tornado diagram. The four critical variables were (1) Capacity of Scheduled Flights, (2) Operating Cost, (3) Hours

Flown, and (4) Charter Price. We only need to think about uncertainty for the first three, because charter price is a decision variable set by Carothers. For the purposes of the example here, let us assume that, in an initial attempt to model the uncertainty, Carothers chooses two values for each variable, one representing an optimistic and one a pessimistic scenario. The influence diagram is shown in Figure 5.5 and shows changes in the model based on the sensitivity analysis so far. Operating Cost, Hours Flown, and Capacity of Scheduled Flights have been changed to chance nodes. The remaining input variables (Interest Rate, Proportion Financed, Price, Insurance, Charter Price, Ticket Price, and Proportion of Chartered Flights) have been left at their base values and hence continue to be represented by rounded rectangles (constants). The decision tree in Figure 5.6 shows the pessimistic and optimistic values for the three uncertain variables.

Now that we have simplified the problem somewhat, we can include considerations regarding the interdependence of the remaining chance variables. In Figures 5.5 and 5.6, the probability distribution for Hours Flown is judged to depend on the Capacity of Scheduled Flights: If Capacity is low, then this may actually result in some flights being canceled and thus fewer total hours. Thus, a relevance arc leads from "Capacity of Scheduled Flights" to "Hours Flown" in the influence diagram, and in the decision tree the value for $r = P(\text{Low Hours} \mid \text{Low Capacity})$ may not be the same as the value for $s = P(\text{Low Hours} \mid \text{High Capacity})$. In fact, our argument suggests that r will be greater than s . On the other hand, Operating Cost is judged to be independent of the other variables.

The next thing to do is to assess some values for probabilities p , q , r , and s . Let us suppose that Carothers is comfortable with an assessment that $p = 0.5$, or that

Figure 5.5
Influence diagram of Eagle Airlines decision. Note that only three variables are considered to be uncertain, and that Hours Flown and Capacity are considered to be probabilistically dependent.

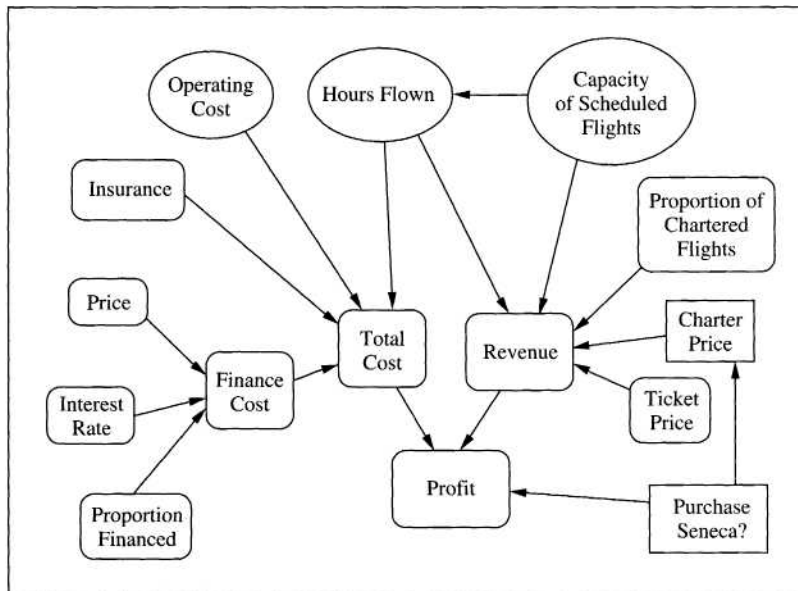
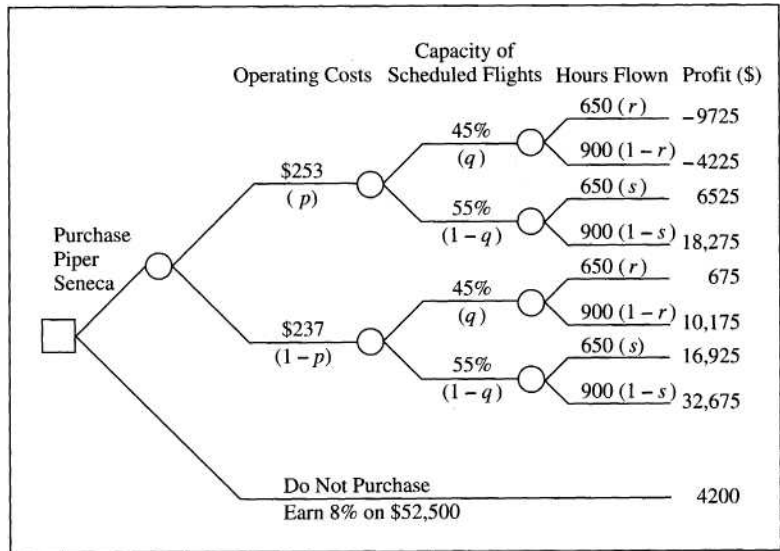
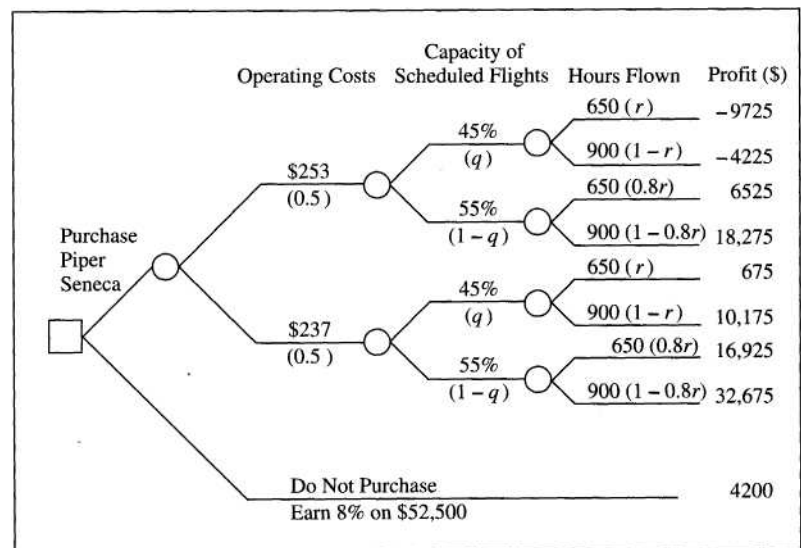


Figure 5.6
 Decision tree for Eagle Airlines with uncertainty for three variables. Profit is calculated with all other variables held at their base values.



Operating Cost is just as likely to be high (\$253) as low (\$237). Furthermore, suppose that Carothers feels that a reasonable way to represent the dependence between Hours and Capacity is to let s be 80% of r . That is, if Capacity is high (55%), then the probability that Hours = 650 is only 80% of the probability that Hours = 650 when Capacity is low. With these two specifications, we now have only two unspecified probabilities left to consider, q and r . Figure 5.7 shows the modified decision tree with $p = 0.5$ and $s = 0.8r$.

Figure 5.7
 Eagle Airlines' decision tree with probabilities substituted for p and s . This decision tree is now ready for a two-way sensitivity analysis on q and r .



We now can create a two-way sensitivity graph for q and r . As with the two-way sensitivity analysis above, the graph will show regions for which the expected value of purchasing the Seneca is greater than investing in the money market.

To create the graph, we first write out the expected value of purchasing the airplane in terms of q and r , including the specifications that $p = 0.5$ and $s = 0.8r$. This equation comes from solving the decision tree:

$$\begin{aligned} \text{EMV}(\text{Purchase}) = & 0.5 \{q[-9725r - 4225(1 - r)] + \\ & (1 - q) [6525(0.8r) + 18,275(1 - 0.8r)]\} + \\ & 0.5 \{q[675r + 10,175(1 - r)] + (1 - q)[16,925(0.8r) \\ & + 32,675(1 - 0.8r)]\} \end{aligned}$$

After algebraic reduction, this expression becomes

$$\text{EMV}(\text{Purchase}) = q(3500r - 22,500) - 11,000r + 25,475$$

We would want to purchase the airplane if $\text{EMV}(\text{Purchase}) > 4200$. Thus, we can solve the following inequality for q in terms of r .

$$q(3500r - 22,500) - 11,000r + 25,475 > 4200$$

$$25,475 - 4200 - 11,000r > q(22,500 - 3500r)$$


This inequality reduces to

$$\frac{21,275 - 11,000r}{22,500 - 3500r} > q$$

Using this inequality, we can create a two-way sensitivity graph for Eagle Airlines (Figure 5.8). The curve separating the two regions represents the values of q and r for which $\text{EMV}(\text{Purchase}) = \4200 . It was plotted by plugging values for r between 0 and 1 into the inequality above. For these values of q and r , Carothers should be indifferent (in terms of EMV) between buying the airplane and not. The area below the line contains points where $q < (21,275 - 11,000r)/(22,500 - 3500r)$; for these (q, r) points, $\text{EMV}(\text{Purchase}) > \4200 . The graph makes sense because q and r are probabilities of the pessimistic scenarios — low Capacity and low number of Hours Flown. If Carothers thinks that the pessimistic scenarios are likely (q and r close to 1), then he would not want to buy the airplane.

The importance of Figure 5.8 is that Carothers may not have especially firm ideas of what the probabilities q and r should be. Suppose, for example, that in the process of coming up with the probabilities he feels that q could be between 0.4 and 0.5 and that r could be between 0.5 and 0.65. These probabilities are represented by the points inside Rectangle A in Figure 5.8. All of these points fall within the "Purchase Seneca" region, and so the conclusion is that the Seneca should be purchased. The decision is not sensitive to the assessment of the probabilities. If, on the other hand, Carothers thinks that reasonable values of q and r fall in Rectangle B, then the optimal choice is not clear. (No wonder the decision is a hard one!) In this situation, he could reflect on the chances associated with Capacity and Hours Flown and try to refine his model of the uncertainty. Decision-analysis tools for modeling uncertainty more carefully are discussed in Chapters 7 through 12.

Figure 5.8
Two-way sensitivity
graph for Eagle
Airlines.



The value of the two-way sensitivity graph is to provide guidance in determining how much effort is needed to model uncertainty in a decision problem. Looking at it another way, the graph can reveal whether the decision is sensitive to the uncertainty in the problem and to the modeling of that uncertainty.

Two-Way Sensitivity Analysis for Three Alternatives (Optional)

As we have analyzed it here, the Eagle Airlines case involves only two alternatives, and so the sensitivity graph has two regions. What happens when there are more than two alternatives? The graph may contain a region for each alternative. Let us consider a stock market-investment problem.

INVESTING IN THE STOCK MARKET

An investor has funds available to invest in one of three choices: a high-risk stock, a low-risk stock, or a savings account that pays a sure \$500. If he invests in the stocks, he must pay a brokerage fee of \$200.

His payoff for the two stocks depends in part on what happens to the market as a whole. If the market goes up (as measured, say, by the Standard and Poor's 500 Index increasing 8% over the next 12 months), he can expect to earn \$1700 from the high-risk stock and \$1200 from the low-risk stock. Finally, if the stock market goes down (as indicated by the index decreasing by 3% or more), he will lose \$800 with the high-risk stock but still gain \$100 with the low-risk stock. If the market stays at roughly the same level, his payoffs for the high- and low-risk stocks will be \$300 and \$400, respectively.

The decision tree is given in Figure 5.9, with unspecified probabilities $t = P(\text{market up})$ and $v = P(\text{market same})$. Of course, $P(\text{market down}) = 1 - t - v$ because the probabilities must sum to 1.

To construct the graph, we must compare the alternatives two at a time. First we have to realize that $t + v$ must be less than or equal to 1. Thus, the graph (see Figure 5.10) is a triangle rather than a rectangle because all of the points above a line from $(t = 1, v = 0)$ to $(t = 0, v = 1)$ are not feasible. To find the strategy regions, begin by finding the area where the savings account would be preferred to the low-risk stock, or

$$\begin{aligned} \text{EMV}(\text{Savings Account}) &> \text{EMV}(\text{Low-Risk Stock}) \\ 500 &> t(1000) + v(200) - (1 - t - v)100 \end{aligned}$$

Solving for v in terms of t , we get

$$v < 2 - \frac{11t}{2}$$

Figure 5.9
Decision tree for a stock market investor.

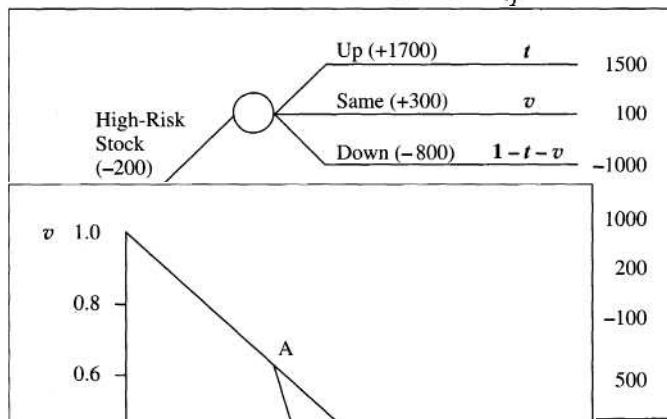


Figure 5.10
Beginning the analysis of the stock market problem. Note that $t + v$ must be less than or equal to 1, and so the only feasible points are within the large triangular region.




Figure 5.10 shows the regions for the savings account and the low-risk stock divided by Line AB.

Now let us find the regions for the high- and low-risk stocks. Begin by setting up the inequality

$$\begin{aligned} \text{EMV}(\text{Low-Risk Stock}) > \text{EMV}(\text{High-Risk Stock}) & t(1000) + v \\ (200) - (1 - t - v)100 > t(1500) + v(100) - (1 - t - v)1000 \end{aligned}$$

This reduces to

$$v < \frac{9}{8} - \frac{7t}{4}$$

Using this inequality, we can add another line to our graph (Figure 5.11). Now Line CDE separates the graph into regions in which EMV(Low-Risk Stock) is greater or less than EMV(High-Risk Stock).

From Figure 5.11 we can tell what the optimal strategy is in all but one portion of the graph. For example, in ADEG, we know that the high-risk stock is preferred to the low-risk stock and that the low-risk stock is preferred to the savings account. Thus, the high-risk stock would be preferred overall. Likewise, in HFBDC the savings account would be preferred, and in DBE the low-risk stock would be preferred. But in CDA, all we know is that the low-risk stock is worse than the other two, but we do not know whether to choose the savings account or the high-risk stock.

If the decision maker is sure that the probabilities t and v do not fall into the region CDA, then the sensitivity analysis could stop here. If some question remains (or even if we feel compelled to finish the job), then we can complete the graph by comparing EMV(Savings Account) with EMV(High-Risk Stock):

Figure 5.11
 Second stage in analysis of the stock market problem. A second inequality has been incorporated. The optimal strategy is clear now for all regions except CDA.





Figure 5.12
 Completed two-way
 sensitivity graph for
 the stock market
 problem. Line ID has
 split region CDA.



$$EMV(\text{Savings Account}) > EMV(\text{High-Risk Stock})$$

$$500 > t(1500) + v(100) - (1 - t - v)1000$$

This inequality reduces to

$$v < \frac{15}{11} - \frac{25t}{11}$$

Incorporating this result into the graph allows us to see that region CDA actually is split between the high-risk stock and the savings account as indicated by Line ID in Figure 5.12.

With the analysis completed, the investor now can think about probabilities t and v . As in the Eagle Airlines case, it should be possible to tell whether the optimal investment decision is sensitive to these probabilities and whether additional effort should be spent modeling the uncertainty about the stock market.

Sensitivity Analysis in Action

Is sensitivity analysis ever used in the real world? Indeed it is. This fundamental approach to modeling is the source of important insights and understanding in many real-world problems. The following example comes from medical decision making, showing how sensitivity-analysis graphs can improve decisions in an area where hard decisions are made even harder by the stakes involved.

HEART DISEASE IN INFANTS


Macartney, Douglas, and Spiegelhalter used decision analysis to study alternative treatments of infants who suffered from a disease known as coarctation of the aorta. Difficult to detect, the fundamental uncertainty is whether the disease is present at all. Three alternative treatments exist if an infant is suspected of having the disease. The first is to do nothing, which may be appropriate if the disease is not present. The second is to operate. The third alternative is to catheterize the heart in an attempt to confirm the diagnosis, although it does not always yield a perfect diagnosis. Moreover, catheterizing the heart of a sick infant is itself a dangerous undertaking and may lead to death. The difficulty of the problem is obvious; with all of the uncertainty and the risk of death from operating or catheterization, what is the appropriate treatment?

Source: F. Macartney, J. Douglas, and D. Spiegelhalter (1984) "To Catheterise or Not to Catheterise?" *British Heart Journal*, 51, 330-338.

In their analysis Macartney et al. created a two-way sensitivity graph (Figure 5.13) showing the sensitivity of the decision to two probabilities. The two probabilities are (1) the disease is present, which is along the horizontal axis, and (2) the mortality rate for cardiac catheterization, which is along the vertical axis. The mortality rate also could be interpreted as the physician's judgment regarding the chance that the infant would die as a result of catheterization.

The graph shows three regions, reflecting the three available alternatives. The location of the three regions makes good sense. If the physician believes that the

Figure 5.13
Two-way sensitivity analysis for the heart disease treatment decision.
Source: Macartney et al. (1984).



chances are low that the disease is present and that the risk of catheterizing the infant is high, then the appropriate response is to do nothing. On the other hand, if the risk of catheterization is high relative to the chance that the disease is present, then operating without catheterizing is the prescribed treatment. Catheterization is recommended only for situations with relatively low risk from the procedure.

Sensitivity Analysis by Computer

Who wants to go through all that algebra to perform a sensitivity analysis? Fortunately, sensitivity analysis can be done easily by computer. The decision-analysis software packages described in Chapters 3 and 4 all have built-in sensitivity-analysis routines. Although every program has its own unique approach to sensitivity analysis, many of the tools are consistent with the discussion in this chapter. In this section, we will look briefly at the sensitivity tools provided by DPL and DATA.

Sensitivity Analysis with DPL

DPL can produce one-way sensitivity graphs with the command "Value Sensitivity Analysis." For example, Figure 5.14 shows Eagle's profit for different values of Hours Flown, and the graph is partitioned into regions that represent different optimal strategies. (On DPL's screen, the regions have different colors, and hence DPL calls this kind of graph a *rainbow diagram*.)

Creating tornado diagrams is easy with DPL. On the "Run" menu, the command "Value Sensitivity Comparison" allows the user to build a tornado diagram one van-

Figure 5.14
Rainbow diagram
from DPL The vertical
axis, labeled "Expected
Value," represents
Eagle's profit.





Figure 5.15
Eagle Airline tornado diagram from DPL.



able at a time, and it is the user's choice as to which variables are included in the diagram. Figure 5.15 shows DPL's version of the tornado diagram for Carothers's decision. As you can see, DPL has the capability of automatically including the extreme values of the input variables and the corresponding value for profit at the ends of the bars. In addition, DPL has automatically truncated the four bars that cross over into the region where profit falls below \$4200. (In these cases, DPL has solved the complete decision problem and found that the optimal choice is to invest in the money market for a return of \$4200. Creating a tornado diagram whose bars are not truncated as in Figure 5.3 is straightforward; delete the "Purchase Seneca" decision node in the influence diagram, and then create the tornado diagram, specifying Profit as the criterion variable.)

The "Value Sensitivity Comparison" is designed specifically to help a decision maker decide which deterministic variables to model probabilistically. Thus, as long as you are moving from a deterministic toward a probabilistic model, this is an appropriate tool. What if you have created a complex decision model with many chance nodes? In this case, you might like to go the other way, deciding which of the

chance variables could be changed to deterministic ones with minimal impact to the model. DPL includes a command, "Event Sensitivity Comparison," that allows the decision maker to examine chance variables, in the context of a tornado diagram, for just this purpose.

Sensitivity Analysis with DATA

In DATA's "Analysis" menu, the user can find a number of sensitivity-analysis commands. Figure 5.16 again shows the one-way sensitivity for Hours Flown in the Eagle Airlines decision. DATA even calculates the threshold value of Hours Flown (664.118) where the two lines intersect and gives Profit (\$4200) at that point. DATA also can construct a tornado diagram (Figure 5.17).

DATA can also perform two- and three-way sensitivity analysis. For example, Figure 5.18 is the two-way sensitivity graph that we created in the probabilistic analysis of Eagle Airlines (see Figure 5.8). The graph identifies the appropriate strategy for each region by using shading. To do a three-way sensitivity analysis, DATA can display an animated sequence of two-way graphs.

This overview of sensitivity analysis by computer serves only as an introduction. This book cannot show all of the available ways to conduct such analysis, but you should now have a better appreciation for the available tools and techniques. With practice using these tools, you should be able to avoid a lot of dreary algebra.

Figure 5.16
One-way sensitivity
analysis on hours
flown from DATA.




Figure 5.17
Tornado diagram from
DATA.





Figure 5.18
DATA's two-way
sensitivity graph for
probabilistic analysis
of Eagle Airlines.



Sensitivity Analysis: A Built-in Irony

There is a strange irony in sensitivity analysis and decision making. We begin by structuring a decision problem, part of which involves identifying several alternatives. Then some alternatives are eliminated on the grounds of dominance. The remaining ones are difficult to choose from. Being difficult to choose from, they lead us to unveil our array of decision-analysis tools. But also being difficult to choose from, they probably are not too different in expected value; and if so, then it does not matter much which alternative one chooses, does it? For the analyst who wants to be quite sure of making the best possible choice, this realization can be terribly frustrating; almost by definition, hard decisions are sensitive to our assessments. For those who are interested in modeling to improve decision making, the thought is comforting; the better the model, the better the decision, but only a small degree of improvement may be available from rigorous and exquisite modeling of each minute detail. Adequate modeling is all that is necessary. The best way to view sensitivity analysis is as a source of guidance in modeling a decision problem. It provides the guidance for each successive iteration through the decision-analysis cycle. You can see now how the cycle is composed of modeling steps, followed by sensitivity analysis, followed by more modeling, and so on. The ultimate objective of this cycle of modeling and analysis is to arrive eventually at a requisite decision model and to analyze it just enough to understand clearly which alternative should be chosen. By the time the decision maker reaches this point, all important issues will be included in the decision model, and the choice should be clear.

SUMMARY

This chapter has presented an approach and several tools for performing sensitivity analysis. We have considered sensitivity analysis in terms of identifying and structuring problems, dominance among alternatives, and probability assessment. Tornado diagrams and one- and two-way sensitivity graphs were developed, and we discussed ways to perform sensitivity analysis using computers. The purpose of sensitivity analysis in the decision-analysis cycle is to provide guidance for the development of a requisite decision model.

EXERCISES

- 5.1 What is the fundamental question that sensitivity analysis answers?
- 5.2 Some friends of yours have been considering purchasing a new home. They currently live 20 miles from town on a two-acre tract. The family consists of the mother, father, and two small children. The parents also are considering having more children, and they realize that as the children grow, they may become more involved in activities in town. As it is, most of the family's outings take place in town. Describe the role that sensitivity analysis could play in your friends' decision. What variables could be subjected to sensitivity analysis?


- 5.3 Over dinner, your father mentions that he is considering retiring from real-estate sales. He has found a small retail business for sale, which he is considering acquiring and running. There are so many issues to think about, however, that he has a difficult time keeping them all straight. After hearing about your decision-analysis course, he asks you whether you have learned anything that might help him in his decision. What kinds of issues are important in deciding whether to buy a retail business? Describe how he might use sensitivity analysis to explore the importance of these issues.
- 5.4 When purchasing a home, one occasionally hears about the possibility of "renting with an option to buy." This arrangement can take various forms, but a common one is that the renter simply pays rent and may purchase the house at an agreed-upon price. Rental payments typically are not applied toward purchase. The owner is not permitted to sell the house to another buyer unless the renter/option holder waives the right to purchase. The duration of the option may or may not be specified.

Suppose that a buyer is considering whether to purchase a house outright or rent it with an option to buy. Under what circumstances would renting with an option be a dominated alternative? Under what circumstances would it definitely not be dominated?
- 5.5 What role does sensitivity analysis play in the development of a requisite decision model?
- 5.6 Explain why the lines separating the three regions in Figure 5.12 all intersect at Point D.

QUESTIONS AND PROBLEMS

5.7 *Cost-to-loss ratio problem.* Consider the decision problem shown in Figure 5.19. This basic decision tree often is called a cost-to-loss ratio problem and is characterized as a decision situation in which the question is whether to take some protective action in the face of possible adverse circumstances. For example, the umbrella problem (Figure 4.9) is a cost-to-loss ratio problem. Taking the umbrella incurs a fixed cost and protects against possible adverse weather. A farmer may face a cost-to-loss ratio problem if there is a threat of freezing weather that could damage a fruit crop. Steps can be taken to protect the orchard, but they are costly. If no steps are taken, the air temperature may or may not become cold enough to damage the crop.

Figure 5.19
Cost-to-loss ratio problem.

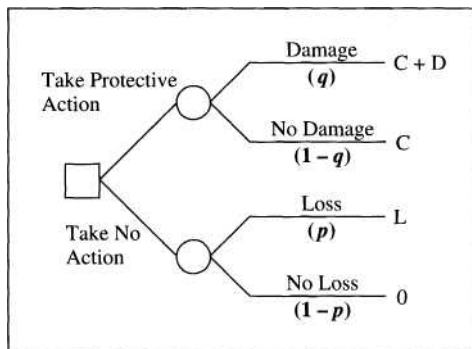


Sensitivity analysis is easily performed for the cost-to-loss ratio problem. How large can the probability p become before "Take Protective Action" becomes the optimal (minimum expected cost) alternative? Given your answer, what kind of information does the decision maker need in order to make the decision? (*Hint*: This is an algebra problem. If that makes you uncomfortable, substitute numerical values for C , L , and p .)

- 5.8** *The cost-to-loss ratio problem continued.* The cost-to-loss ratio problem as shown in Figure 5.19 may be considered a simplified version of the actual situation. The protective action that may be taken may not provide perfect protection. Suppose that, even with protective action, damage D will be sustained with probability q . Thus, the decision tree appears as Figure 5.20. Explain how sensitivity analysis could be used to determine whether it is important to include the upper chance node with probability q and damage D .

Figure 5.20

More general version of the cost-to-loss problem.



- 5.9** An orange grower in Florida faces a dilemma. The weather forecast is for cold weather, and there is a 50% chance that the temperature tonight will be cold enough to freeze and destroy his entire crop, which is worth some \$50,000. He can take two possible actions to try to alleviate his loss if the temperature drops. First, he could set burners in the orchard; this would cost \$5000, but he could still expect to incur damage of approximately \$15,000 to \$20,000. Second, he could set up sprinklers to spray the trees. If the temperature drops, the water would freeze on the fruit and provide some insulation. This method is cheaper (\$2000), but less effective. With the sprinklers he could expect to incur as much as \$25,000 to \$30,000 of the loss with no protective action.

Compare the grower's expected values for the three alternatives he has, considering the various possible loss scenarios for the burners and the sprinklers. Which alternative would you suggest the grower take? Why?

- 5.10** An important application of sensitivity analysis occurs in problems involving multiple attributes. Many decision makers experience difficulty in assessing trade-off weights. A sensitivity analysis of the trade-off weight, though, can reveal whether a decision maker must make a more precise judgment. Reconsider the summer-job example described and analyzed in Chapter 4 (pages 129-136). In the analysis, we used trade-off weights of $k_s = 0.60$ for salary and $k_f = 0.40$ for fun (see Figure 4.36).

Suppose Sam Chu is uncomfortable with the precise assessment that $k_s = 0.60$. Sam does believe, though, that k_s could range from 0.50 up to 0.75. (Recall that k_s and k_f add up to 1, so by implication, k_f can range from 0.50 to 0.25, depending on the value of

- k_s .) Perform a sensitivity analysis on the expected overall score for the two jobs by varying k_s over this range. Is the forest job preferred for all values of k_s between 0.50 and 0.75?
- 5.11** A friend of yours can invest in a multiyear project. The cost is \$14,000. Annual cash flows are estimated to be \$5000 per year for six years but could vary between \$2500 and \$7000. Your friend estimates that the cost of capital (interest rate) is 11%, but it could be as low as 9.5% and as high as 12%. The basis of the decision to invest will be whether the project has a positive net present value. Construct a tornado diagram for this problem. On the basis of the tornado diagram, advise your friend regarding either (1) whether to invest or (2) what to do next in the analysis.
- 5.12** Reconsider Hugh Liedtke's decision as diagrammed in Figure 4.2. Note that three strategies are possible: (1) accept \$2 billion, (2) counteroffer \$5 billion and then accept \$3 billion if Texaco counteroffers, and (3) counteroffer \$5 billion and then refuse \$3 billion if Texaco counteroffers. Suppose that Liedtke is unsure about the probabilities associated with the final court outcome. Let $p = P(10.3)$ and $q = P(5)$ so that $1 - p - q = P(0)$. Create a two-way sensitivity graph that shows optimal strategies for Liedtke for possible values of p and q . (*Hint*: What is the constraint on $p + q$?) If Liedtke thinks that p must be at least 0.15 and q must be more than 0.35, can he make a decision without further probability assessment?

CASE STUDIES

THE STARS AND STRIPES

In 1987, the United States won the prestigious America's Cup sailing race, winning the trophy from Australia. The race normally is run every four years, but in 1988 New Zealand invoked an obscure provision in the race charter and challenged the U.S. team to a match. Furthermore, New Zealand proposed to race with the largest boat permitted rather than a standard 12-meter craft. The new yacht, dubbed the *New Zealand*, was 133 feet long, designed and built using space-age material, and equipped with state-of-the-art computer equipment to monitor performance.

Not to be outdone, the U.S. team countered by designing and building a catamaran, the *Stars and Stripes*. A conventional sailboat like the *New Zealand* drags a heavy keel through the water to maintain stability. A catamaran relies on two long, narrow hulls for stability and thus can be considerably lighter and faster. Furthermore, the *Stars and Stripes* was outfitted with a rigid sail designed like an airplane wing. With slots and flaps controlled by wires, the sail could be adjusted precisely for optimum airflow.

New Zealand counterattacked with a lawsuit claiming that the vague deed that established the competition implied that the match was to be between similar boats. But the New York Supreme Court ruled that the race should go on and that protests should be filed afterward. The race began on September 7, 1988. The *Stars and Stripes* won easily, by 18 minutes in the first race and 21 minutes in the second.

Questions

- 1 Designing world-class racing sailboats involves thousands of decisions about shape, size, materials, and countless other details. What are some objectives that might be reasonable in designing such a sailboat?
- 2 What are some specific design decisions that must be made (for example, the shape of the sail)?
- 3 How can sensitivity analysis be used to decide which design decisions are more important than others?

DUMOND INTERNATIONAL, PART I

"So that's the simplified version of the decision tree based on what appear to be the critical issues," Nancy Milnor concluded. "Calculating expected values, it looks as though we should introduce the new product. Now, I know that we don't all agree on the numbers in the tree, so why don't we play around with them a little bit. I've got the data in the computer here. I can make any changes you want and see what effect they have on the expected value."

Nancy had just completed her presentation to the board of directors of DuMond International, which manufactured agricultural fertilizers and pesticides. The decision the board faced was whether to go ahead with a new pesticide product to replace an old one or whether to continue to rely on the current product, which had been around for years and was a good seller. The problem with the current product was that evidence was beginning to surface which showed that the chemical's use could create substantial health risks, and there even was some talk of banning the product. The new product still required more development, and the question was whether all of the development issues could be resolved in time to meet the scheduled introduction date. And once the product was introduced, there was always the question of how well it would be received. The decision tree (Figure 5.21) that Nancy had presented to the board captured these concerns.

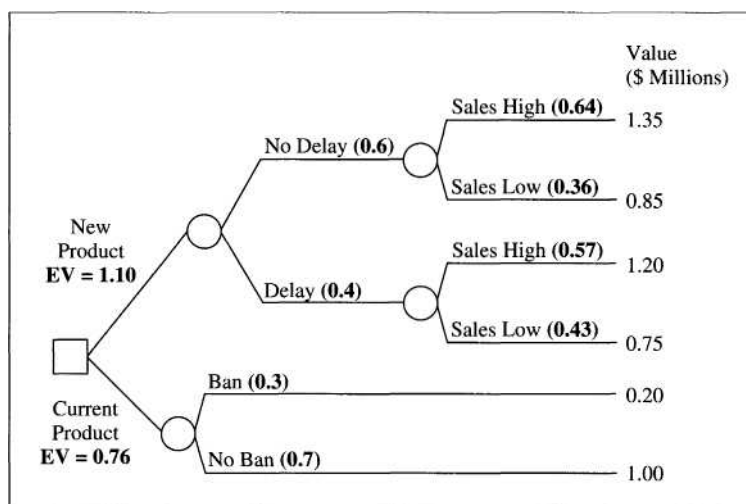
The boardroom was beginning to get warm. Nancy sat back and relaxed as she listened to the comments.

"Well, I'll start," said John Dilts. "I don't have much trouble with the numbers in the top half of the tree. But you have the chance of banning the current product pinned at 30 percent. That's high. Personally, I don't think there's more than a 10 percent chance of an out-and-out ban."

"Yeah, and even if there were, the current product ought to be worth \$300,000 at least," added Pete Lillovich. "With a smaller chance of a ban and a higher value, surely we're better off with the old product!"

"Well, I don't know about you two," said Maria Jenkins. "I think we have a pretty good handle on what's going on with the current product. But I'd like to play the new product a little more conservatively. I know that the values at the ends of the branches on the top half of the tree are accounting's best guesses based on a complete analysis,

Figure 5.21
DuMond's new
product decision.



but maybe they should all be reduced by \$100,000 just to play it safe. And maybe we should just set the probability of high sales equal to 50 percent regardless of the delay."

Steven Kellogg had been involved in the preliminary development of the new product more than anyone else. He piped up, "And the delay is actually more likely than no delay. I'd just reverse those probabilities so that there's a 60 percent chance of a delay. But I wouldn't make any changes on the lower part of the tree. I agree with Maria that we have a good idea about the performance of the current product and the prospects for a ban."

"I don't think it matters," countered Lillovich. "The changes John and I suggest make the current product look better than it does in Nancy's analysis. Maria's and Steven's changes make the new product look worse. Either way, the effect is the same."

Nancy had kept track of the comments and suggested changes. She sat down at the computer and started to enter the changes. After a few moments, she grinned and turned to the board. "In spite of your changes," she said, "I believe I can persuade you that DuMond should go with the new product."

Question

- 1 Explain why Nancy believes that DuMond should go with the new product.

STRENLAR, PART II

Question

- 1 The Strenlar case study at the end of Chapter 4 (page 148) required substantial modeling. Use sensitivity analysis to refine your model. In particular, you might consider (1) the interest rate used to calculate net present value, (2) legal fees, (3)

the eventual price of PI's stock, (4) Strenlar's gross sales, (5) Fred's profits if Strenlar is successful (remembering that profits are linked to sales), (6) the probability of Strenlar being successful, and (7) the probability of winning the lawsuit. Do you think that Fred's decision is sensitive to any of these variables? Try wiggling one variable at a time away from its base value (the value given in the case) while holding everything else at base value. How much can you wiggle the variable before the decision changes? At the end of your analysis, discuss your results and the implications for Fred's decision model. If he were to refine his model, what refinements should he make?

FACILITIES INVESTMENT AND EXPANSION

Spetzler and Zamora describe a decision analysis concerning whether a major U.S. corporation should make a \$20 million investment in a new plant with the possibility of a \$5 million expansion later. The product was a "brightener," and there was some chance that the process also could generate significant quantities of a valuable by-product. Unfortunately, the primary chemical reaction was difficult to control. The exact yields were uncertain and were subject to the amounts of impurities in the raw material. Other substantial uncertainties surrounded the decision, such as raw material costs, inflation effects, federal regulatory intervention, the development of a full-scale production process based on the pilot project, and so on.

At an early stage of the analysis, the focus was on the numerous uncertain variables that could affect the value of the investment. Table 5.2 shows the effect of 18 different variables on the project's present value.

Question

- 1 Which variables definitely should be kept in the model? What additional information is needed in order to decide whether any of the other variables can be eliminated?

Source: C. S. Spetzler and R. M. Zamora (1984) "Decision Analysis of a Facilities Investment and Expansion Decision." In R. A. Howard and J. E. Matheson (eds.) *The Principles and Applications of Decision Analysis*. Menlo Park, CA: Strategic Decisions Group.

JOB OFFERS, PART II

Questions

- 1 Reconsider Robin Pinelli's dilemma in choosing from among three job offers (page 150). Suppose that Robin is unwilling to give a precise probability for disposable income with the Madison Publishing job. Conduct a sensitivity analysis on the expected value of the Madison Publishing job assuming that the probability of dispos-

Variable	Base Case	Tested Range		PV Range (\$ Million)		Change In PV (\$ Millions)
		From	To	From	To	
1 By-Product Production	36.00	0.00	80.00	-30	35	65
	lb/ton	lb/ton	lb/ton			
2 Market Price of Brightener in 1980	\$0.27/lb	\$0.15/lb	\$0.35/lb	-12	45	57
3 Raw Material Cost Growth	5.00%/yr	0.00%/yr	8.00%/yr	-7	40	47
4 Raw Material Costs	\$7.00/ton	\$2.00/ton	\$18.00/ton	-9	35	44
5 Impurities in Raw Material	4.00	2.00	6.00	-10	30	40
	lb/ton	lb/ton	lb/ton			
6 Cost Multiplier on Investment	90.00%	70.00%	125.00%	-12	25	37
7 Brightener Price Growth after 1980	4.00%/yr	2.00%/yr	6.00%/yr	⁻⁵	28	33
8 Cost Multiplier on Operations	110.00%	80.00%	150.00%	3	32	29
Expenses						
9 Cost Multiplier on Maintenance Expenses	100.00%	70.00%	120.00%	2	30	28
10 By-Product Price Growth	\$0.03/yr	\$0.01/yr	\$0.06/yr	3	30	27
11 Water Reclamation Costs	\$0.03/gal	\$0.02/gal	\$0.04/gal	4	31	27
12 By-Product Price in 1970	\$0.50/lb	\$0.40/lb	\$0.60/lb	3	29	26
13 Plant Efficiency	75.00%	50.00%	110.00%	0	26	26
14 Brightener Production	45.00	40.00	50.00	4	30	26
	lb/ton	lb/ton	lb/ton			
15 Market Price of Brightener in 1974	\$0.25/lb	\$0.15/lb	\$0.30/lb	4	28	24
16 Government Regulation Costs	\$0.02/lb	\$0.01/lb	\$0.03/lb	3	25	22
17 Market Price of Other By-Products	\$0.10/lb	\$0.07/lb	\$0.12/lb	4	24	20
18 Other By-Products Produced	84.00	70.00	96.00	4	24	20
	lb/ton	lb/ton	lb/ton			

Source: Spetzler and Zamora (1984).

Table 5.2

Results of deterministic sensitivity analysis on 18 different variables in facilities investment decision.

able income being \$1500 could range anywhere from zero to 1. Does the optimal choice — the job with the highest expected overall score — depend on the value of this probability? [Hint: Remember that probabilities must add up to 1, so $P(\text{Disposable Income} = \$1300)$ must equal $1 - P(\text{Disposable Income} = \$1500)$.] Suppose Robin is unable to come up with an appropriate set of trade-off weights. Assuming that any combination of weights is possible as long as all three are positive and add up to 1, conduct a sensitivity analysis on the weights. Create a graph that shows the regions for which the different job offers are optimal (have the high-

est expected overall score). (*Hint: Given that the three weights must sum to 1, this problem reduces to a two-way sensitivity analysis like the stock market example on pages 169-172.*)

REFERENCES

Sensitivity analysis is one of the more recent additions to the decision analyst's bag of tricks. As more complicated decision problems have been tackled, it has become obvious that sensitivity analysis plays a central role in guiding the analysis and interpreting the results. Recent overviews of sensitivity analysis can be found in Samson (1988), von Winterfeldt and Edwards (1986), and Watson and Buede (1987). In particular, Watson and Buede use real-world examples to show how sensitivity analysis is a central part of the decision-analysis modeling strategy. Phillips (1982) describes an application in which sensitivity analysis played a central part in obtaining consensus among the members of a board of directors.

Howard (1988) presents tornado diagrams and gives them their name. This approach to deterministic sensitivity analysis, along with other sensitivity-analysis tools, is discussed by McNamee and Celona (1987). Another sensitivity tool worth mentioning is the spiderplot (Eschenbach, 1992). The spiderplot is similar to the tornado diagram in that it allows simultaneous comparison of the impact of several different variables on the consequence or expected value.

Eschenbach, T. G. (1992) "Spiderplots versus Tornado Diagrams for Sensitivity Analysis." *Interfaces*, 22(6), 40-46.

Howard, R. A. (1988) "Decision Analysis: Practice and Promise." *Management Science*, 34, 679-695.

McNamee, P., and J. Celona (1987) *Decision Analysis for the Professional with Supertree*. Redwood City, CA: Scientific Press.

Phillips, L. D. (1982) "Requisite Decision Modelling." *Journal of the Operational Research Society*, 33, 303-312.

Samson, D. (1988) *Managerial Decision Analysis*. Homewood, IL: Irwin.

von Winterfeldt, D., and W. Edwards (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press.

Watson, S., and D. Buede (1987) *Decision Synthesis*. Cambridge: Cambridge University Press.

EPILOGUE

After the America's Cup race, New Zealand filed its lawsuit against the U.S. team, claiming that the deed of gift that established the race called for a "fair match." The New York Yacht Club, which had held the trophy from 1851 until 1983, even filed an affidavit with the court supporting New Zealand's contention. In April 1989, the court awarded the trophy to New Zealand. But in September 1989, the New York State Supreme Court issued a 4-1 decision that the cup should go back to the United States. (*Sources: "The Cup Turneth Over," Time*, April 10, 1989, p. 42; *Raleigh News and Observer*, September 20, 1989.)

Creativity and Decision Making

The majority of businessmen are incapable of original thinking, because they are unable to escape from the tyranny of reason. Their imaginations are blocked. (David Ogilvy, *Confessions of an Advertising Man*)

How ironic that we have spent so much effort in discussing and demonstrating the systematic analysis of decisions only to find this quote. Decision analysis *is* reason exemplified, is it not? If we relentlessly practice the form of reason that we have learned as decision analysis, will we find ourselves trapped by that reason, unable to be creative?

The emphatic answer is no. From a decision-making perspective, we need not be stuck with the alternatives that present themselves to us; in fact, good decision making includes active creation of new and useful alternatives. Moreover, the very process of decision analysis—especially the specification of objectives—provides an excellent basis for developing creative new alternatives.

Although it may seem unusual, a chapter on creativity in a decision-making text makes good sense. Everyone is frustrated occasionally by an inability to think creatively and thus can use a hand in being more creative. Perhaps not so clear, however is the increasing need for creative and innovative solutions. In *Thriving on Chaos* Tom Peters (1988) depicts the modern business climate as one in which conditions change rapidly. Modern managers must do more than simply cope with radical transformations: They must be on the attack. To be successful, a manager must learn to view new situations as opportunities for beneficial change rather than as problems to overcome somehow without rocking the boat too much. Indeed, Peters argues that the core paradox a manager faces is building an organization that is stable in its ability to innovate rapidly and flexibly. Solving problems creatively must

become part of a manager's and a firm's essence. And corporate America appears to take this message seriously; Bazerman (1994) reports that a majority of large U.S. corporations engaged in formal creativity training during the early 1980s.

This chapter presents a short overview of creativity. The literature on creativity is large and growing; interested readers will have no problem locating additional material. We frame the discussion first by defining creativity in decision making and looking at some of the psychological theories that have been developed to help us understand the creative process. What follows is a discussion of the many different ways in which creativity can be blocked. Finally, we discuss ways to enhance the creative process, especially the process of generating alternatives in a decision-making situation. One important technique promised back in Chapter 3 is the use of the means-objectives network. Other techniques use fundamental objectives, and still others (many developed by decision analysts) can be used to break through various blocks to find creative new alternatives.

What Is Creativity?

One thing is certain regarding the definition of creativity — it is much easier to identify creative acts than it is to define the term itself (Barron 1965). We readily recognize creative acts, and we often use adjectives like novel, insightful, clever, unique, different, or imaginative. But coming up with a coherent and useful definition of the term *creativity* is not easy.

Many different scholars have attempted to define creativity. All definitions include some aspect of novelty. But there is also an element of effectiveness that must be met. Slings buckets of mud at customers as they arrive at a used-car lot is indeed a novel greeting but may not be very effective in selling cars. But offering coupons for a cosmetic mud pack, an evening at the local mud-wrestling arena, or a therapeutic and relaxing mud bath might be very effective as well as novel. And what about creating a television advertisement in which the car lot's owner offers customers the opportunity to dunk him in a mud bath set up at the lot specifically for this purpose? ("My name is mud, because I ordered too many cars! So come on down and make the name stick! Throw me into the pit! And after you do that, check out some of these great deals...." Well, it would be no worse than many similar ads.)

For our purposes in this chapter, we are particularly concerned with the development of creative alternatives in decision problems. To be sure, creativity arises in many different situations; a novel and elegant proof of a mathematical theorem, an artist's creativity in painting or music, and a storyteller's clever retelling of an old tale are a few examples. When we think of creativity in decision making, though, we will be looking for new alternatives with elements that achieve fundamental objectives in ways previously unseen. Thus, a creative alternative has both elements of novelty and effectiveness, where effectiveness is thought of in terms of satisfying objectives of a decision maker, a group of individuals, or even the diverse objectives held by different stakeholders in a negotiation.

Theories of Creativity

Why do creative thoughts seem to come more readily to some people than to others? Or in certain kinds of situations? Many scholars have tried to understand the creative process, and in this section we review some of the psychological bases of creativity.

Perhaps the most basic approach relates creativity to Maslow's (1954) concept of self-actualization. For example, Davis (1986) describes self-actualization as, among other things, being able to perceive reality accurately and compare cultures objectively, having a degree of genuine spontaneity, and being able to look at things in a fresh, naive, and simple way. Davis claims that these and other qualities help people, even those without special talent, to act creatively, and he reviews some recent psychological evidence to support this proposition. Thus, Davis's position is good news for many of us. Self-actualization, happy lives, and creativity all seem to go hand in hand and to some extent can be developed by anyone. One need not have the special talent of an Einstein, Mozart, or Alexander the Great to reap the creative rewards that follow from self actualization.

Others have attempted to delve more deeply into the process of creative thought itself. Psychoanalytic theories (Kris 1952, Kubie 1958, Rugg 1963) generally maintain that creative productivity is the result of preconscious mental activity. These theories suggest that our brain is processing information at a level that is not accessible to our conscious thoughts. Behavioristic theories (Maltzman 1960, Skinner 1972) argue that our behavior, including creative behavior, is simply a conglomerate of responses to environmental stimuli. Appropriate rewards (stimuli) can lead to more creative behavior.


A cognitive approach suggests that creativity stems from a capacity for making unusual and new mental associations of concepts (Campbell 1960, Mednick 1962, Staats 1968). Campbell proposes that creative thought is just one manifestation of a general process by which people acquire new knowledge and thereby learn about the world. This process includes as the first step the production of "variations," a result of mentally associating elements of a problem in new ways. People who are more creative are better at generating a wider range of variations as they think about the problems they face. Having a broader range of life experiences and working in the right kind of environment can facilitate the production of variations. Finally, some people simply are better at recognizing and seizing appropriate creative solutions as they arise; the ability to come up with creative solutions is not very helpful if one ignores those solutions later.

Before going on, try the following problem. Do not read further until you have devoted some genuine effort to finding a solution.

CHAINS OF THOUGHT

You have four three-link chain segments as shown in Figure 6.1. A jeweler has offered to connect the segments to make a complete circle but to do so must open and then resolder some of the links. Opening and closing a link costs \$50. When you

Figure 6.1.
Chains of thought.
 Connect the segments
 by opening and closing
 only three links.



point out that you have only \$150, the jeweler says the job can be done for that amount. How can the jeweler connect the segments by opening and closing only three links? [Source: Adapted from Winkelgren (1974).]

Most people see the obvious solution of opening one link on each of the four segments and connecting to the adjacent segment, but this approach requires opening and closing four links. The creative and efficient solution, however, opens all three links of one segment and uses those three to connect the remaining three segments. This is an example of making a new association of the elements of the problem. Rather than thinking about the four segments as the elements of the problem, the creative solution considers open links as one kind of element, connected segments of closed links as another kind of element, and associates these by specifying that an open link must be used to connect two segments.

Phases of the Creative Process

A number of authors have identified phases of the individual creative thought process. For example, Wallas (1926) identified preparation, incubation, illumination, and verification.

1 Preparation In this first stage, the individual learns about the problem. This includes understanding the elements of the problem and how they relate to each other. It may include looking at the problem from different perspectives or asking other people what they know or think about the problem. From a decision-making point of view, this stage is very similar to problem structuring as delineated in Chapters 2 and 3. Spending effort understanding fundamental objectives, decisions that must be made (along with the immediately available set of alternatives), uncertainties inherent in the situation, and how these elements relate to each other prepares the decision maker for creative identification of new alternatives.

2 Incubation In the second stage, the prepared decision maker explores, directly or indirectly, a multitude of different paths toward new alternatives. We might also

use the terms *production* or *generation* of alternatives. The decision maker may do many things that seem to have a low chance of generating a new alternative, such as eliminating assumptions or adopting an entirely different perspective. Apparently frivolous activities may evoke the idea of the decision maker "playing" with the decision.

Many authors have included in this phase unconscious processing of information known about the decision. The literature on creativity contains several well-known and oft-quoted examples of this unconscious process, including Kekule's dream that led to the discovery of the chemical structure of the benzene ring and Poincare's discovery of the fuchsian family of groups and functions in mathematics. More recently, Gay Balfour of Colorado made the U.S. national news (*Associated Press*, January 24, 1992) by literally dreaming up a way to control prairie dog populations on ranches; he sucks them up with a giant vacuum machine created from a septic-system cleaning truck. (He hastens to explain that sucking the critters out of their holes causes them no harm at all; even though they may be slightly confused, shortly after landing in the holding area they begin burrowing in the soft dirt sucked up with them!) Balfour's new technique has reportedly been quite successful in the southwestern United States.

One explanation of unconscious incubation as a valid element of the creative process has been suggested by researchers in artificial intelligence. The explanation is based on a "blackboard" model of memory in the human brain. When the brain is in the process of doing other things — when a problem is incubating — parts of the blackboard are erased and new items put up. Every so often, the new information just happens to be pertinent to the original problem, and the juxtaposition of the new and old information suggests a creative solution; in other words, the process of coming up with a new and unusual association can result simply from the way the brain works. An attractive feature of this theory is that it explains why incubation works only a small percentage of the time. Too bad it works so infrequently!

Beyond the literature's examples and the above speculation about how the brain works, however, there is little hard evidence that unconscious processes are at work searching for creative solutions (Baron 1988). Still, it is a romantic thought, and if it provides an excuse for many of us to relax a little, perhaps the reduced stress and re-focus of our attention is enough to help us be more creative when we are trying to be.

3 Illumination This is the instant of becoming aware of a new candidate solution to a problem, that flash of insight when all the pieces come together, either spontaneously (Aha!) or as the result of careful study and work. Wallas described illumination as a separate stage, but you can see that illumination is better characterized as the culmination of the incubation stage.

4 Verification In the final step the decision maker must verify that the candidate solution does in fact have merit. (How many times have you thought you had the answer to a difficult problem, only to realize later — sometimes moments, sometimes much later — that your "dream solution" turned out to be just that: an impossible dream?) The verification stage requires the careful thinker to turn back to the hard

logic of the problem at hand to evaluate the quality of the candidate solution. In our decision-making context, this means looking very carefully at a newly invented alternative in terms of whether it satisfies the constraints of the problem and how well it performs relative to the fundamental objectives.

Although there are many ways to think about the creative thought process, the cognitive approach described above, including the stages of creativity, can help us to frame the following discussion. We turn now to ways in which our creativity can be hindered, and we follow with suggestions about how to reduce or eliminate such blocks and thereby increase our creativity.

Blocks to Creativity

With a clear understanding of the creative process, we are now in a position to discuss ways in which that process can be derailed, albeit inadvertently. This section describes three kinds of creativity blocks, drawing heavily from work by Adams (1979), Baron (1988), Bazerman (1994), Hogarth (1987), and Kleindorfer, Kunreuther, and Schoemaker (1993). All of these blocks interfere with the creativity process by hindering the generation and recognition of new and unusual solutions to a problem or alternatives in a decision situation.

Framing and Perceptual Blocks

These blocks arise because of the ways in which we tend to perceive, define, and examine the problems and decisions that we face. To get a feel for these blocks, consider the following two problems. As with the earlier chain puzzle, give serious effort to solving these problems before reading on.

THE MONK AND THE MOUNTAIN

At dawn one day, a monk begins to walk along a path from his home to the top of a mountain. Never straying from the path, he takes his time, traveling at various speeds, stopping to rest here and there, and arrives at the top of the mountain as the sun sets. He meditates at the top of the mountain overnight and for the next full day. At dawn the following morning, he begins to make his way back down the mountain along the same path, again relaxing and taking his time, and arrives home in the afternoon. Prove that there is a spot along the path that the monk occupies at the same time of day going up and coming down. [*Source*: Adapted from Hogarth (1987, p. 161).]

MAKING CIGARS

Lonesome Molly loves to smoke cigars, and she has learned to make one out of five cigar butts. Suppose she collects 25 butts. How many cigars can she make? [*Source: Adapted from Bartlett (1978).*]

Did you solve either one? The monk problem is difficult to solve if you try to maintain the frame in which it is cast: that the same monk travels up and down on two different days. Change the frame, though, and imagine identical monks taking the uphill and downhill journeys on the same day, each one starting at dawn. Now it is easy to see that the monks will meet somewhere along the path at some time during the day. At that instant, they are at the same spot in the path; hence in the original problem there must be a point along the path that the monk in the problem occupies at the same time going up and coming down.

Lonesome Molly can obviously make at least five cigars out of the 25 butts that she finds. But you were right to suspect that the obvious answer is not correct! The answer lies in thinking not about the gross requirement of butts to make a cigar but to frame the problem in terms of net usage. Molly indeed requires five butts to make a cigar. For each one she makes (and smokes), however, she has one butt left. The net consumption per cigar is four butts. So if she has 25 butts to start with, she can make six cigars and have one butt left over: She makes five cigars out of the original 25 butts and smokes those five cigars, which yields five butts from which she can make a sixth cigar. After she smokes the last one, she has one butt left (and needs to find only four more for her next smoke).

Here are some specific blocks relating to framing and perception that hinder our creative potential:

1 Stereotyping Suppose you are a personnel manager, and an individual with long hair and no necktie applies for a job as an engineer. Imagine your reaction. What would you think about the person? A typical mental strategy that most people use is to fit observations (people, things, events, and so on) into a standard category or stereotype. Much of the time this strategy works well because the categories available are rich enough to represent most observations adequately. But when new phenomena present themselves, stereotyping and associated preconceived notions can interfere with good judgment.

2 Tacit Assumptions Consider the classic nine-dot puzzle. Lay out nine dots in a square, three dots per row (Figure 6.2), and then, without lifting your pencil, draw four straight lines that cross all nine dots. Try it before you read on. The epilogue to this chapter gives the standard solution as well as many surprising solutions that Adams (1979) has collected from creative readers.

Figure 6.2*Nine-dot puzzle.*

Connect the dots using four straight lines without lifting your pencil.

[Source: From *Conceptual*

Blockbusting: A Guide to Better Ideas by


James L. Adams. © 1974, 1976, 1979,

1986 by James L.

Adams. Stanford

Alumni Association, Stanford, CA and

James L. Adams.]



The nine-dot puzzle is a nice example, but what does this block have to do with decision making? People often look at problems with tacitly imposed constraints, which are sometimes appropriate and sometimes not. Suppose you believe you need more warehouse space for your business, so you have your real-estate agent look for warehouses of a specific size to rent. The size, however, may be an inappropriate constraint. Perhaps so much space is not necessary or may be divided among several smaller warehouses. Perhaps some characteristic of your product would permit a smaller warehouse to be modified in a clever way to provide adequate storage.

3 Inability to Understand a Problem at Different Levels This block can be manifest in different ways. First is the familiar issue of isolating the precise decision context that requires attention. Suppose you are a national sales manager for a line of boots. Sales in the Rocky Mountain states are down. Knowing your regional sales person, you suspect that the problem is motivational. The "obvious" solution is to threaten or cajole the salesperson into better sales. But is the problem just what you think? Could it be a marketing problem — for example, competition with a regional

brand that has been developed specifically for the area? What about a distribution problem? Perhaps it is difficult for the one warehouse in the region to supply the area's special needs. Perhaps customers in the region finally are getting tired of the same old style that has been the company's cash cow for many years. Even if the problem does lie with the salesperson, other possibilities exist, such as personal problems or personality conflicts with local business owners.

Another manifestation is focusing too much on detail and not being able to reframe the decision in a broader context, a problem commonly called "not seeing the forest for the trees." Many decisions require attention to a large amount of detailed information. For example, consider the issues involved in deciding whether to attempt a takeover of another firm, or where to site a new manufacturing plant. The sheer volume of information to be processed can keep the decision maker from seeing new and promising alternatives.

4 Inability to See the Problem from Another Person's Perspective Where the previous block relates to seeing the problem itself in different ways, this one relates to seeing the problem through someone else's eyes and with their values. When a decision involves multiple stakeholders, it is always important to understand the values, in-

terests, and objectives of other parties. Really creative solutions incorporate and satisfy as many competing objectives as possible, and an inability to understand others' values can interfere with the development of such solutions. For example, finding a meaningful way to achieve peace in the Middle East requires the parties to consider the interests of both Israelis and Palestinians, as well as other nations in the region.

Value-Based Blocks

Blocks in this category relate to the values we hold. In many cases our values and objectives can interfere with our ability to seek or identify truly creative alternatives in a decision situation.

1 Fear of Taking a Risk To get a feel for this block, try the following game at a party with a lot of friends. Each person is assigned to be a particular kind of barnyard animal: cow, donkey, chicken, goat, sheep, or whatever else you designate. The more people, the better. After everyone has been assigned to be an animal, the organizer counts to three. On the count of three, each person looks directly at his or her nearest neighbor and makes the sound of his or her animal as loudly as possible. For obvious reasons, this is called the Barnyard Game (Adams 1979). Almost all participants feel some reluctance to play because they risk appearing silly in front of their friends.

There is nothing inherently wrong with being afraid to take a risk. In fact, the idea of risk aversion is a basic concept in decision making under uncertainty; we have seen, for example, that the basic risky decision as described in Chapter 3 requires the decision maker to determine whether the risk of a loss (relative to a sure thing) is justified by a possible but uncertain gain. It may be counterproductive, though, not to offer a creative alternative for consideration in a decision problem because you risk others thinking your idea is impossible, too "far out," or downright silly. What are the consequences of presenting a far-out idea that turns out to be unacceptable? The worst that might happen is that the idea is immediately determined to be infeasible. (Making far-out suggestions can have a more subtle value. Outsiders often have a difficult time understanding exactly what the problem is. Presenting far-out ideas for action is a sure way to get a clear statement of the problem, couched in an explicit and often supercilious explanation of why the idea will not work. Although this technique cannot be used in every situation, when it works the result is a better understanding of the decision situation.)

2 Status Quo Bias Decision making automatically means that the decision maker is considering at least one alternative that is different from the status quo. As indicated in the opening of this chapter, the ability to deal with change is becoming increasingly important for managers and decision makers. Studies show, however, that many people have a built-in bias toward the status quo. The stronger that bias, the more difficulty one may have coming up with creative problem solutions and alternatives.

3 Reality versus Fantasy An individual may place a lot of value on being realistic and a low value on fantasizing. Creative people must be able to control their

imagination, and they need complete access to it. Many exercises are available for developing an enhanced imagination and the ability to fantasize. Richard de Mille's *Put Your Mother on the Ceiling* (1976) has many imagination games. Although designed primarily for children, going through one of these games as an exercise in using fantasy can provide a remarkable experience for anyone. An excerpt from one of these games is reproduced in *Breathing* (see below). For the best effect, have a friend read this to you, pausing at the slash marks, while you sit quietly with your eyes closed.

4 Judgment and Criticism This block arises from applying one's values too soon in the creative process. Rather than letting ideas flow freely, some individuals tend to find fault with ideas as they arise. Such fault finding can discourage the creation of new ideas and can prevent ideas — one's own or someone else's — from maturing and gathering enough detail to become usable. Making a habit of judging one's own thoughts inevitably sacrifices some creative potential.

Breathing

Let us imagine that we have a goldfish in front of us. / Have the fish swim around. / Have the fish swim into your mouth. / Take a deep breath and have the fish go down into your lungs, into your chest. / Have the fish swim around in there. / Let out your breath and have the fish swim out into the room again.

Now breathe in a lot of tiny goldfish. / Have them swim around in your chest. / Breathe them all out again.

Let's see what kinds of things you can breathe in and out of your chest. / Breathe in a lot of rose petals. / Breathe them out again. / Breathe in a lot of water. / Have it gurgling *in* your chest. / Breathe it out again. / Breathe in a lot of dry leaves. / Have them blowing around in your chest. / Breathe them out again. / Breathe in a lot of raindrops. / Have them pattering in your chest. / Breathe them out again. / Breathe in a lot of sand. / Have it blowing around in your chest. / Breathe it out again. / Breathe in a lot of little firecrackers. / Have them all popping in your chest. / Breathe out the smoke and bits of them that are left. / Breathe in a lot of little lions, / Have them roaring in your chest. / Breathe them out again.

Breathe in some fire. / Have it burning and crackling in your chest. / Breathe it out again. / Breathe in some logs of wood. / Set fire to them in your chest. / Have them roaring as they burn up. / Breathe out the smoke and ashes...

Be a fish. / Be in the ocean. / Breathe the water of the ocean, in and out. / How do you like that? / Be a bird. / Be high in the air. / Breathe the cold air, in and out. / How do you like that? / Be a camel, / Be on the desert. / Breathe the hot wind of the desert, in and out. / How does that feel? / Be an old-fashioned steam locomotive, / Breathe out steam and smoke all over everything. / How is that? / Be a stone. / Don't breathe. / How do you like that? / Be a boy (girl). / Breathe the air of this room in and out. / How do you like that?

Source: de Mille, Richard (1976). *Put Your Mother on the Ceiling*. New York: Viking Penguin.

Cultural and Environmental Blocks

All decisions are made in some sort of social and cultural environment. The blocks that we describe here represent ways in which that environment may hinder the production and recognition of creative alternatives in decision situations.

1 Taboos This type of block has to do with what is "proper behavior" or "acceptable" in a cultural sense; taboos may exist for no apparently good reason. The following problem (adapted from Adams 1979) demonstrates this block. As before, think about this problem before reading further.

PING-PONG BALL IN A PIPE

You are in a room with six other people. The room is entirely barren except for a steel pipe embedded solidly in the concrete floor and extending 25 centimeters above the floor. The upper end of the pipe is open, and a Ping-Pong ball is at the bottom of the pipe as in Figure 6.3. Your job is to get the Ping-Pong ball out of the pipe without damaging the ball, the pipe, or the concrete. You have only the following items available that you can use to extricate the ball:


- 10 feet of cotton string
- 5 ounces of dry cereal
- A wire coat hanger
- A steel file
- A chisel
- A hammer with a wooden handle
- A monkey wrench
- A light bulb

Often people come up with solutions like fashioning a long set of tweezers out of the coat hanger or smashing the hammer handle and using the splinters to retrieve the ball. But occasionally a few intrepid individuals cross a cultural borderline and realize that the simplest way to get the ball out is to have people in the group urinate in the pipe! Of course, urinating in such a public situation is somewhat taboo in our western culture, effective though it may be for solving the problem.

For a more realistic example, suppose one of your co-workers has a new baby and wishes to bring the child to work so that she can continue to nurse the child. Certain taboos are involved here, including nursing in public and having a child in the workplace during "serious" work time. Should the taboos be violated? A creative alternative would find a way to accommodate the mother without grossly violating the taboos.

Figure 6.3

Ping-Pong ball in a pipe. Extract the Ping-Pong ball from the steel pipe without damaging the concrete, pipe, or ball.



2 Strength of Tradition As we mentioned previously, individuals can resist change because of a bias toward the status quo. There is a cultural counterpart to this block; in many cases, the sociocultural environment in which a decision maker operates places a high value on maintaining tradition. Adopting change can be difficult in such a situation, which in turn can hinder the production of creative suggestions in the first place. For example, the musical *Fiddler on the Roof* describes the tradition-bound culture of Russian Jews in the early twentieth century and a father's difficulty in dealing with his daughters' new ways of finding husbands.

3 Reason and Logic versus Humor, Fantasy, and Artistic Thinking There is a clear block against using feelings, intuitions, and emotions in business problem solving. Certainly valuable insights and understanding come from analytical treatments of any given problem; indeed these skills are important in decision making, and a course in decision analysis offers to teach such skills. However, valuable cues and ideas can also arise by admitting and examining feelings, intuitions, and emotions. For example, doing so can help understand the values of others who may have a stake in a decision. In the example above regarding bringing a child to work, reluctance to allow a child in the workplace may be due to the values of the workers who feel that the playfulness, fantasy, and humor that children represent should not displace reason and logic at work.

In a decision-making course much of the emphasis is on the development of analytical thinking. Unfortunately, little effort is put into more artistically oriented thinking skills such as using imagery, being playful, storytelling, or expressing and appreciating feelings. Such activities tend to be culturally blocked because of the stress placed on analysis. From the discussion in this chapter, it would appear that artistic thinking can play an important role in the development of creative alternatives. The best possible arrangement is for an individual to be "mentally ambidextrous," or good at switching between analytical and artistic thinking styles. This enhances creative development of potential alternatives without sacrificing subsequent careful analysis.

Organizational Issues

Without a doubt, different organizations have different characteristics or cultures, and organizational culture can have a strong influence on decision making. Many of the issues that we have already discussed can be a part of an organization's decision-making culture. For example, an organization may have a culture that in subtle ways promotes criticism and judging of ideas, stereotyping, or being risk-averse. Humor, playfulness, or artistic thinking may be frowned upon, or change may be resisted in order to preserve company traditions. For all of the reasons discussed above, such characteristics can reduce the creative potential of individuals in the organization.

By their very nature, organizations can impede creative thought. As Adams (1979, p. 143) points out, "the natural tendency of organizations to routinize, decrease uncertainty, increase predictability, and centralize functions and controls is certainly at odds with creativity." Other features of organizations also can hinder creativity. Examples include excessive formal procedures (red tape) or lack of cooperation and trust among co-workers. Hierarchical organizational structures can hinder creativity, which in turn can be exacerbated by supervisors who tend to be autocratic.

Teresa Amabile has studied creativity and organizations for over twenty years. Her work has led to a detailed model of individual creativity in the organizational context (Amabile, 1988). First, individual creativity requires three ingredients: expertise in the domain, skill in creative thinking, and intrinsic motivation to do the task well. In other words, we need someone who is good at what he or she does, who likes to do it just because it is interesting and fun, and who has some skill in creative thinking, perhaps along the lines of the creativity-enhancing techniques we discuss later in this chapter.

Amabile's work shows how the organizational environment can influence individual creativity. In particular, she warns that expecting detailed and critical evaluation, being closely watched, focusing on tangible rewards, competing with other people, and having limited choices and resources for doing the job all can hinder one's creativity. When she compared high- and low-creativity scenarios in organizations, though, the results indicated that a delicate balance must be maintained. For example, workers need clear overall goals, but at the same time they need latitude in how to achieve those goals. Likewise, evaluation is good as long as it focused on the work itself (as opposed to the person) and provides informative and constructive help. Such evaluation ideally involves peers as well as supervisors. Although a focus on tangible rewards can be detrimental, knowing that one's successful creative efforts will be recognized is important. A sense of urgency can create a challenging atmosphere, particularly if individuals understand the importance of the problem on which they are working. If the challenge is viewed as artificial, however, such as competing with another division in the company or having an arbitrary deadline, the effect can be to decrease creativity. Thus, although creativity is essentially an individual phenomenon, managers can have a significant impact on creativity in their organizations through goal setting, evaluation, recognition and rewards, and creating pressure that reflects a genuine need for a creative solution.

Finally, even though managers can help individuals in their organizations be more creative, one can develop a "blind spot" because of a long-term association with a particular firm; it becomes difficult to see things in a new light simply because certain procedures have been followed or perspectives adopted for a long time. The German word *betriebsblind* for this situation literally means "company-blind." One of the important roles that consultants serve is bringing a new perspective to the client's situation.

Value-Focused Thinking for Creating Alternatives

Keeney (1992, Chapters 7 and 8) describes a number of different ways in which fundamental and means objectives can be used as a basis for creating new alternatives for decision alternatives. In this section we review some of these techniques.

Fundamental Objectives

The most basic techniques use the fundamental objectives directly. For example, take one fundamental objective and, ignoring the rest, invent a (possibly hypothetical) alternative that is as good as it could be on that one objective. Do this for each fundamental objective one at a time, and keep track of all of the alternatives you come up with. Now go back and consider pairs of objectives; what are good alternatives that balance these two objectives? After doing this for various combinations of objectives, look at the alternatives you have listed. Could any of them be modified so that they would be feasible or perhaps satisfy the remaining objectives better? Can any of the alternatives be combined?

A related approach is to consider all of the fundamental objectives at once and imagine what an alternative would look like that is perfect in all dimensions; call this the *ideal* alternative. Most likely it is impossible, but what makes it impossible? If the answer is constraints, perhaps some of those constraints can be removed or relaxed.

Still another possibility is to go in the opposite direction. Find a good alternative and think of ways to improve it. The fact that the alternative is a good one in the first place can reduce the pressure of finding a better one. In searching for a better one, examine the alternative carefully in terms of the objectives: On which objectives does it perform poorly? Can it be improved in these dimensions? For example, Keeney (1992) describes an analysis of possible sites for a hydroelectric power plant. One of the potential sites was very attractive economically but had a large environmental impact, which made it substantially less desirable. On further study, however, the design of the facility at the site was modified to reduce the environmental impact while maintaining the economic advantage.

Means Objectives

We mentioned back in Chapter 3 that the means objectives can provide a particularly fruitful hunting ground for new alternatives. The reason for this is simply that the

means objectives provide guidance on what to do to accomplish the fundamental objectives. In complicated problems with many fundamental objectives and many related means objectives, this approach can generate many possible courses of action. For example, consider the following decision situation.

TRANSPORTATION OF NUCLEAR WASTE

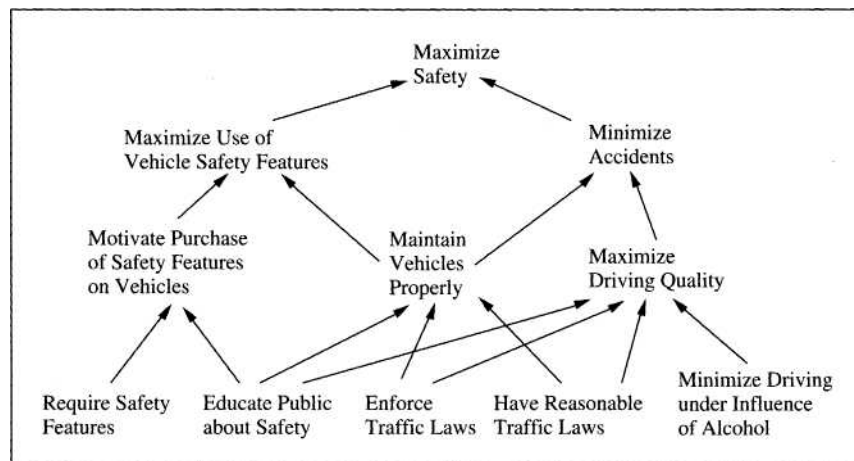
One of the problems with the use of fission reactors for generating electricity is that the reactors generate substantial amounts of radioactive waste that can be highly toxic for extremely long periods of time. Thus, management of the waste is necessary, and one possibility is to place it in a storage facility of some sort. Transporting the waste is itself hazardous, though. In describing the problem, Keeney (1992) notes that the decision situation includes the selection of a type of storage cask in which the material will be shipped, followed by the selection of a transportation route and a choice as to how many casks to ship at once. The uncertainties include whether an accident occurs, the amount of radiation released, and whether an efficient evacuation plan exists when and if an accident occurs.

Means objectives are associated with each of the decisions and uncertainties. For example, a means objective is to select the best possible cask, and that might include designing a special kind of cask out of a particular material with appropriate size and wall thickness specifications. Selecting a transportation route that travels through a sparsely populated area is a means objective to reduce potential exposure in the case of an accident. In selecting the number of casks to ship at once, one would want to balance the chance of smaller accidents with more frequent but smaller shipments against the chance of a larger accident with larger and less frequent shipments.

In examining the uncertainties, obvious means objectives come to mind. For example, an important means objective is to reduce the chance of an accident, which in turn suggests strict rules for nuclear-waste transportation (slow speeds, driving during daylight hours, special licensing of drivers, additional maintenance of roads along the route, and so on). Reducing the amount of radiation released in an accident and increasing the chance of an efficient evacuation plan being in place suggest the development of special emergency teams and procedures at all points along the transportation route. [*Source*: Keeney (1992, pp. 205-207).]

In another example, we can use a means network directly. Take the automotive-safety example that we discussed in Chapter 3; the problem is to find ways to maximize safety of automotive travel, and according to the fundamental objectives (Figure 3.1), maximizing safety means minimizing minor injuries, serious injuries, and deaths of both adults and children. The means-objectives network is reproduced here as Figure 6.4. You can see that each node in the network suggests particular alternatives. The objective of maximizing use of vehicle-safety features suggests that currently available features might be mandated, or incentives might be provided to

Figure 6.4
A means-objectives network for improving automotive safety.
Source: Keeney (1992, p. 70).



manufacturers. Likewise, incentives (e.g., a reduced insurance premium) might be given to automobile drivers for using safety features. Farther down in the network we find the objective of educating the public. We already have widespread driver-education programs for new drivers, but what about providing incentives for adult drivers to take periodic refresher courses? To combat drunken driving in particular, many states in the United States have implemented a combination of tougher traffic laws and public-education programs.

You can see from these examples how useful the means objectives can be for identifying new alternatives. Moreover, using a means-objectives network can ensure that as many aspects of the decision problem as possible are covered; the decision maker can see exactly what the new alternatives help achieve and can perhaps further develop the alternatives to attain a level of balance among the fundamental and means objectives. For example, in examining a set of safety proposals, matching the proposals against means and fundamental objectives might reveal that safety of children has not been explicitly addressed or that the set of proposals currently under consideration do nothing to encourage proper vehicle maintenance.

The Decision Context

Finally, it is always possible to broaden the decision context as part of the search for new ideas. We mentioned this possibility back in Chapter 3 in discussing decision structuring. As argued above, part of the creative process requires that the decision maker look at a problem from as many different perspectives as possible, and considering a broader context is guaranteed to reveal a different view of a decision situation. Take the automotive-safety example, and suppose the issue is broadened from automotive safety to transportation safety in general. At a microlevel in a city, this might imply a comprehensive education program for local cyclists, pedestrians, and motorists. At a county or state level, one might consider incentives to reduce urban sprawl and increase the use of alternative transportation methods.

For another example, reconsider the nuclear-waste transportation problem. Deciding how to transport the waste may arise from a prior decision to place a storage site in a designated location. But suppose we broaden the decision context to consider a variety of storage options. Doing so suggests possibilities like storing waste at the nuclear power plants themselves, which in turn suggests finding ways to reduce waste production or to store it more efficiently at these sites. Another possibility is to store the waste in smaller but more numerous waste facilities rather than in fewer and larger sites. Although some of these alternatives may not be reasonable to pursue, broadening the decision context can help uncover new ways to address the decision at hand and may suggest solutions of a more general nature.

Other Creativity Techniques

Fluent and Flexible Thinking

Fluency and flexibility of thinking are important in enhancing creativity. Fluency is the ability to come up with many new ideas quickly. Flexibility, on the other hand, stimulates variety among these new ideas. An individual who can write down many ideas quickly, regardless of what they may be, would be a fluent thinker. The flexible thinker might have a shorter list of ideas, but the ideas would tend to cover a broader range of possibilities. An individual who is both fluent and flexible can write down many different ideas quickly. One useful analogy compares thinking with digging holes. Fluent thinking is seen as the ability to dig one hole very deep and very quickly by taking a lot of dirt from one place. Flexible thinking, however, is more like the ability to dig many smaller holes in many different places.

A common exercise that is used to demonstrate these ideas is to think of new uses for common red construction bricks. The goal is to come up with marketing ideas that a brickyard owner could pursue to get out of financial difficulties. A list with many uses that are variations on a common theme indicates fluency. For example, many of the uses may be ways to use bricks to build things, or uses that take advantage of a brick's weight (for example, as a paperweight or ballast for hot-air balloons). A list with a lot of variety in different attributes of the bricks indicates flexibility. For example, flexible thinking would use many different attributes of a brick: strength, weight, color, texture, hardness, shape, and so on. Understanding the difference between the two different kinds of thinking is the first step toward developing enhanced thinking skills in either fluency or flexibility or even both.

Idea Checklists

One classic technique for enhancing creativity uses checklists that cover many potential sources of creative solutions to problems. Most of us use lists in a rather natural way. The yellow pages provide a simple and ubiquitous list that can provide

many ideas for solutions to specific problems. (Are any plumbers available on Sunday? Any do-it-yourself stores close by?) Mail-order and retail catalogs are examples of lists that we may use to solve a gift-giving problem. Many authors have written general-purpose lists for problem solving. The best known is Osborn's (1963) 73 *Idea-Spurring Questions*, reproduced below. Of course, Osborn's list can be extended with other descriptive verbs such as multiply, squeeze, lighten, propel, and flatten.

Another creativity-enhancing technique is to write down attributes of a problem, list alternative options under each attribute, and then consider various combinations and permutations of the alternatives. This use of lists was suggested by Koberg and Bagnall (1974), who dubbed the technique *morphological forced connections* to emphasize the combination of morphological attributes in design problems. The technique is not limited to designing new products. In fact, an early variant of this technique was used to think about objectives in a decision situation. Dole et al. (1968 a,b) describe an application by the National Aeronautics and Space Administration (NASA) for determining the scientific objectives of space-exploration missions. Action phrases (such as "measure tidal deformations in . . . ") were combined with target features (such as "the interior of . . . ") and target subjects (Jupiter, for example) to create a possible scientific objective. The candidate objective then was considered to determine whether it was a valid scientific objective. If so, it was included

Osborn's 73 Idea-Spurring Questions

Put to other uses? New ways to use as is? Other uses if modified?

Adapt? What else is like this? What other idea does this suggest? Does the past offer a parallel? What could I copy? Whom could I emulate?

Modify? New twist? Change meaning, color, motion, sound» odor, form, shape? Other changes?

Magnify? What to add? More time? Greater frequency? Stronger? Higher? Longer? Thicker? Extra value? Plus ingredient? Duplicate? Multiply? Exaggerate?

Minify? What to subtract? Smaller? Condensed? Miniature? Lower? Shorter? Lighter? Omit? Streamline? Split up? Understate?

Substitute? Who else instead? What else instead? Other ingredient? Other material? Other process? Other power? Other place? Other approach? Other tone of voice?

Rearrange? Interchange components? Other pattern? Other layout? Other sequence? Transpose cause and effect? Change pace? Change schedule?

Reverse? Transpose positive and negative? How about opposites? Turn it backward? Turn it upside down? Reverse roles? Change shoes? Turn tables? Turn other cheek?

Combine? How about a blend, an alloy, an assortment, an ensemble? Combine units? Combine purposes? Combine appeals? Combine ideas?

Source: From A. F. Osborn, *Applied Imagination*, 3rd edition. Copyright © 1963. All rights reserved. Adapted by permission of Allyn & Bacon.

in the list. For example, "Establish the structure of the interior of the sun" was an objective, but "Determine the characteristic circulation patterns in the photosphere of the space environment" was not.

The value of morphological forced connections is not so much to find all possible combinations as much as to provide a framework within which all imaginable combinations can be screened easily to determine the most appropriate candidates. In the NASA example, candidate objectives were generated readily and then screened for validity. Another example, reported by Howard (1988), comes from an advertising claim by a fast food hamburger chain that, with its custom service, one can order 1024 different kinds of hamburgers; that is, a customer could order a burger with or without each of 10 possible ingredients. Not all of the possible combinations are reasonable, however; for example, one combination is a "burger" that consists only of lettuce, or only of ketchup. One is the "nullburger": nothing at all. Most individuals would agree that many of these unusual combinations must be screened out. In fact, a burger may not be a burger without the beef patty and a bun.

A more serious example also comes from Howard (1988), who suggests the *strategy-generation table*. Figure 6.5 shows a typical strategy-generation table for an energy conglomerate that is considering possible expansion. Dividend payout and dividend-to-equity ratio also are important attributes for the conglomerate to consider in strategic planning. For the most part, the table is self-explanatory. An overall strategy is one that has individual elements in each column. As with the hamburger, however, not all combinations make sense. For example, because of cash constraints it is not likely that the conglomerate could pursue aggressive expansion in each of its five existing businesses, acquire another firm, and also have a high dividend payout. The strategy outlined in Figure 6.5 is one that might be described by executives as a "service business" strategy. Other feasible strategies could be assembled using the strategy-generation table. The point here is that the morphological forced connections technique facilitates both generation and screening to come up with a list of reasonable alternatives.

Figure 6.5

Strategy-generation table.

Source: Reprinted by permission of Ron Howard, "Decision Analysis: Practice and Promise," *Management Science*, 34, No. 6, June 1988, pp. 679-695. Copyright 1988, The Institute of Management Sciences.

Utility	Exploration and Production	Oil Field Services	Forest Products	Coal	Acquisition	Dividends	D/E Ratio
Aggressive Supply Buildup	Increase Exploration Budget	Aggressive Expansion	Hold-Improve Earnings	Purchase Additional Reserves	None	70%	1/1
Hold/Restricted Investment	\$800M Investment	Modest Expansion	Add Timberlands	Joint Venture Synfuels	Service Business	50%	1.5/1
Severe Capital Constraint	Sell/Milk	Hold	Sell	Hold	Resource Business	25%	2/1
		Milk		Milk		0%	

Brainstorming

Brainstorming is another popular way of generating a long list of ideas quickly. To be effective, a brainstorming session should include at least two people, and probably no more than 8 or 10 (it can be difficult to keep up with all of the ideas generated by a group that is very large). The rules for a brainstorming session are simple:

- 1 No evaluation of any kind is permitted.
- 2 All participants should think of the wildest ideas possible.
- 3 The group should be encouraged to come up with as many ideas as possible.
- 4 Participants should try to build upon or modify ideas of others.

Brainstorming works well for several reasons. The most important is probably the lack of any judgment, which eliminates an important block for many people. That so many ideas are created rapidly reassures those who have little faith in their own creative potential: The enthusiasm of a few individuals tends to be contagious, and a "one-upmanship" game usually develops as participants try to top previous ideas. Practitioners report that the technique tends to generate a first rash of ideas utilizing common solutions. After this initial phase, participants must come up with new concepts. Naturally, the newer concepts are the most valuable result of the brainstorming exercise.

Metaphorical Thinking

Given our earlier discussion of fantasy and imagination, it should come as no surprise that creativity can be enhanced by the use of metaphors. Much of the work on metaphorical thinking comes from William J. J. Gordon (1961, 1969), who founded Synectics, Inc., a consulting corporation that specializes in creative problem solving. Three kinds of metaphors can be used systematically to enhance creative potential: direct analogy, personal analogy, and fantasy analogy.

Direct analogy involves thinking about how others have solved problems similar to the one under consideration. Often the most productive approach is to examine solutions found in nature. For example, wet leaves, which pack together snugly, suggested metaphorically how potato chips might be packaged. The result was Pringles potato chips, which are stacked and sold in a can. The inventor of Velcro was inspired while removing burrs from his dog's fur. Suppose the problem is to design a home security system. What kinds of security systems do plants and animals have? Many animals make a lot of noise when threatened. Perhaps a security system could be designed that could sense intruders and then make noise — by turning on a stereo or TV inside a house — to frighten away intruders.

Personal analogy is closely related to the kinds of games that are played in de Mille's *Put Your Mother on the Ceiling*. The idea is to imagine a variety of personal situations that are pertinent in some way to the problem at hand. Personal analogy can be extremely helpful in computer programming, for example. Imagining the precise steps (and perhaps shortcuts) that would be taken to solve a problem by hand can help in designing a computer program to solve the same kind of problem.

In *fantasy analogy*, the group tries to come up with truly far-fetched, fantastic, and ideal solutions. The classic example arose when a group was trying to design an airtight closure for space suits that would be easily operated. Here is a transcript of part of the session, taken from Gordon (1961):

G: Okay. That's over. Now what we need here is a crazy way to look at this mess. A real insane viewpoint... a whole new room with a viewpoint!

T: Let's imagine you could will the suit closed... and it would do just as you wanted by wishing... [fantasy analogy]

G: Wishing will make it so...

F: Ssh, okay. Wish fulfillment. Childhood dream... you wish it closed, and invisible microbes, working for you, cross hands across the opening and *pull* it tight...

B: A zipper is kind of a mechanical bug [direct analogy]. But not airtight... or strong enough...

G: How do we build a psychological model of "willing-it-to-be-closed"?

R: What are you talking about?

B: He means that if we could conceive of how "willing-it-to-be-closed" might happen in an actual model—then we...

R: There are two days left to produce a working model—and you guys are talking about childhood dreams! Let's make a list of all the ways there are of closing things.

F: I hate lists. It goes back to my childhood and buying groceries...

R: F, I can understand your oblique approach when we have time, but now, with this deadline... and you still talking about wish fulfillment.

G: All the crappy solutions in the world have been rationalized by deadlines.

T: Trained insects?

D: What?

B: You mean, train insects to close and open on orders? 1-2-3 Open! Hup! 1-2-3 Close!

F: Have two lines of insects, one on each side of the closure—on the order to close they all clasp hands... or fingers... or claws... whatever they have... and then closure closes tight...

G: I feel like a kind of Coast Guard insect [personal analogy]. *D*:

Don't mind me. Keep talking...

G: You know the story... worst storm of the winter—vessel on the rocks... can't use lifeboats... some impatient hero grabs the line in his teeth and swims out...

B: I get you. You've got an insect running up and down the closure, manipulating the little latches...

G: And I'm looking for a demon to do the closing for me. When I will it to be closed [fantasy analogy], Presto! It's closed!

B: Find the insect — he'd do the closing for you!

R: If you used a spider... he could spin a thread... and sew it up [direct analogy].

T: Spider makes thread... gives it to a flea... Little holes in the closure... flea runs in and out of the holes closing as he goes...

G: Okay. But those insects reflect a low order of power ... When the Army tests this thing, they'll grab each lip in a vise one inch wide and they'll pull 150 pounds on it ... Those idiot insects of yours will have to pull steel wires behind them in order ... They'd have to stitch with steel. *Steel*

B: I can see one way of doing that. Take the example of that insect pulling a thread up through the holes ... You could do it mechanically ... Same insect... put holes in like so ... and twist a spring like this ... through the holes all the way up to the damn closure ... twist, twist, twist, ... Oh, crap! It would take hours! And twist your damn arm off!

G: Don't give up yet. Maybe there's another way of stitching with steel. ...

B: Listen ... I have a picture of another type of stitching ... That spring of yours ... take two of the ... let's say you had a long demon that forced its way up ... like this ...

R: I see what he's driving at. ...

B: If that skinny demon were a wire, I could poke it up to where, if it got a start, it could pull the whole thing together ... the springs would be pulled together closing the mouth ... Just push it up ... push — and it will pull the rubber lips together ... Imbed the springs in rubber ... and then you've got it stitched with steel!

Source: Gordon (1961).

Other Techniques

Several other techniques are available to enhance a group's creative potential. Many rely on methods for improving group interaction in general. For example, Nominal Group Technique (NGT) (Delbecq, Van de Ven, and Gustafson 1975) begins with no interaction. Individuals in the group each write down as many ideas as they can on pieces of paper. Then each individual in turn presents one of his or her ideas. The group leader records these ideas on a flipchart or chalkboard. Discussion begins after ideas from each participant are written down. At the end, each individual writes down his or her ranking or rating of the ideas. These are then combined mathematically to arrive at a group decision.

The main advantage of NGT is that the group leader manages the interaction of the group in such a way that certain blocks are avoided and the environment is enhanced. For example, discussion is not permitted until after the ideas are presented, thus creating a more supportive environment.

Other techniques actually use a more adversarial approach. Devil's advocacy and dialectical inquiry are techniques in which individuals take sides in a debate. On the surface, this might appear to hamper creative thought, but when such techniques are used only after ideas have been generated and a healthy creative environment has been established, they can work well. It also helps if all participating members understand what the techniques are meant to do. The main advantage of this kind of approach is that it can help a group of individuals consider a problem from multiple perspectives. Being forced into an alternative viewpoint can lead to new creative ideas. Another advantage is that the group is less likely to overlook basic issues that may be hidden from certain vantage points.

The role of the leader in group discussion techniques is paramount. A good leader sets the tone of the session, and a positive tone promotes an atmosphere that is conducive to healthy discussion and that encourages the free flow of ideas. It is easy to see that a group with such a leader probably will have more success in generating creative ideas and solving problems.

Creating Decision Opportunities

Creativity in decision making can be much more than generating new alternatives. A really creative decision maker is one who creates decision opportunities. Keeney (1992) stresses that an individual, group, or organization which understands its values and objectives clearly is perfectly positioned to look for decision opportunities proactively rather than merely reacting to decision problems served up by life. As you may agree, life does not always generously provide decision situations with many attractive options, but instead often seems to pose difficult decision problems that must be addressed, often under trying circumstances.

In Chapter 9 of *Value-Focused Thinking*, Keeney describes a wide variety of ways to use one's objectives to help identify new decision opportunities. The more obvious ones are to look at fundamental and means objectives and find ways to achieve them; simply choosing to look for ways to achieve ones objective is in itself creating an opportunity. As we discussed above, broadening a decision context can be worthwhile. Simply learning to say no to those who ask for time, money, or assistance can create decision opportunities for those who have a difficult time turning down meritorious requests. And, as Keeney explains, when you just do not know what to do, that is the perfect occasion to create a decision opportunity in which you can tailor the alternatives to match your objectives. The decision maker who searches actively and creatively for new opportunities can look forward to many exciting possibilities and a full life.

SUMMARY Creativity is important in decision making because the available alternatives determine the boundaries of the decision. We discussed various theories of creativity, including the phases of creative thought. We then introduced many different kinds of creativity blocks that hinder creative efforts. In discussing creativity-enhancing techniques, we began with the use of one's objectives as a basis for developing new ideas and showed how means objectives in particular can provide a fertile ground for generation of new alternatives. Other creativity techniques include becoming proficient in both flexible and fluent thinking, list making, brainstorming, and metaphorical thinking. Group discussion techniques can promote creativity through appropriate management of interaction in the group and can also enhance the creative environment. Finally, we argued that a truly creative decision maker goes beyond the creation of alternatives all the way to the development of new decision opportunities.

QUESTIONS AND PROBLEMS

- 6.1** In talking about cultural and environmental blocks, we discussed the matter of a woman bringing an infant into the workplace and the possible reactions of others. Consider the situation in which a classmate wants to bring her child to school in order to continue nursing. Are the taboos the same? What are the differences between the two situations?
- 6.2** Choose a decision that you currently face. What are your objectives in this situation? List your means objectives, and for each means objective list at least one alternative that could help achieve that objective.
- 6.3** An important objective of most academic programs is to get students to be more creative. How can this be done? What specific exercises can you think of that would help? What kinds of things can an instructor do? Any changes to the curriculum or requirements? Try some of the creativity-enhancing techniques that we discussed. For example, you might try any of the following:
- a** List fundamental and means objectives that you think would be appropriate for your academic program. Use these objectives as a basis for generating ways to make students more creative.
 - b** List making, especially morphological forced connections. What are important attributes to list? What are alternatives under those attributes? Can you use Osborn's list? Can you think of extensions to his list?
 - c** Brainstorming. To do this, you need to find a small group of people to brainstorm with. The exact size of the group is not important, but you need at least one more besides yourself, and probably no more than 8 or 10 altogether.
 - d** Metaphorical thinking. Direct analogy is difficult here. How do other disciplines teach students to be creative? How do animals teach offspring to deal with new situations? Personal analogy is somewhat more straightforward. Imagine that you are the boss. What do you want your employee to be able to do? Imagine that you are a professor. Would you want your students to be creative? How would you encourage it? Can you imagine that you are a problem that needs to be solved? How would you like to be solved? Try fantasy analogy. Try to imagine the most fantastic and ideal creativity training program. Describe it in detail; exercise your imagination! What kinds of resources would be required? How many people? How much time? What kinds of interactions with others? Would there be instructors and students, or just participants, all on the same level? Think of other questions.
- You are not required to use any of these techniques, nor are you limited to them. Your job is to generate some creative ideas any way you can.
- 6.4** How would you design an organization so that it could, in Tom Peters's words, "thrive on chaos"? What characteristics would such an organization have? What kind of people would you try to hire? How would the role of the managers differ from the traditional view of what a manager does?
- 6.5** In discussing the perceptual block of stereotyping, we used the example of a person with long hair and no necktie applying for a job as an engineer. Did you imagine this person as male or female? Why? Was there a block involved in your perception? Describe it.
- 6.6** The point is often made that formal schooling can actually discourage young children from following their natural curiosity. That curiosity is an important element of creativ-

ity, and so it may be the case that schools indirectly and inadvertently are causing children to become less creative than they might be. What does this imply for those of us who have attended school for many years? What can you suggest to today's educators as ways to encourage children to be curious and creative?

- 6.7 Use the means-objectives network in Figure 6.4 to create a list of alternatives for improving automotive safety. Try to create at least one alternative for each objective listed.
- 6.8 Reconsider the summer-intern decision discussed in Chapter 3. Figure 3.4 shows the fundamental and means objectives. With a group of friends, use these objectives along with brainstorming to generate a list of at least 20 things PeachTree can do to help locate and hire a good summer intern.
- 6.9 The lists that were discussed focused on the generation of alternatives. However, the same technique can be used to determine objectives in a decision situation. Write a list of questions that you might ask yourself when searching for a job. What do you want to accomplish? For example, do you want to save money? Save time? Improve your lifestyle? What else is important?
- 6.10 Describe a situation in which unconscious incubation worked for you. Describe one in which it did not. Can you explain why it worked in the first case but not in the second?
- 6.11 One of the technological problems that we face as a society is the increasing use of plastics in disposable items. Landfills are becoming more and more expensive and difficult to maintain, and land for new ones can only be obtained at premium prices. Furthermore, the plastics that are dumped in the landfills may release dioxins (toxic chemical substances) into the soil.

Of course, many different kinds of plastic exist. Spend 10 minutes writing down possible ways to recycle one-gallon plastic milk jugs. You can assume that the milk jugs are received rinsed out and reasonably clean, but not sterile. Look at your list. Does it reflect fluent thinking, flexible thinking, or both?

C A S E S T U D I E S

MODULAR OLYMPICS

Seoul, Korea, was the site of the 1988 Summer Olympic Games. However, the *Wall Street Journal* (June 29, 1987) reported that because of political unrest in Korea, there was some concern about having the games there. Would it have been possible to change sites as late as 1987? It seems unlikely because of the expense and planning associated with putting the games on. Wouldn't it be nice to be able to disassemble the Olympics from one site, ship them to a new site, reassemble, and proceed? Modular Olympics.

Question

- 1 Use whatever creativity-enhancing techniques you can to help think about ways to modify the Olympics to make it possible to change locations more easily than is currently the case.

BURNING GRASS-SEED FIELDS

Grass-seed farmers in the Willamette Valley in western Oregon have been burning their fields since the 1940s. After the grass seed is harvested, the leftover straw is burned to remove it and to sterilize the fields. However, burning large fields of straw creates a lot of pollution. Another burning method is available that uses a propane torch to burn the stubble in a field after the straw has been removed.

In the face of controversy over the negative effects of burning, the grass-seed industry is considering alternative ways to sterilize the fields. The propane method is the most promising, but it first requires removal of the straw.

Question

- 1 What can be done with the straw? Spend 10 minutes writing down all the possible uses you can think of for straw. If possible, form a group and have a brainstorming session. Or try any other creativity technique that you can think of. Does your list demonstrate flexibility or fluency of ideas?

REFERENCES

It is rather unusual for a book on decision theory to have a chapter on creativity. The "management science" tradition would suggest that there is a set way of attacking a specific problem and that all the decision maker must do is apply the appropriate technique. However, this is too simple for the complex problems that managers face these days. Hence, this chapter is meant to dispel the notion that analytical thinking is all that is required to solve real problems.

Very little literature exists within management science on creativity. Keller and Ho (1989) review and summarize this literature. Some of Keeney and Raiffa's (1976, Chapter 2) discussion of the structuring of objectives is pertinent to creativity in decision making. As indicated in the text, Keeney (1992) devotes many pages to creativity. Kleindorfer, Kunreuther, and Schoemaker (1993) also discuss creativity from a management science perspective.

Some of the literature on creativity, like Adams (1979), comes from an engineering/design/inventing perspective. (In fact a recent creativity book by Adams (1986) contains a chapter on decision analysis!) Most of the literature, however, comes from psychology, as the many references at the beginning of the chapter indicate. Baron (1988), Bazerman (1994), and Hogarth (1987) provide good reviews of the creativity literature from a psychological perspective. Johnson, Parrott, and Stratton (1968) provide some insights on the value of brainstorming.

Adams, J. L. (1979) *Conceptual Blockbusting: A Guide to Better Ideas*, 2nd ed. Stanford, CA: Stanford Alumni Association.

Adams, J. L. (1986) *The Care and Feeding of Ideas*. Stanford, C A: Stanford Alumni Association.

Amabile, T. (1988) "A Model of Creativity and Innovation in Organizations." *Research in Organizational Behavior*, 10, 123-167.

- Baron, J. (1988) *Thinking and Deciding*. Cambridge, Cambridge University Press.
- Barron, F. (1965) "The Psychology of Creativity." In T. M. Newcomb (ed.) *New Directions in Psychology*, Vol. 2, pp. 1-134. New York: Holt.
- Bartlett, S. (1978) "Protocol Analysis in Creative Problem Solving," *Journal of Creative Behavior*, 12, 181-191.
- Bazerman, M. H. (1994) *Judgment in Managerial Decision Making*, 3rd ed. New York: Wiley.
- Campbell, D. T. (1960) "Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes," *Psychological Review*, 67, 380-400.
- Davis, G. (1986) *Creativity Is Forever*, 2nd ed. Dubuque, IA: Kendall/Hunt.
- Delbecq, A. L., A. H. Van de Ven, and D. H. Gustafson (1975) *Group Techniques for Program Planning*. Glenview, IL: Scott, Foresman.
- De Mille, R. (1976) *Put Your Mother on the Ceiling*. Middlesex, England: Penguin.
- Dole, S. H., H. G. Campbell, D. Dreyfuss, W. D. Gosch, E. D. Harris, D. E. Lewis, T. M. Parker, J. W. Ranftl, and J. String, Jr. (1968a) Methodologies for Analyzing the Comparative Effectiveness and Costs of Alternate Space Plans. *RM-5656-NASA*, Vol. 1 (Summary). Santa Monica, CA: The Rand Corporation.
- Dole, S. H., H. G. Campbell, D. Dreyfuss, W. D. Gosch, E. D. Harris, D. E. Lewis, T. M. Parker, J. W. Ranftl, and J. String, Jr. (1968b) Methodologies for Analyzing the Comparative Effectiveness and Costs of Alternate Space Plans. *RM-5656-NASA*, Vol. 2. Santa Monica, CA: The Rand Corporation.
- Gordon, W. J. J. (1961) *Synectics*. New York: Harper & Row.
- Gordon, W. J. J. (1969) *The Metaphorical Way of Learning and Knowing*. Cambridge, MA: SES Associates.
- Hogarth, R. (1987) *Judgement and Choice*, 2nd ed. Chichester, England: Wiley.
- Howard, R. A. (1988) "Decision Analysis: Practice and Promise," *Management Science*, 34, 679-695.
- Johnson, D. M., G. L. Parrott, and R. P. Stratton (1968) "Production and Judgment of Solutions to Five Problems," *Journal of Educational Psychology Monograph Supplement*, 59, 1-21.
- Keeney, R. (1992) *Value-Focused Thinking*. Cambridge, MA: Harvard University Press.
- Keeney, R., and H. Raiffa (1976) *Decisions with Multiple Objectives*. New York: Wiley.
- Keller, L. R., and J. L. Ho (1989) "Decision Problem Structuring: Generating Options," *IEEE Transactions on Systems, Man, and Cybernetics*, 18, 715—728.
- Kleindorfer, P. R., H. C. Kunreuther, and P. J. H. Schoemaker, (1993) *Decision Sciences: An Integrated Perspective*. Cambridge, Cambridge University Press.
- Koberg, D., and J. Bagnall (1974) *The Universal Traveler*. Los Altos, CA: William Kaufmann.
- Kris, E. (1952) "On Preconscious Mental Processes." In *Psychoanalytic Explorations in Art*. New York: International Universities Press.
- Kubie, L. S. (1958) *Neurotic Distortion of the Creative Process*. Lawrence: University of Kansas Press.
- Maltzman, I. (1960) "On the Training of Originality." *Psychological Review*, 67, 229-242.

- Maslow, A. (1954) *Motivation and Personality*. New York: Harper & Row.
- Mednick, S. A. (1962) "The Associative Basis of the Creative Process." *Psychological Review*, 69, 220-232.
- Osborn, A. F. (1963) *Applied Imagination*, 3rd ed. New York: Scribner.
- Peters, T. (1988) *Thriving on Chaos*. New York: Knopf.
- Rugg, H. (1963) *Imagination: An Inquiry into the Sources and Conditions that Stimulate Creativity*. New York: Harper & Row.
- Skinner, B. F. (1972) *Cumulative Record: A Selection of Papers*, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall.
- Staats, A. W. (1968) *Learning, Language, and Cognition*. New York: Holt.
- Wallas, G. (1926) *The Art of Thought*. New York: Harcourt, Brace.
- Winkelgren, W. A. (1974) *How to Solve Problems*. San Francisco: Freeman.

EPILOGUE

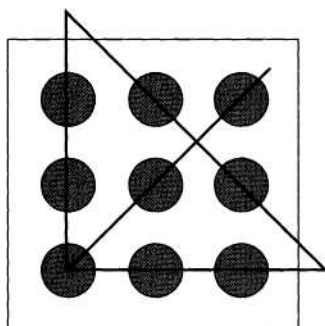
The basic solution to the nine-dot puzzle is shown in Figure 6.6. Many people tacitly assume that the lines may not go beyond the square that is implied by the dots, and so they fail to solve the puzzle. Figure 6.7 shows how to connect nine fat dots with three straight lines, removing the block that the line has to go through the centers of the dots.

It is possible, with enough effort to remove the necessary blocks, to connect the dots with one line. Some solutions include:


- Fold the paper in a clever way so that the dots line up in a row. Then just draw one straight line through all nine dots.
- Roll the paper up and tape it so that you can draw a spiral through all of the dots.
- Cut the paper in strips and tape it together so that the dots are in one row.
- Draw large dots, wad the paper up and stab it with a pencil. Unfold it and see if the pencil went through all the dots. If not, try again. "Everybody wins!"

Figure 6.6

The standard solution to the nine-dot puzzle.



**Figure
6.7**
*An unblocked three-
line solution.*



- Fold the paper carefully so that the dots are lying one on top of the other. Now stab your pencil through the nine layers of paper. It helps to sharpen your pencil first.
- Draw very small dots very close together in the square pattern and then draw a fat line through all of them at once. This one is courtesy of Becky Buechel when she was 10 years old.

Source: Adams (1979).

SECTION 2

Modeling Uncertainty

As we have seen, uncertainty is a critical element of many decisions that we face. In the next five chapters we will consider a variety of ways to model uncertainty in decision problems by using probability. We begin with a brief introduction to probability in Chapter 7. This introduction has three objectives: to remind you of probability basics, to show some ways that probability modeling can be useful in decision problems, and to give you a chance to polish your ability to manipulate probabilities. The problems and cases at the end of Chapter 7 are recommended especially to help you accomplish the last goal.

Chapter 8 introduces the topic of subjective probability. In the introduction to this text the claim was made that subjective judgements are critical in a decision-analysis approach. Most of us are comfortable making informal statements that reflect our uncertainty. For example, we use terms such as "there's a chance that such-and-such will happen." In a decision-analysis approach, however, there is a need for more precision. We can use probability to model subjective beliefs about uncertainty, and chapter 8 presents techniques to do this. Thus, our beliefs and feelings about uncertainty can be translated into probability numbers that can be included in a decision tree or influence diagram.

In many cases the uncertainty that we face has characteristics that make it similar to certain prototypical situations. In these cases, it may be possible to represent

the uncertainty with a standard mathematical model and then derive probabilities on the basis of the mathematical model. Chapter 9 presents a variety of theoretical probability models that are useful for representing uncertainty in many standard situations.

Chapter 10 discusses the use of historical data as a basis for developing probability distributions. If data about uncertainty in a decision situation are available, a decision maker would surely want to use them. First, we see how to use data alone to create histograms and continuous distributions. Second, we discuss the use of data to model relationships among variables. Finally, we develop an approach that uses data to update a decision maker's probability beliefs via Bayes' theorem.

It is also possible to "create data" through computer simulation, or by what is known as Monte Carlo simulation. That is, one can construct a model of a complex decision situation and use a computer to simulate the situation many times. By tracking the outcomes, the decision maker can obtain a fair idea of the probabilities associated with various outcomes. Monte Carlo simulation techniques are discussed in Chapter 11.

When faced with uncertainty most decision makers do their best to reduce it. The basic strategy that we follow is to collect information. In this day and age of extensive telecommunications and computer data bases, information is anything but scarce, but determining what information is appropriate and then processing it can be costly. How much is information worth to you? Moreover, in a problem with many sources of uncertainty, calculating the value of information can help to guide the decision analysis, thus indicating where the decision maker can best expend resources to reduce uncertainty. Chapter 12 explores the value of information within the decision-analysis framework.

Probability Basics

One of the central principles of decision analysis is that we can represent uncertainty of any kind through the appropriate use of probability. We already have been using probability in straightforward ways to model uncertainty in decision trees and influence diagrams. This chapter presents some of the basic principles for working with probability and probability models. Our objective is to be able to create and analyze a model that represents the uncertainty faced in a decision. The nature of the model created naturally depends on the nature of the uncertainty faced, and the analysis required depends on the exigencies of the decision situation.

We have used the term *chance event* to refer to something about which a decision maker is uncertain. In turn, a chance event has more than one possible *outcome*. When we talk about probabilities, we are concerned with the chances associated with the different possible outcomes. For convenience, we will refer to chance events with boldface letters (e.g., Chance Events **A** and **B**), and to outcomes with lightface letters (e.g., Outcomes A_1 , B_j , or C). Thus Chance Event **B**, for example, might represent a particular chance node in an influence diagram or decision tree, and Outcomes B_1 , B_2 , and B_3 would represent **B**'s possible outcomes.*

After reading the chapter and working the problems and cases, you should be (1) reasonably comfortable with probability concepts, (2) comfortable in the use of probability to model simple uncertain situations, (3) able to interpret probability statements in terms of the uncertainty that they represent, and (4) able to manipulate and analyze the models you create.

*The terminology we adopt in this book is slightly unconventional. Most authors define an *outcome space* that includes all possible elemental outcomes that may occur. For example, if the uncertain event is how many orders arrive at a mail-order business in a given day, the outcome space is composed of the integers 0, 1, 2, 3, . . . , and each integer is a possible elemental outcome (or simply an outcome). An *event* is then defined as a set of possible outcomes. Thus, we might speak of the event that the number of orders equals zero or the event that more than 10 orders arrive in a day.

In this book we use the term *outcome* to refer to what can occur as the result of an uncertain event. Such occurrences, which we represent as branches from chance nodes in decision trees, can be either events or elemental outcomes in the conventional terminology. Thus, our usage of *outcome* includes the conventional *event* as well as *outcome*. Why the change? For our purposes it is not necessary to distinguish between elemental outcomes and sets of outcomes. Also, we avoid the potential confusion that can arise by using *uncertain event* to refer to a process and *event* to refer to a result of that process.

A Little Probability Theory

Probabilities must satisfy the following three requirements.

1 Probabilities Must Lie Between 0 and 1 Every probability (p) must be positive, and between 0 and 1, inclusive ($0 < p < 1$). This is a sensible requirement. In informal terms it simply means nothing can have more than a 100% chance of occurring or less than a 0% chance.

2 Probabilities Must Add Up Suppose two outcomes are mutually exclusive (only one can happen, not both). The probability that one or the other occurs is then the sum of the individual probabilities. Mathematically, we write $F(A_1 \text{ or } A_2) = P(A_1) + P(A_2)$ if A_1 and A_2 cannot both happen. For example, consider the stock market. Suppose there is a 30% chance that the market will go up and a 45% chance that it will stay the same (as measured by the Dow Jones average). It cannot do both at once, and so the probability that it will either go up or stay the same must be 75%.

3 Total Probability Must Equal 1 Suppose a set of outcomes is mutually exclusive and collectively exhaustive. This means that one (and only one) of the possible outcomes must occur. The probabilities for this set must sum to 1. Informally, if we have a set of outcomes such that one of them has to occur, then there is a 100% chance that one of them will indeed come to pass.


We have seen this in decision trees; the branches emanating from a chance node must be such that one and only one of the branches occurs, and the probabilities for all the branches must add to 1. Consider the stock market example again. If we say that the market can go up, down, or stay the same, then one of these three outcomes must happen. The probabilities for these outcomes must sum to 1—that is, there is a 100% chance that one of them will occur.

Venn Diagrams

Venn diagrams provide a graphic interpretation of probability. Figure 7.1 shows a simple Venn diagram in which two Outcomes, A_1 and A_2 , are displayed. Think of the diagram as a whole representing a chance event $(A)_x$ and areas represent possible outcomes. The circle labeled A_1 thus represents outcome A_1 . Because the areas of A_1 and A_2 do not overlap, A_1 and A_2 cannot both occur at the same time; they are mutually exclusive.

We can use Figure 7.1 to interpret the three requirements of probability mentioned above. The first requirement is that a probability must lie between 0 and 1. Certainly an outcome cannot be represented by a negative area. Furthermore, an outcome cannot be represented by an area larger than the entire rectangle. For the second requirement, we see that A_1 and A_2 are mutually exclusive because they do not overlap. Thus, the prob-

Figure 7.1
A Venn diagram.
Outcomes A_1 and A_2
are mutually exclusive
events.



ability of A_1 or A_2 occurring must be just the sum of the probability of A_1 plus the probability of A_2 . For the third requirement, label the shaded portion of the rectangle as C . This is what must happen if neither A_1 nor A_2 happens. Because A_1 , A_2 , and C together make up the whole rectangle, then one of the three must occur. Moreover, only one of them can occur. The upshot is that their probabilities must sum to 1. Alternatively, there is a 100% chance that A_1 , A_2 , or C will occur.

More Probability Formulas


The following definitions and formulas will make it possible to use probabilities in a wide variety of decision situations.

4 Conditional Probability Suppose that you are interested in whether a particular stock price will increase. You might use a probability $P(\text{Up})$ to represent your uncertainty. If you find out that the Dow Jones industrial average rose, however, you would want to base your probability on this condition. Figure 7.2 represents the situation with a Venn diagram. Now the entire rectangle actually represents two chance events, what happens with the Dow Jones index and what happens with the stock price. For the Dow Jones event, one possible outcome is that the index goes up, represented by the circle. Likewise, the oval represents the possible outcome that the stock price goes up. Because these two outcomes can happen at once, the two areas overlap in the diagonally shaded area. When both outcomes occur, this is called the *joint outcome* or *intersection*. It also is possible for one to rise while the other does not, which is represented in the diagram by the nonoverlapping portions of the "Dow Jones Up" and "Stock Price Up" areas. And, of course, the gray area surrounding the circle and oval represents the joint outcome that neither the Dow Jones index nor the stock price goes up.

Once we know that the Dow Jones has risen, then the entire rectangle is no longer appropriate. At this point, we can restrict our attention to the "Dow Jones Up" circle. We want to know what the probability is that the stock price will increase given that the Dow Jones average is up, and so we are interested in the probability associated with the area "Stock Price Up" in the restricted space.

Given that we are looking at the restricted space, the conditional probability of "Stock Price Up given Dow Jones Up" would be represented by the proportion of "Dow Jones Up" in the original diagram that is the joint outcome "Stock Price Up" and "Dow Jones Up" (the diagonally shaded area). This intuitive approach leads to

Figure 7.2
A Venn-diagram
representation of
conditional probability.



the conclusion that $P(\text{Stock Price Up given Dow Jones Up}) = \frac{P(\text{Stock Price Up and Dow Jones Up})}{P(\text{Dow Jones Up})}$.

Mathematically, we write $P(A | B)$ to represent the conditional probability of A given that B occurs. Read it as "Probability of A given B." The definition is

$$P(A | B) = \frac{P(A \text{ and } B)}{P(B)}$$

Informally, we are looking only at the occasions when Outcome B occurs, and $P(A | B)$ is the proportion of those times that Outcome A also occurs. The probability $P(A \text{ and } B)$ is often called a *joint* probability.

5 Independence The definition of probabilistic independence is as follows:

Chance Events **A** (with Outcomes A_1, \dots, A_n) and **B** (with Outcomes B_1, \dots, B_m) are *independent* if and only if

$$P(A_i | B_j) = P(A_i) \text{ for all possible Outcomes } A_i \text{ and } B_j.$$

In words, knowing which of **B**'s outcomes occurred will not help you find probabilities for **A**'s outcomes. Furthermore, if Chance Events **A** and **B** are independent, then we can write

$$P(A_i) = P(A_i | B_j) = \frac{P(A_i \text{ and } B_j)}{P(B_j)}$$

From this we can see that $P(A_i \text{ and } B_j) = P(A_i) P(B_j)$. Thus, when two events are independent, we can find the probability of a joint outcome by multiplying the probabilities of the individual outcomes.

Independence between two chance events is shown in influence diagrams by the absence of an arrow between chance nodes. This is fully consistent with the definitions given in Chapter 3. If one event is not relevant in determining the chances associated with another event, there is no arrow between the chance nodes. An arrow from chance node B to chance node A would mean that the probabilities associated with A are conditional probabilities that depend on the outcome of B.

As an example of independent chance events, consider the probability of the Dow Jones index increasing and the probability of the Dow Jones increasing, *given* that it rains tomorrow. It seems reasonable to conclude that

$$P(\text{Dow Jones Up}) = P(\text{Dow Jones Up} | \text{Rain})$$

because knowing about rain does not help to assess the chance of the index going up.

Independent chance events are not to be confused with mutually exclusive outcomes. Two outcomes are mutually exclusive if only one can happen at a time. Clearly, however, independent chance events can have outcomes that happen together. For example, it is perfectly possible for the Dow Jones to increase and for rain to occur tomorrow. Rain and No Rain would constitute mutually exclusive outcomes; only one occurs at a time.

Finally, if two chance events are probabilistically dependent, this does *not* imply a causal relationship. As an example, consider economic indicators. If a leading economic indicator goes up in one quarter, then it is unlikely that a recession will occur in the next quarter; the change in the indicator and the occurrence of a recession are dependent events. But this is not to say that the indicator going up causes the recession not to happen, or vice versa. In some cases there may be a causal chain linking the events, but it may be a very convoluted one. In general, dependence does not imply causality.

Conditional Independence. This is an extension of the idea of independence. Conditional independence is best demonstrated with an influence diagram. In Figure 7.3, Events A and B are conditionally independent given Event C. Note that C can be a chance event or a decision, as in Figure 7.3b. The only connection between A and B goes through C; there is no arrow directly from A to B, or vice versa. Mathematically, we would write


Events **A** and **B** are conditionally independent given **C** if and only if

$$P(A_i | B_j, C_k) = P(A_i | C_k), \text{ for all}$$

possible Outcomes $A_i, B_j,$ and $C_k.$

In words, suppose we are interested in Event **A**. If **A** and **B** are conditionally independent given **C**, then learning the outcome of **B** adds no new information regarding **A** if

Figure 7.3



Conditional independence in an influence diagram. In these influence diagrams, A and B are conditionally independent given C. As shown, the conditioning event can be either (a) a chance event or (b) a decision.

the outcome of C already is known. Alternatively, conditional independence means that

$$P(A_i \text{ and } B_j | C_k) = P(A_i | C_k) P(B_j | C_k)$$

Conditional independence is the same as normal (unconditional) independence except that every probability has the same conditions to the right side of the vertical bar. When constructing influence diagrams, identification of conditional independence can ease the burden of finding probabilities for the chance events.

As an example of conditional independence, consider a situation in which you wonder whether or not a particular firm will introduce a new product. You are acquainted with the CEO of the firm and may be able to ask about the new product. But one of your suppliers, who has chatted with the CEO's assistant, reports that the company is on the brink of announcing the new product. Your probability that the company will indeed introduce the product thus would change:

$$P(\text{Introduce Product} | \text{Supplier's Report}) \neq P(\text{Introduce Product})$$

Thus, these two events are not independent when considered by themselves. Consider, however, the information from the CEO. Given that information, the supplier's report might not change your probability:

$$\begin{aligned} &P(\text{Introduce Product} | \text{Supplier's Report, Information from CEO}) \\ &= P(\text{Introduce Product} | \text{Information from CEO}) \end{aligned}$$

Thus, given the information from the CEO, the supplier's report and the event of the product introduction are conditionally independent.

Conditional independence is an important concept when thinking about causal effects and dependence. As mentioned above, dependence between two events does not necessarily mean that one causes the other. For example, more drownings tend to be associated with more ice cream consumption, but it seems unlikely that these two are causally related. The explanation lies in a common cause; both tend to happen during summer, and it might be reasonable to assume that drownings and ice cream consumption are conditionally independent given the season. Another example is that high skirt hemlines tend to occur when the stock market is advancing steadily (a

bull market), although no one believes that headlines cause the stock market to advance, or vice versa. Perhaps both reflect a pervasive feeling of confidence and adventure in our society. If we had an adequate index of "societal adventuresomeness," headline height and stock market activity might well be conditionally independent given the index.

6 Complements Let \bar{B} ("Not B" or "B-bar") represent the outcome that is the complement of B. This means that if B does not occur, then \bar{B} must occur. Because probabilities must add to 1 (requirement 3 above),

$$P(\bar{B}) = 1 - P(B)$$

The Venn diagram in Figure 7.4 demonstrates complements. If the area labeled B represents Outcome B, then everything outside of the oval must represent what happens if B does not happen. For another example, the lightly shaded area in Figure 7.2 represents the complement of the Outcome "Dow Jones Up or Stock Price Up," the union of the two individual outcomes.

7 Total Probability of an Event A convenient way to calculate $P(A)$ is with this formula:

$$P(A) = P(A \text{ and } B) + P(A \text{ and } \bar{B})$$

$$= P(A | B) P(B) + P(A | \bar{B}) P(\bar{B})$$

To understand this formula, examine Figure 7.5. Clearly, Outcome A is composed of those occasions when A and B occur and when A and \bar{B} occur. Because the joint Outcomes "A and B" and "A and \bar{B} " are mutually exclusive, the probability of A must be the sum of the probability of "A and B" plus the probability of "A and \bar{B} ." As an example, suppose we want to assess the probability that a stock price will increase. We could use its relationship with the Dow Jones index (Figure 7.2) to help make the assessment:

$$P(\text{Stock Price Up}) = P(\text{Stock Price Up} | \text{Dow Jones Up}) \times P(\text{Dow Jones Up}) + P(\text{Stock Price Up} | \text{Dow Jones Not Up}) \times P(\text{Dow Jones Not Up})$$

Figure 7.4
Venn diagram illustrating the idea of an outcome's complement.





Figure 7.5

Total probability.
Outcome A is made up
of Outcomes "A and
B" and "A and B̄."



Although it may appear that we have complicated matters by requiring three probabilities instead of one, it may be quite easy to think about (1) the probability of the stock price movement conditional on the change in the Dow Jones index and (2) the probabilities associated with changes in the index.

8 Bayes' Theorem Because of the symmetry of the definition of conditional probability, we can write

$$P(B | A) P(A) = P(A | B) P(B)$$

$$P(B | A) = \frac{P(A | B) P(B)}{P(A)}$$

Now expanding $P(A)$ with the formula for total probability, we obtain

from which we can derive

$$P(B | A) = \frac{P(A | B)P(B)}{P(A | B) P(B) + P(A | \bar{B}) P(\bar{B})}$$

This formula often is referred to as *Bayes' theorem*. It is extremely useful in decision analysis, especially when using information. We will not bother with an example showing its application here, but we will see numerous applications of Bayes' theorem in the examples at the end of this chapter, as well as in Chapters 9, 10, and 12.

Uncertain Quantities

Many uncertain events have quantitative outcomes. For example, we already have mentioned stock prices and the level of the Dow Jones index. Another example would be tomorrow's maximum temperature. If an event is not quantitative in the first place, we might define a variable that has a quantitative outcome based on the

original event. For example, we might be concerned about precipitation. If we consider the amount of precipitation (in centimeters) that falls tomorrow, this uncertain event is quantitative. Moreover, we could define an uncertain quantity X : Let $X = 1$ if precipitation occurs, and $X = 0$ if not.

The set of probabilities associated with all possible outcomes of an uncertain quantity is called its *probability distribution*. For example, consider the probability distribution for the number of raisins in an oatmeal cookie, which we could denote by Y . We might have $P(Y = 0) = 0.02$, $P(Y = 1) = 0.05$, $P(Y = 2) = 0.20$, $P(Y = 3) = 0.40$, and so on. Of course, the probabilities in a probability distribution must add to 1 because the events — numerical outcomes — are mutually exclusive. Uncertain quantities (often called *random variables*) and their probability distributions play a central role in decision analysis.

In general, we will use capital letters to represent uncertain quantities. Thus, we will write $P(X = 3)$ or $P(Y > 0)$, for example, which are read as "the probability that the uncertain quantity X equals 3," and "the probability that the uncertain quantity Y is greater than 0." Occasionally we will need to use a more general form. Lowercase letters will denote outcomes or realizations of an uncertain quantity. An example would be $P(X = x)$, where capital X denotes the uncertain quantity itself and lowercase x represents the actual outcome.

In general, it is helpful to distinguish between *discrete* and *continuous* uncertain quantities. In the next section we will describe in detail discrete uncertain quantities, their probability distributions, and certain characteristics of those distributions. Then we will turn to some continuous quantities and show how their probability distributions and their characteristics are analogous to the discrete case.

Discrete Probability Distributions

The discrete probability distribution case is characterized by an uncertain quantity that can assume a finite or countable number of possible values. We already have seen two examples. The first was the precipitation example; we defined a discrete uncertain quantity that could take only the values 0 and 1. The other example was the number of raisins in an oatmeal cookie. Other examples might be the number of operations a computer performs in any given second or the number of games that will be won by the Chicago Cubs next year. Strictly speaking, future stock prices quoted on the New York Stock Exchange are discrete uncertain quantities because they can take only values that are in eighths: $10^{5/8}$, $11^{3/4}$, or $12^{1/2}$, for example.

When we specify a probability distribution for a discrete uncertain quantity, we can express the distribution in several ways. Two approaches are particularly useful. The first is to give the *probability mass function*. This function lists the probabilities for each possible discrete outcome. For example, suppose that you think that no cookie in a batch of oatmeal cookies could have more than five raisins. A possible probability mass function would be:

$$\begin{aligned}
 P(Y = 0 \text{ raisins}) &= 0.02 \\
 P(Y = 3 \text{ raisins}) &= 0.40 \quad P(Y = 1 \text{ raisin}) = 0.05 \\
 P(Y = 4 \text{ raisins}) &= 0.22 \quad P(Y = 2 \text{ raisins}) = 0.20 \quad P(Y = 5 \text{ raisins}) = 0.11
 \end{aligned}$$

This mass function can be displayed in graphical form (Figure 7.6). Such a graph often is called a *histogram*.

The second way to express a probability distribution is as a *cumulative distribution function* (CDF). A cumulative distribution gives the probability that an uncertain quantity is less than or equal to a specific value: $P(X \leq x)$. For our example, the CDF is given by

$$\begin{aligned} P(Y \leq 0 \text{ raisins}) &= 0.02 & P(Y \leq 3 \text{ raisins}) &= 0.67 \\ P(Y \leq 1 \text{ raisins}) &= 0.07 & P(Y \leq 4 \text{ raisins}) &= 0.89 \\ P(Y \leq 2 \text{ raisins}) &= 0.27 & P(Y \leq 5 \text{ raisins}) &= 1.00 \end{aligned}$$

Cumulative probabilities can be graphed; the CDF for the oatmeal cookie is graphed in Figure 7.7.

Note that the graph actually covers all the points along the horizontal axis. That is, we can read from the graph not only $P(Y \leq 3)$, but also $P(Y \leq 4.67)$, for example. In fact, $P(Y \leq 4.67) = P(Y \leq 4) = 0.89$, because it is not possible for a cookie to have a fractional number of raisins (assuming whole raisins, of course).

You may recognize this idea of a probability mass function and a CDF. When we constructed risk profiles and cumulative risk profiles in Chapter 4, we were working with these two representations of probability distributions.

Figure 7.6
A probability mass function displayed as a histogram.





Figure 7.7
Cumulative distribution function (CDF) for number of raisins in an oatmeal cookie.



Expected Value

A discrete uncertain quantity's *expected value* is the probability-weighted average of its possible values. That is, if X can take on any value in the set $\{x_1, x_2, \dots, x_n\}$, then the expected value of X is simply the sum of x_1 through x_n , each weighted by the probability of its occurrence. Mathematically,

$$\begin{aligned} \text{Expected value of } X &= x_1P(X = x_1) + x_2P(X = x_2) + \dots + x_nP(X = x_n) \\ &= \sum_{i=1}^n x_iP(X = x_i). \end{aligned}$$

The expected value of X also is referred to as the average or mean of X and is denoted by $E(X)$ or occasionally μ_x (Greek mu).

The expected value can be thought of as the "best guess" for the value of an uncertain quantity or random variable. If it were possible to observe many outcomes of the random variable, the average of those outcomes would be very close to the expected value. We already have encountered expected values in calculating EMVs to solve influence diagrams and decision trees. The expected monetary value is the expected value of a random variable that happens to be the monetary outcome in a decision situation.

Suppose that X is used to calculate some other quantity, say Y . Then it is possible to talk about the expected value of Y , or the expected value of this function of X :

$$\begin{aligned} \text{If} \\ & Y = f(X) \\ \text{then} \\ & E(Y) = E[f(X)] \end{aligned}$$

A particularly useful function is a linear function of X , one in which X is multiplied by a constant and has a constant added to it. If Y is a linear function of X , then $E(Y)$ is particularly easy to find:

$$\begin{aligned} \text{If} \\ & Y = a + bX \\ \text{then} \\ & E(Y) = a + bE(X) \end{aligned}$$

That is, plug $E(X)$ into the linear formula to get $E(Y)$. Unfortunately, this does not hold for nonlinear functions (log, square root, and so on); in general, it applies only

to linear ones. That is, suppose you have some function $f(X)$ and you want to find the expected value of $E[f(X)]$. In general, you cannot plug $E(X)$ into the function and get the correct answer; $E[f(X)] \neq f[E(X)]$, unless $f(X)$ is a linear function like $a + bX$.

To go one step further, suppose we have several uncertain quantities, X_1, \dots, X_k , and we add them to get uncertain quantity Y . Then the expected value of Y is the sum of the expected values:

<p>If</p> $Y = X_1 + X_2 + \dots + X_k$ <p>then</p> $E(Y) = E(X_1) + E(X_2) + \dots + E(X_k)$

For instance, if we know the expected amount of precipitation for each of the next seven days, the expected amount of precipitation for the entire week is simply the sum of the seven daily expected values.

Variance and Standard Deviation

Another useful measure of a probability distribution is the *variance*. The variance of uncertain quantity X is denoted by $\text{Var}(X)$ or σ_X^2 (Greek sigma) and is calculated mathematically by

$\begin{aligned} \text{Var}(X) &= [x_1 - E(X)]^2 P(X = x_1) + [x_2 - E(X)]^2 P(X = x_2) \\ &\quad + \dots + [x_n - E(X)]^2 P(X = x_n) \\ &= \sum_{i=1}^n [x_i - E(X)]^2 P(X = x_i) \\ &= E[X - E(X)]^2 \end{aligned}$

In words, calculate the difference between the expected value and x_i and square that difference. Do this for each possible x_i . Now find the expected value of these squared differences.

As with expected values, we can find variances of functions for X . In particular, the variance of a linear function of X is easily found:

<p>If</p> $Y = a + bX$ <p>then</p> $\text{Var}(Y) = b^2 \text{Var}(X)$
--

For example, suppose that a firm will sell an uncertain number (X) of units of a product. The expected value of X is 1000 units, and the variance is 400. If the price is \$3 per unit, then the revenue (Y) is equal to $3X$. Because this is a linear function of X , $E(Y) = 3E(X) = \$3000$, and $\text{Var}(Y) = 3^2 \text{Var}(X) = 9(400) = 3600$.

We also can talk about the variance of a sum of independent uncertain quantities. That is, *as long as the uncertain quantities are probabilistically independent* — the probability distribution for one does not depend on the others — then the variance of the sum is the sum of the variances:

If

$$Y = X_1 + X_2 + \dots + X_k$$

then

$$\text{Var}(Y) = \text{Var}(X_1) + \text{Var}(X_2) + \dots + \text{Var}(X_k)$$

So, if our firm sells two products, one for \$3 and another for \$5, and the variance for these two products is 400 and 750, respectively, then the variance of the firm's revenue is $\text{Var}(Y) = 3^2 \text{Var}(X_1) + 5^2 \text{Var}(X_2) = 9(400) + 25(750) = 22,350$.

The *standard deviation* of X , denoted by σ_x , is just the square root of the variance. Because the variance is the expected value of the squared differences, the standard deviation can be thought of as a "best guess" as to how far the outcome of X might lie from $E(X)$. A large standard deviation and variance means that the probability distribution is quite spread out; a large difference between the outcome and the expected value is anticipated. For this reason, the variance and standard deviation of a probability distribution are used as measures of variability. A large variance or standard deviation would indicate a situation in which the outcomes are highly variable.

To illustrate the ideas of expected value, variance, and standard deviation, consider the double-risk dilemma depicted in Figure 7.8. Choices A and B both lead to uncertain dollar outcomes. Given Decision A, for example, there are three possible profit outcomes having probability mass function $P(\text{Profit} = \$20 \mid A) = 0.24$, $P(\text{Profit} = \$35 \mid A) = 0.47$, and $P(\text{Profit} = \$50 \mid A) = 0.29$. Likewise, for B we have $P(\text{Profit} = -\$9 \mid B) = 0.25$, $P(\text{Profit} = \$0 \mid B) = 0.35$, and $P(\text{Profit} = \$95 \mid B) = 0.40$. Now we can calculate the expected profits conditional on choosing A or B:

$$E(\text{Profit} \mid A) = 0.24(\$20) + 0.47(\$35) + 0.29(\$50) = \$35.75$$

$$E(\text{Profit} \mid B) = 0.25(-\$9) + 0.35(\$0) + 0.40(\$95) = \$35.75$$

These two uncertain quantities have exactly the same expected values. (Note that the expected profit does not have to be one of the possible outcomes!)

We also can calculate the variances and standard deviations (σ) for A and B:

$$\begin{aligned}\text{Var}(\text{Profit} \mid A) &= (20 - 35.75)^2(0.24) + (35 - 35.75)^2(0.47) \\ &\quad + (50 - 35.75)^2(0.29) \\ &= 118.69 \text{ "dollars squared"} \\ \sigma_A &= \sqrt{118.69} = \$10.89 \\ \text{Var}(\text{Profit} \mid B) &= (-9 - 35.75)^2(0.25) + (0 - 35.75)^2(0.35) \\ &\quad + (95 - 35.75)^2(0.40) \\ &= 2352.19 \text{ "dollars squared"} \\ \sigma_B &= \sqrt{2352.19} = \$48.50\end{aligned}$$

The variance and standard deviation of B are much larger than those for A. The outcomes in B are more spread out or more variable than the outcomes for A, which are clustered fairly closely around the expected value.


The example also points out the fact that variance, being a weighted sum of squared terms, is expressed in squared units. In this case, the variance is in "dollars squared" because the original outcomes are in dollars. Taking the square root to find the standard deviation brings us back to the original units, or, in this case, dollars. For this reason, the standard deviation is interpreted more easily as a measure of variability.

You might have noticed that, while the variance and standard deviation can be used to gauge the riskiness of an option, the cumulative probabilities also can be useful in this respect. For example, we have $P(\text{Profit} \leq 0 \mid A) = 0$, but $P(\text{Profit} \leq 0 \mid B) = 0.6$, and B thus looks somewhat riskier than A. At the other extreme, of course, B looks better. Project A cannot produce a profit greater than \$50: $P(\text{Profit} \leq \$50 \mid A) = 1.00$. For B, however, $P(\text{Profit} \leq \$50 \mid B) = 0.60$.

Covariance and Correlation for Measuring Dependence (Optional)

We have discussed the notion of probabilistic dependence above and indicated that dependence is defined in terms of conditional distributions. In some cases, though, the use of conditional distributions can be difficult, and another approach to measuring

Figure 7.8
A choice between two uncertain prospects.



dependence is worthwhile. *Covariance* is a quantity that is closely related to the idea of variance. Covariance and its close relative *correlation* can be used to measure certain kinds of dependence.

The covariance between two uncertain quantities X and Y is calculated mathematically by

$$\begin{aligned}
 \text{Cov}(X, Y) &= [x_1 - E(X)][y_1 - E(Y)]P(X = x_1 \text{ and } Y = y_1) \\
 &\quad + [x_1 - E(X)][y_2 - E(Y)]P(X = x_1 \text{ and } Y = y_2) \\
 &\quad + \cdots + [x_1 - E(X)][y_m - E(Y)]P(X = x_1 \text{ and } Y = y_m) \\
 &\quad + [x_2 - E(X)][y_1 - E(Y)]P(X = x_2 \text{ and } Y = y_1) \\
 &\quad + \cdots + [x_2 - E(X)][y_m - E(Y)]P(X = x_2 \text{ and } Y = y_m) \\
 &\quad + \cdots + [x_n - E(X)][y_1 - E(Y)]P(X = x_n \text{ and } Y = y_1) \\
 &\quad + \cdots + [x_n - E(X)][y_m - E(Y)]P(X = x_n \text{ and } Y = y_m) \\
 &= \sum_{i=1}^n \{ [x_i - E(X)][y_1 - E(Y)]P(X = x_i \text{ and } Y = y_1) \\
 &\quad + \cdots + [x_i - E(X)][y_m - E(Y)]P(X = x_i \text{ and } Y = y_m) \} \\
 &= \sum_{i=1}^n \sum_{j=1}^m [x_i - E(X)][y_j - E(Y)]P(X = x_i \text{ and } Y = y_j) \\
 &= E[(X - E(X))(Y - E(Y))]
 \end{aligned}$$

Although this is a complicated formula, with a little interpretation we can get some insight into it. First, it really is similar to the formula for variance. There we calculated an average of squared deviations of X from its expected value $E(X)$. Here, instead of squaring the deviations, we multiply X 's deviation times Y 's deviation. This is sometimes called a *cross product* because it is a product of deviations for two different quantities (X and Y).

Although it may not be evident on the surface, the covariance can be either positive or negative. Suppose that large values of X tend to occur with large values of Y , and small with small. On the other hand, the probability that high values of X and low values of Y (and vice versa) occur together is low. Now consider what happens with the cross products in the formula. When both X and Y are high, they will both be above their corresponding expected values, making both deviations and their cross product positive. When both quantities are low, they will both be below their expected values. Both deviations will be negative, but the cross product will again be positive. When one is high and one is low (that is, one above its expected value and one below), the cross product will be negative. If large X 's tend to go with large Y 's, and small with small, then the positive cross products will get more weight (higher

probabilities) than negative cross products in the formula, and the overall calculation will yield a positive covariance. On the other hand, if large X 's tend to occur with small Y 's, and vice versa, then the negative cross products will get more weight, resulting in a negative covariance. Thus, a positive covariance reflects quantities that tend to move in the same direction, and this is called a *positive* relationship. Likewise, a negative covariance indicates that the quantities tend to move in opposite directions, which is called a *negative* relationship.

A simple example will help clarify the idea of covariance and its calculation. Suppose an investor is considering purchasing shares in American Rivets Corporation (ARC). The investor already has shares of Sundance Solar Power (SSP). One of the things the investor would like to accomplish is to stabilize the rate of return of the portfolio; ideally, when the return on one stock goes down, the other would go up. On the other hand, a positive relationship would be bad because the returns would tend to go up and down together, making the overall return on the portfolio vary considerably.

A simple model of the returns on the two stocks involves only two possible outcomes for each one. ARC could have returns of 10% or -5%. SSP, on the other hand, could have returns of 12% or -8%. The probabilities of the possible outcomes are

$$P(\text{ARC} = 10\% \text{ and } \text{SSP} = 12\%) = 0.35$$

$$P(\text{ARC} = 10\% \text{ and } \text{SSP} = -8\%) = 0.10$$

$$P(\text{ARC} = -5\% \text{ and } \text{SSP} = 12\%) = 0.15$$

$$P(\text{ARC} = -5\% \text{ and } \text{SSP} = -8\%) = 0.40$$

You can see that the two stocks tend to move in the same direction. There is a 75% chance that they are both high or both low. Likewise there is only a 25% chance that one is high and the other low. Thus, we might expect the covariance to be positive.

Calculating the covariance requires first calculating the expected values for each stock. Considering ARC first, we can use the law of total probability:

$$\begin{aligned} P(\text{ARC} = 10\%) &= P(\text{ARC} = 10\% \text{ and } \text{SSP} = 12\%) \\ &\quad + P(\text{ARC} = 10\% \text{ and } \text{SSP} = -8\%) = 0.35 \\ &\quad + 0.10 = 0.45 \end{aligned}$$

Thus, $P(\text{ARC} = -5\%) = 1 - P(\text{ARC} = 10\%) = 0.55$, and we can calculate

$$\begin{aligned} E(\text{ARC}) &= 0.45(10\%) + 0.55(-5\%) \\ &= 1.75\% \end{aligned}$$

Likewise, we can calculate $E(\text{SSP})$:

$$\begin{aligned} P(\text{SSP} = 12\%) &= P(\text{ARC} = 10\% \text{ and } \text{SSP} = 12\%) \\ &\quad + P(\text{ARC} = -5\% \text{ and } \text{SSP} = 12\%) = 0.35 \\ &\quad + 0.15 = 0.50 \end{aligned}$$

$$\begin{aligned} P(\text{SSP} = -8\%) &= 1 - P(\text{SSP} = 12\%) \\ &= 0.50 \end{aligned}$$

$$\begin{aligned} E(\text{SSP}) &= 0.50(12\%) + 0.50(-8\%) \\ &= 2\% \end{aligned}$$

Now we can calculate the covariance between ARC and SSP:

$$\begin{aligned} \text{Cov}(\text{ARC}, \text{SSP}) &= [10\% - 1.75\%][12\% - 2\%]P(\text{ARC} = 10\% \text{ and } \text{SSP} = 12\%) \\ &\quad + [10\% - 1.75\%][-8\% - 2\%]P(\text{ARC} = 10\% \text{ and } \text{SSP} = -8\%) + \\ &\quad [-5\% - 1.75\%][12\% - 2\%]P(\text{ARC} = -5\% \text{ and } \text{SSP} = 12\%) + \\ &\quad [-5\% - 1.75\%][-8\% - 2\%]P(\text{ARC} = -5\% \text{ and } \text{SSP} = -8\%) \\ &= 8.25\% \times 10\% \times 0.35 \\ &\quad + 8.25\% \times (-10\%) \times 0.10 \\ &\quad + (-6.75\%) \times 10\% \times 0.15 \\ &\quad + (-6.75\%) \times (-10\%) \times 0.40 \\ &= 37.5(\% \text{ squared}) \end{aligned}$$

As expected, the covariance is positive. The problem, however, is that the magnitude of the covariance is not very meaningful because it depends on the range of variation in the two quantities. Also, as with the variance, the covariance carries units that are meaningful. In the case of the two stock returns, the units are % squared. But suppose we wanted to calculate the covariance between hemline height and stock market return; the calculation would involve multiplying inches times percentage points of return, and so the units would be % inches. What in the world is a % inch?

To solve these two problems, we often transform the covariance to get a standardized measure of dependence. This standardized measure is called the *correlation coefficient*, and the Greek symbol ρ (rho) is used to represent it. To calculate ρ , divide the covariance of X and Y by the standard deviations of these two uncertain quantities:

$$\rho_{XY} = \frac{\text{Cov}(X, Y)}{\sigma_X \sigma_Y}$$

The correlation ρ has a number of useful properties. First, it ranges between +1 (perfect positive dependence) and -1 (perfect negative dependence). A correlation of zero suggests no relationship, although certain kinds of dependence are possible even though the correlation is zero (see Exercise 7.11). Also, ρ has no units. In the hemline-stock market example, we would divide the covariance (% inches) by the standard deviation of return (%) and the standard deviation of hemline height (inches), and the units would cancel each other out. An implication is that the correlation is useful for comparing the strength of the relationship in one case with the strength of the relationship in another that involves different variables altogether.

To continue the example, we can calculate the correlation between the returns for stocks ARC and SSP. To do this calculation, we must first calculate the standard deviation for each of the two individual stocks. These are $\sigma_{ARC} = 7.46\%$ and $\sigma_{SSP} = 10\%$. Thus,

$$\rho_{ARC,SSP} = \frac{37.5}{7.46 \times 10} = 0.503$$

This correlation of 0.503 gives the investor an indication of the extent to which the returns are related. By comparing the correlations of different pairs of stocks, the investor can try to locate those with lower (or even negative) correlations in order to accomplish the objective of stabilizing the return of the overall portfolio. This is the principle of portfolio diversification.

One final warning is in order before leaving the ideas of covariance and correlation. Although these measures of dependence are widely used, they only provide insight into a certain kind of dependence. That is, as long as the relationship is such that an increase in one variable always suggests an increase (or always a decrease) in the other, then the covariance and correlation will reflect this relationship. If the relationship is more complex, however, such a relationship may not be adequately reflected in the covariance and correlation. For example as X increases up to a certain point, Y might be expected to increase. As X continues to increase, though, Y might be expected to decrease. This kind of *nonmonotonic* relationship may not be reflected by the covariance and correlation.


Continuous Probability Distributions

The discussion above has focused on discrete uncertain quantities. Now we turn briefly to continuous uncertain quantities. In this case, the uncertain quantity can take any value within some range. For example, the temperature tomorrow at O'Hare Airport in Chicago at noon is an uncertain quantity that can be anywhere between, say, -50°F and 120°F . The length of time until some anticipated event (for example, the next major earthquake in California) is a continuous uncertain quantity, as are locations in space (the precise location of the next earthquake), as well as various measurements such as height, weight, and speed (for example, the peak "ground acceleration" in the next major earthquake).

With continuous uncertain quantities, it is not reasonable to speak of the probability that a specific value occurs. In fact, the probability of a particular value occurring is equal to zero: $P(Y = y) = 0$. Intuitively, there are infinitely many possible values, and so the probability of any particular value must be infinitely small. Instead, we typically speak of interval probabilities: $P(a \leq Y \leq b)$. The CDF for a continuous uncertain quantity can be constructed on the basis of such intervals.

The easiest way to understand this process is through a simple example. Let us suppose we are interested in a movie star's age. For a variety of reasons, we may be certain that she is older than 29 and no older than 65. We can translate these into probability statements: $P(\text{Age} \leq 29) = 0$, and $P(\text{Age} \leq 65) = 1$. Now, suppose that

Figure 7.9
Cumulative distribution function for movie star's age.



you decide that she most likely is between 40 and 50 years old, and that $P(40 < \text{Age} \leq 50) = 0.8$. Also suppose you figure that $P(\text{Age} \leq 40) = 0.05$, and $P(\text{Age} > 50) = 0.15$. Naturally, if you were so inclined, you also could make several other judgments. For example, you might figure that her age is just as likely to be 44 or less as it is to be greater than 44. This would translate into $P(\text{Age} \leq 44) = 0.50$.

We can take the probabilities from this last example, transform them, and create the following table of cumulative probabilities:

- $P(\text{Age} \leq 29) = 0.00$
- $P(\text{Age} \leq 40) = 0.05$
- $P(\text{Age} \leq 44) = 0.50$
- $P(\text{Age} \leq 50) = 0.85$
- $P(\text{Age} \leq 65) = 1.00$

These probabilities then can be displayed in a graph, with years along the horizontal axis and $P(\text{Age} \leq \text{Years})$ along the vertical axis. The graph, shown in Figure 7.9, allows you to select a number of years on the horizontal axis and then read off the probability that the movie star's age is less than or equal to that number of years. As with the discrete case, this graph represents the cumulative distribution function, or CDF.

As before, the CDF allows us to calculate the probability for any interval. For example, $P(40 < \text{Age} \leq 44) = P(\text{Age} \leq 44) - P(\text{Age} \leq 40) = 0.45$. If we were to make more assessments, we could refine the curve in Figure 7.9, but it should always slope upward. If it were to slope downward, it would imply that some interval had a negative probability!

Stochastic Dominance Revisited

Chapter 4 introduced the concept of stochastic dominance as it related to cumulative risk profiles for discrete uncertain quantities. The same principles also, hold true in the continuous case.

Figure 7.10
CDFs for three investment alternatives.
 Investment B stochastically dominates Investment A.

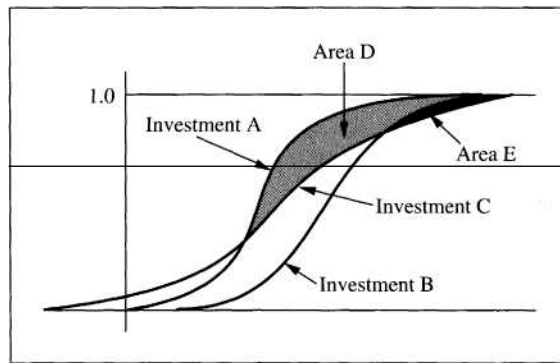


Figure 7.10 shows CDFs for three investment alternatives, each of which is modeled as a continuous uncertain quantity. Investment B stochastically dominates Investment A because the CDF for B lies entirely to the right of the CDF for A. On the other hand, Investment C neither dominates nor is dominated by the other two alternatives. For instance, Area E represents the portion of B's CDF that lies to the left of the CDF for Investment C. Because C's CDF crosses the other two, no decision can be made between C and B (or between C and A) on the basis of stochastic dominance. There is a sense in which B dominates C, however. Imagine an investor who is averse to risk; such an investor would gladly trade any risky prospect for its expected value. B is better than C for such an investor because it has less risk (the distribution is less spread out) and a greater expected value (the distribution is shifted to the right).

Stochastic Dominance and Multiple Attributes (Optional)

Now we are prepared to follow up the discussion of stochastic dominance with multiple attributes in Chapter 4. In the case of multiple attributes, one must consider the joint distribution for all of the attributes together. To develop this, we need to introduce some notation.

First, let $F(X)$ denote the CDF for a variable X . That is, $F(X) = P(X \leq x)$. For example, in the case of the movie star's ages, $F(50) = P(\text{Age} \leq 50)$. Now we can write the condition for stochastic dominance in terms of the F 's. Considering the investments in Figure 7.10, B dominates A because $F_B(x) \leq F_A(x)$ for all values of x on the horizontal axis. This condition asserts that the CDF for B must lie to the right of the CDF for A.

When there are more attributes, the CDF must encompass all of the attributes. For example, recall the summer-job example from Chapter 4, in which we discussed uncertainty about both salary and summer fun. We would have to look at the CDF for both uncertain quantities. The CDF would be denoted by $F(x_s, x_f) = P(\text{Salary} \leq x_s \text{ and Fun} \leq x_f)$. An alternative B (a specific job like the in-town job) dominates alternative A if $F_B(x_s, x_f) \leq F_A(x_s, x_f)$ for all values of x_s and x_f and is strictly less for some x_s and x_f . If we could draw the picture of the graph in three-dimensional space, you could see that this means that the CDF for B must be entirely below the CDF for A and shifted toward larger values for both Salary and Fun.

In Chapter 4, we made the claim that if the uncertain quantities are independent, then stochastic dominance on each of the individual attributes implies overall stochastic dominance. To see this, consider the summer-job example further. Stochastic dominance requires that

$$F_B(x_s, x_f) \leq F_A(x_s, x_f)$$

for all values of x_s and x_f , which is the same as

$$P_B(\text{Salary} \leq x_s \text{ and Fun} \leq x_f) \leq P_A(\text{Salary} \leq x_s \text{ and Fun} \leq x_f)$$

for all values of x_s and x_f . If Salary and Fun are independent for both alternatives A and B, we now know that the joint probabilities in this condition can be rewritten as the product of the individual (or *marginal*) probabilities:

$$P_B(\text{Salary} \leq x_s)P_B(\text{Fun} \leq x_f) \leq P_A(\text{Salary} \leq x_s)P_A(\text{Fun} \leq x_f)$$

Now, suppose that B dominates A individually on each attribute. This means that

$$P_B(\text{Salary} \leq x_s) \leq P_A(\text{Salary} \leq x_s)$$

and

$$P_B(\text{Fun} \leq x_f) \leq P_A(\text{Fun} \leq x_f)$$


If this is true, it is certainly the case that the overall stochastic-dominance condition is met, because the product $P_B(\text{Salary} \leq x_s)P_B(\text{Fun} \leq x_f)$ must be less than or equal to $P_A(\text{Salary} \leq x_s)P_A(\text{Fun} \leq x_f)$.

A final word of caution is in order here. The reasoning above only goes in one direction. That is, *if* the attributes are independent and *if* the individual stochastic-dominance conditions are met, *then* the overall stochastic-dominance condition is also met. That is, we have identified sufficient conditions for overall stochastic dominance. However, it is possible for overall stochastic dominance to exist even though the uncertain quantities are not independent or do not display stochastic dominance in the individual attributes. In other words, in some cases, you might have to go back to the definition of overall stochastic dominance [$F_B(x_1, \dots, x_n) \leq F_A(x_1, \dots, x_n)$ for all x_1, \dots, x_n] in order to determine whether B dominates A.

Probability Density Functions

The CDF for a continuous uncertain quantity corresponds closely to the CDF for the discrete case. Is there some representation that corresponds to the probability mass function? The answer is yes, and that representation is called a *probability density function*, which, if we are speaking of uncertain quantity X , would be denoted typically as $f(x)$. The density function $f(x)$ can be built up from the CDF. It is a function in which the area under the curve within a specific interval represents the probability that the uncertain quantity will fall in that interval. For example, the density function $f(\text{Age})$ for the movie star's age might look something like the graph in Figure 7.11. The total area under the curve equals 1 because the uncertain quantity must take on

Figure 7.11
Probability density
function for movie
star's age.



some value. The shaded area in Figure 7.11 corresponds to $P(40 < \text{Age} < 50)$ and so this area must be equal to 0.80.

Expected Value, Variance, and Standard Deviation: The Continuous Case

As in the discrete case, a continuous probability distribution can have an expected value, variance, and standard deviation. But the definition is not as easy as it was before because now we do not have probabilities for specific values, only probabilities for intervals. Without going into a lot of detail, these characteristics of a continuous probability distribution are defined by using calculus. The definitions for a continuous uncertain quantity X correspond to the discrete case, except that the summation sign is replaced with an integral sign and the density function is used in place of the probabilities:


$$E(X) = \int_{x^-}^{x^+} x f(x) dx$$

$$\text{Var}(X) = \sigma_X^2 = \int_{x^-}^{x^+} [x - E(X)]^2 f(x) dx$$

where x^- and x^+ represent the lower and upper bounds for the uncertain quantity X . As before, the standard deviation σ_x is the square root of the variance.

The interpretation of these formulas also corresponds closely to the summation in the discrete case. It turns out that integration is really the continuous equivalent of the summing operation. Consider the formula for the expected value, for example. Each possible value x between x^- and x^+ is multiplied by (weighted by) the height of the density function $f(x)$. The integration then adds all of these $xf(x)$ products to find $E(X)$. It is as if we had carved up the density function into a very large number of quite narrow but equally wide intervals (Figure 7.12). The relative likelihood of X falling in these different intervals corresponds to the relative height of the density function in each interval. That is, if the density function is (on average) twice as high in one interval as in another, then X is twice as likely to fall into the first interval as the second.

Figure 7.12
A density function $f(x)$ in narrow intervals.
 The probability that X falls in interval A is about twice as great as for B because $f(x)$ is about twice as high in A as it is in B.



As the intervals become infinitesimally narrow, the height of the interval becomes equal to the value $f(x)$. In the limit, as the width of each interval approaches 0, we take each x , multiply by $f(x)$, add these products (by integration), and — bingo! — we get $E(X)$.

Rest assured — you will not be required to perform any integration (in this book anyway). In general, the integration of density functions to obtain expected values and variances is a difficult task and beyond the technical scope of this textbook. Fortunately, mathematicians have studied many different kinds of probability distributions and have performed the integration for you, providing you with formulas for the expected values and variances. We will encounter several of these in Chapter 9.

What about all of those formulas for the expected value and variance of Y , a function of X , or for linear combinations of random variables? Fortunately, *all of those formulas carry over to the continuous case*. If you know the expected value and variance of several uncertain quantities, you can apply the formulas given above regardless of whether the uncertain quantities are continuous or discrete. For example, if $Y = a + bX$, then $E(Y) = a + b E(X)$. If $Y = aX_1 + bX_2$, and X_1 and X_2 are independent, then $\text{Var}(Y) = a^2 \text{Var}(X_1) + b^2 \text{Var}(X_2)$. It does not matter whether the X 's are discrete or continuous.

Covariance and Correlation: The Continuous Case (Optional)

Covariance and correlation also have counterparts when the uncertain quantities are continuous. As with expected value and variance, the definition of covariance uses an integral sign instead of a summation:

$$\text{Cov}(X, Y) = \int_{x^-}^{x^+} \int_{y^-}^{y^+} [x - E(X)][y - E(Y)]f(x, y) dx dy$$

As before, the correlation ρ_{XY} is calculated by dividing $\text{Cov}(X, Y)$ by $\sigma_x \sigma_y$. The double integral in the formula above replaces the double summation in the previous formula for the covariance of two discrete uncertain quantities. The term $f(x, y)$ refers to the *joint density function* for uncertain quantities X and Y . This joint density function is a natural extension of the density function for a single variable; it can be interpreted as a function that indicates the relative likelihood of different (x, y) pairs occurring, and the probability that X and Y fall into any given region can be calculated from $f(x, y)$.

Examples

The formulas and definitions given above are the essential elements of probability theory. With these few tools we will be able to go quite a long way in the construction of uncertainty models for decision situations. The intent has been to present the tools and concepts so that they are easily grasped. It may be worthwhile to memorize the formulas, but most important is understanding them at an intuitive level. To help you cement the concepts in your mind and to show the formulas in action, we turn now to specific examples.

OIL WILDCATTING

An oil company is considering two sites for an exploratory well. Because of budget constraints, only one well can be drilled. Site 1 is fairly risky, with substantial uncertainty about the amount of oil that might be found. On the other hand, Site 2 is fairly certain to produce a low level of oil. The characteristics of the two sites are as follows:

Site 1: Cost to Drill \$100K

Outcome	Payoff
Dry	-100K
Low producer	150K
High producer	500K

If the rock strata underlying Site 1 are characterized by what geologists call a "dome" structure (see Figure 7.13), the chances of finding oil are somewhat greater than if no dome structure exists. The probability of a dome structure is $P(\text{Dome}) = 0.6$. The conditional probabilities of finding oil at Site 1 are given in Table 7.1.

Site 2 is considerably less complicated. A well there is not expected to be a high producer, so the only outcomes considered are a dry hole and a low producer. The cost, outcomes, payoffs, and probabilities are as follows:

Site 2: Cost to Drill \$200K

Outcome	Payoff	Probability
Dry	-200K	0.2
Low producer	50K	0.8

Figure 7.13
Rock strata forming a dome structure. Oil tends to pool at the top of the dome in an oil-bearing layer if the layer above is impermeable.




Table 7.1
 Conditional probabilities of outcomes at site 1.

If Dome Structure Exists

Outcome	P(Outcome Dome)
Dry	0.600
Low	0.250
High	0.150

If No Dome Structure Exists


Outcome	P(Outcome No Dome)
Dry	0.850
Low	0.125
High	<u>0.025</u>
	1.000

The decision tree is shown in Figure 7.14. The problem with it as drawn, however, is that we cannot assign probabilities immediately to the outcomes for Site 1. To find these probabilities, we must use the conditional probabilities and the law of total probability to calculate P(Dry), P(Low), and P(High):

$$\begin{aligned}
 P(\text{Dry}) &= P(\text{Dry} \mid \text{Dome}) P(\text{Dome}) + P(\text{Dry} \mid \text{No Dome}) P(\text{No Dome}) \\
 &= 0.6(0.6) + 0.85(0.4) = 0.70
 \end{aligned}$$

$$\begin{aligned}
 P(\text{Low}) &= P(\text{Low} \mid \text{Dome}) P(\text{Dome}) + P(\text{Low} \mid \text{No Dome}) P(\text{No Dome}) \\
 &= 0.25(0.6) + 0.125(0.4) = 0.20
 \end{aligned}$$

Figure 7.14
Decision tree for oil-wildcatting problem.



$$P(\text{High}) = P(\text{High} \mid \text{Dome}) P(\text{Dome}) + P(\text{High} \mid \text{No Dome}) P(\text{No Dome})$$

$$= 0.15(0.6) + 0.025(0.4) = 0.10$$

Everything works out as it should. The three probabilities are for mutually exclusive and collectively exhaustive outcomes, and they add to 1, just as they should. Folding back the decision tree, we find that Site 1 has the higher EMV (expected payoff):

$$\text{EMV}(\text{Site 1}) = 0.7(-100\text{K}) + 0.2(150\text{K}) + 0.1(500\text{K})$$

$$= 10\text{K}$$

$$\text{EMV}(\text{Site 2}) = 0.2(-200\text{K}) + 0.8(50\text{K}) = 0$$

We also can calculate the variance and standard deviation of the payoffs for each site:

$$\sigma_1^2 = 0.7(-100 - 10)^2 + 0.2(150 - 10)^2 + 0.1(500 - 10)^2 =$$

$$= 0.7(-110)^2 + 0.2(140)^2 + 0.1(490)^2 = 36,400\text{K}^2$$

$$\sigma_1 = 190.79\text{K}$$

$$\sigma_2^2 = 0.2(-200 - 0)^2 + 0.8(50 - 0)^2 =$$

$$= 0.2(-200)^2 + 0.8(50)^2 =$$

$$= 10,000\text{K}^2$$

$$\sigma_2 = 100.00\text{K}$$

If we treat these numbers as measures of variability, then it is clear that Site 1, with its higher variance and standard deviation, is more variable than Site 2. In this sense, Site 1 might be considered to be riskier than Site 2.

Note that we could have drawn the decision tree as in Figure 7.15, with "stacked" probabilities. That is, we could have drawn it with a first chance node indicating

whether or not a dome is present, followed by chance nodes for the amount of oil. These chance nodes would include the conditional probabilities from Table 7.1.

We also could have created a probability table as in Table 7.2. The probabilities in the cells of the table are the joint probabilities of both outcomes happening at the same time. Calculating the joint probabilities for the table requires the definition of conditional probability. For example,

$$P(\text{Low and Dome}) = P(\text{Dome}) P(\text{Low} \mid \text{Dome}) = 0.60(0.25) = 0.15$$

The probability table is easy to construct and easy to understand. Once the probabilities of the joint outcomes are calculated, the probabilities of the individual outcomes then are found by adding across the rows or down the columns. For example, from Table 7.2, we can tell that $P(\text{Dry}) = 0.36 + 0.34 = 0.70$. (You can remember that these are called *marginal* probabilities because they are found in the margins of the table!)

Suppose that the company drills at Site 1 and the well is a high producer. In light of this evidence, does it seem more likely that a dome structure exists? Can we figure out $P(\text{Dome} \mid \text{High})$? This question is part of a larger problem. Figure 7.16 shows a "probability tree" (a decision tree without decisions) that reflects the information that we have been given in the problem. We know $P(\text{Dome})$ and $P(\text{No Dome})$, and we know the conditional probabilities of the amount of oil given the presence or absence of a dome. Thus, our probability tree has the chance node representing the presence or absence of a dome on the left and the chance node representing the amount of oil on the right. Finding $P(\text{Dome} \mid \text{High})$ is a matter of "flipping the tree" so that the chance node for the amount of oil is on the left and the node for the presence or absence of a dome is on the right, as in Figure 7.17.

Flipping a probability tree is the same as reversing an arrow between two chance nodes in an influence diagram. In Figure 7.18a, the direction of the arrow represents the probabilities as they are given in Table 7.1. Because the arrow points to the "Oil

Figure 7.15
An alternative decision tree for the oil-wildcatting problem.

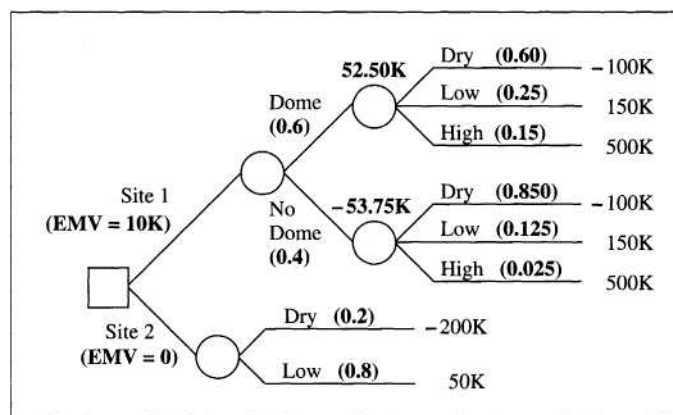


Table 7.2
Calculating the probabilities for Site 1.

	Dome	No Dome
Dry	0.36	0.34
Low	0.15	0.05
High	<u>0.09</u>	<u>0.01</u>
	0.60	<hr/> 0.40 <hr/> 1.00

Figure 7.16
Probability tree for the uncertainty faced at Site 1.

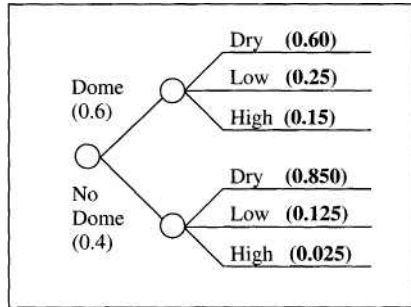



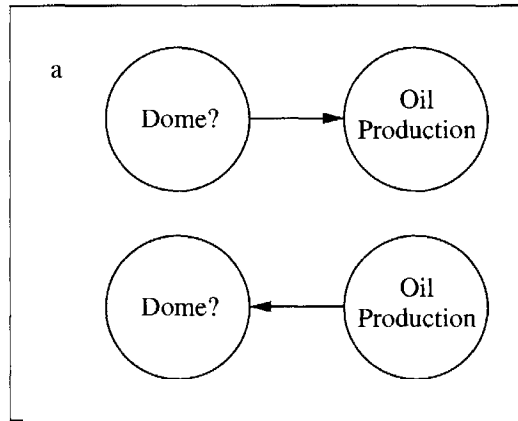
Figure 7.17
Flipping the probability tree in Figure 7.16.



Production" node, the probabilities for the level of oil production are conditioned on whether or not there is a dome. Figure 7.18b shows the arrow reversed. Now the probability of a dome or no dome is conditioned on the amount of oil produced. Changing the direction of the arrow has the same effect as flipping the probability tree.

Simply flipping the tree or turning the arrow around in the influence diagram is not enough. The question marks in Figure 7.17 indicate that we do not yet know what the probabilities are for the flipped tree. The probabilities in the new tree must be consistent with the probabilities in the original version. Consistency means that if we were to calculate $P(\text{Dome})$ and $P(\text{No Dome})$ using the probabilities in the new tree, we would get 0.60 and 0.40, the probabilities we started with in Figure 7.16. How do we ensure consistency?

Figure 7.18
Reversing arrows between chance nodes in an influence diagram.




We can find $P(\text{High})$, $P(\text{Low})$, and $P(\text{Dry})$ by using the law of total probability. We did these calculations above for Figure 7.14, and we found that $P(\text{High}) = 0.10$, $P(\text{Low}) = 0.20$, and $P(\text{Dry}) = 0.70$. How about finding the new conditional probabilities? Bayes' theorem provides the answer. For example:

$$\begin{aligned}
 P(\text{Dome} \mid \text{High}) &= \frac{P(\text{High} \mid \text{Dome}) P(\text{Dome})}{P(\text{High})} \\
 &= \frac{P(\text{High} \mid \text{Dome}) P(\text{Dome})}{P(\text{High} \mid \text{Dome}) P(\text{Dome}) + P(\text{High} \mid \text{No Dome}) P(\text{No Dome})} \\
 &= \frac{0.15(0.6)}{0.15(0.6) + 0.025(0.4)} \\
 &= \frac{0.15(0.6)}{0.10} = 0.90
 \end{aligned}$$

Probabilities that have the same conditions must add to 1, and so $P(\text{No Dome} \mid \text{High})$ must be equal to 0.10. Likewise, we can calculate the conditional probabilities of a dome or no dome given a dry hole or a low producer. These probabilities are shown in Figure 7.19.

If you did not enjoy using Bayes' theorem directly to calculate the conditional probabilities required in flipping the tree, you may be pleased to learn that the probability table (Table 7.2) has all the information needed. Recall that the entries in the cells inside the table are the joint probabilities. For example, $P(\text{High and Dome}) = 0.09$. We need

Figure 7.19
The flipped probability tree with new probabilities.



$P(\text{Dome} \mid \text{High}) = P(\text{High and Dome})/P(\text{High})$, which is just $0.09/0.10 = 0.90$. That is, we take the joint probability from inside the table and divide it by the probability of the outcome on the right side of the vertical bar. This probability is found in the margin of the probability table.

Whether we use Bayes' theorem directly or the probability-table approach to flip the probability tree, the conditional probabilities that we obtain have an interesting interpretation. Recall that we started with $P(\text{Dome}) = 0.60$. After finding that the well was a high producer, we were able to calculate $P(\text{Dome} \mid \text{High})$. This probability sometimes is called a *posterior* probability, indicating that it is the result of revising the original probability after gathering data. In contrast, the probability with which we started, in this case $P(\text{Dome}) = 0.60$, sometimes is called the *prior* probability. One way to think about Bayes' theorem is that it provides a mechanism to update prior probabilities when new information becomes available. [Source: The oil-wildcatting example and analysis are adapted from C.A. Holloway (1979) *Decision Making under Uncertainty: Models and Choices*, pp. 195-200. Englewood Cliffs, NJ: Prentice-Hall.]

JOHN HINCKLEY'S TRIAL

In 1982 John Hinckley was on trial, accused of having attempted to kill President Reagan. During Hinckley's trial, Dr. Daniel R. Weinberger told the court that when individuals diagnosed as schizophrenics were given computerized axial tomography (CAT) scans, the scans showed brain atrophy in 30% of the cases compared with only 2% of the scans done on normal people. Hinckley's defense attorney wanted to introduce as evidence Hinckley's CAT scan, which showed brain atrophy. The defense argued that the presence of atrophy strengthened the case that Hinckley suffered from mental illness.

We can use Bayes' theorem easily to analyze this situation. We want to know the probability that Hinckley was schizophrenic given that he had brain atrophy. Approximately 1.5% of people in the United States suffer from schizophrenia. This is the *base rate*, which we can interpret as the prior probability of schizophrenia before we find out about the condition of an individual's brain. Thus, $P(S) = 0.015$, where S means schizophrenia. We also have $P(A \mid S) = 0.30$ (A means atrophy) and $P(A \mid \bar{S}) = 0.02$. We want $P(S \mid A)$. Bayes' theorem provides the mathematical mechanism for flipping the probability from $P(A \mid S)$ to $P(S \mid A)$:

$$\begin{aligned} P(S \mid A) &= \frac{P(A \mid S) P(S)}{P(A \mid S) P(S) + P(A \mid \bar{S}) P(\bar{S})} \\ &= \frac{0.30(0.015)}{0.30(0.015) + 0.02(0.985)} \\ &= 0.186 \end{aligned}$$

Thus, given that his brain showed such atrophy, Hinckley still has less than a 1-in-5 chance of being schizophrenic. Given the situation, this is perhaps a surprisingly low probability. The intuition behind this result is that there are many false-positive tests. If we tested 100,000 individuals, some 1500 of them would be schizophrenic and 98,500 would be normal (or at least not schizophrenic). Of the 1500, only 30%, or approximately 450, would show atrophy. Of the 98,500, 2% (some 1970) would show brain atrophy. If a single individual has atrophy, is he one of the 450 with schizophrenia or one of the 1970 without? Note that

$$0.186 = \frac{450}{450 + 1970}$$

The real question is whether this is good news or bad news for Hinckley. The prosecution might argue that the probability of schizophrenia is too small to make any difference; even in light of the CAT-scan evidence, Hinckley is less than one-fourth as likely to be schizophrenic as not. On the other hand, the defense would counter, 0.186 is much larger than 0.015. Thus, the CAT-scan results indicate that Hinckley was more than 12 times as likely to be schizophrenic than a randomly chosen person on the street.

Now, however, consider what we have done in applying Bayes' theorem. We have used a prior probability of 0.015, which essentially is the probability that a randomly chosen person from the population is schizophrenic. *But Hinckley was not randomly chosen.* In fact, it does not seem reasonable to think of Hinckley, a man accused of attempted assassination, as the typical person on the street.


If 0.015 is not an appropriate prior probability, what is? It may not be obvious what an appropriate prior probability should be, so let us consider a sensitivity-analysis approach and see what different priors would imply. Imagine a juror who, before encountering the CAT-scan evidence, believes that there is only a 10% chance that Hinckley is schizophrenic. For most of us, this would be a fairly strong statement; Hinckley is nine times as likely to be normal as schizophrenic. Now consider the impact of the CAT-scan evidence on this prior probability. We can calculate this juror's posterior probability:

$$\begin{aligned} P(S | A) &= \frac{0.30(0.10)}{0.30(0.10) + 0.02(0.90)} \\ &= 0.63 \end{aligned}$$

$p(S | A) = 0.63$ is a substantial probability. We can do this for a variety of values for the prior probability. Figure 7.20 shows the posterior probability $P(S | A)$ graphed as a function of the prior probability that Hinckley was schizophrenic.

As a result of this discussion, it is clear that a juror need not have a very strong prior belief for the CAT-scan evidence to have an overwhelming effect on his or her posterior belief of Hinckley's mental illness. Furthermore, no matter what the juror's prior belief, the CAT-scan result must increase the probability that Hinckley was schizophrenic. [Source: A. Barnett, I. Greenberg, and R. Machol (1984) "Hinckley and the Chemical Bath." *Interfaces*, 14, 48-52.]

Figure 7.20 Graph of the posterior probability that Hinckley was schizophrenic plotted against the prior probability.



Decision-Analysis Software and Bayes' Theorem

As we have noted above, flipping conditional probabilities by means of Bayes' theorem can seem like a complicated task. In fact, the rules are straightforward: Calculate the joint probability and divide by the appropriate marginal probability. The problems are being sure that the joint probability is calculated correctly and selecting the appropriate marginal probability. Fortunately, if all of the probabilities are defined (for example, in a probability tree or table), applying Bayes' theorem is something that computers can do much more readily than humans!

Many decision-analysis software packages have the ability to perform Bayes' theorem calculations automatically. For example, DPL will automatically flip probabilities as needed in analyzing a decision model. DATA permits the user to enter probability data that must be flipped before analyzing an appropriately constructed decision tree. Finally, an electronic spreadsheet like Microsoft Excel can be programmed to flip probabilities based on a probability table like that shown in Table 7.2.

SUMMARY

The definitions and formulas at the beginning of the chapter are the building blocks of probability theory. Venn diagrams provide a way to visualize these probability "laws." We also discussed uncertain quantities, their probability distributions, and characteristics of those distributions such as expected value, variance, and standard deviation. The use of probability concepts and the manipulation of probabilities were demonstrated in the oil-wildcatting and Hinckley trial examples.

EXERCISES

- 7.1 Explain why probability is important in decision analysis.
- 7.2 Explain in your own words what an uncertain quantity or random variable is. Why is the idea of an uncertain quantity important in decision analysis?
- 7.3 You are given the following probability table:

	A	\bar{A}	
B	0.12	0.53	0.65
\bar{B}	0.29	0.06	0.35
	0.41	0.59	1.00

Use the probability table to find the following:

$$P(A \text{ and } B), P(A \text{ and } \bar{B}), P(A), P(B), P(\bar{B}), P(B | A), P(A | B), P(\bar{A} | \bar{B})$$

- 7.4 Use the probability table in Exercise 7.3 to find $P(A \text{ or } B)$, or the outcome where either A occurs or B occurs (or both).
- 7.5 The Outcome A or B sometimes is called the union of A and B. The union event occurs if either Outcome A or Outcome B (or both) occurs. Suppose that both A and B can happen at the same time (that is, their areas overlap in a Venn diagram). Show that $P(A \text{ or } B) = P(A) + P(B) - P(A \text{ and } B)$. Use a Venn diagram to explain. Why is this result consistent with the second requirement that probabilities add up?
- 7.6 Often it is difficult to distinguish between the probability of an intersection of outcomes (joint probability) and the probability of a conditional outcome (conditional probability). Classify the following as joint probability statements or conditional probability statements. [If in doubt, try to write down the probability statement; for example, $P(\text{Crash Landing} | \text{Out of Fuel})$ or $P(\text{Dow Jones Up and Stock Price Up})$.]
- Eight percent of the students in a class were left-handed and red-haired. **b** Of the left-handed students, 20% had red hair.
 - If the Orioles lose their next game, then the Cubs have a 90% chance of winning the pennant.
 - Fifty-nine percent of the people with a positive test result had the disease.
 - For 78% of patients, the surgery is a success and the cancer never reappears.
 - If the surgery is a success, the cancer is unlikely to reappear.
 - Given the drought, food prices are likely to increase.
 - There is an even chance that a farmer who loses his crop will go bankrupt.
 - If the temperature is high and there is no rain, farmers probably will lose their crops.

j John probably will be arrested because he is trading on insider information. **k** John probably will trade on insider information and get caught.

7.7 Calculate the variance and standard deviation of the payoffs for Products B and C in the GPC case at the end of Chapter 4 (page 145).

7.8 $P(A) = 0.42$, $P(B | A) = 0.66$, and $P(B | \bar{A}) = 0.25$. Find the following:

$$P(\bar{A}), P(\bar{B} | A), P(\bar{B} | \bar{A}), P(B), P(\bar{B}), P(A | B), P(\bar{A} | B), P(A | \bar{B}), P(\bar{A} | \bar{B}).$$

7.9 $P(A) = 0.10$, $P(B | A) = 0.39$, and $P(B | \bar{A}) = 0.39$. Find the following:

$$P(\bar{A}), P(\bar{B} | A), P(\bar{B} | \bar{A}), P(B), P(\bar{B}), P(A | B), P(\bar{A} | B), P(A | \bar{B}), P(\bar{A} | \bar{B}).$$

How would you describe the relationship between Outcomes A and B?

7.10 Consider the following probabilities:

$$P(X = 2) = 0.3$$

$$P(X = 4) = 0.7$$

$$P(Y = 10 | X = 2) = 0.9$$

$$P(Y = 20 | X = 2) = 0.1$$

$$P(Y = 10 | X = 4) = 0.25$$

$$P(Y = 20 | X = 4) = 0.75$$

Calculate the covariance and correlation between X and Y .

7.11 Consider the following joint probability distribution for uncertain quantities X and Y :

$$\begin{aligned} P(X = -2 \text{ and } Y = 2) &= \\ 0.2 P(X = -1 \text{ and } Y = 1) &= \\ = 0.2 P(X = 0 \text{ and } Y = 0) &= \\ = 0.2 P(X = 1 \text{ and } Y = 1) &= \\ 0.2 P(X = 2 \text{ and } Y = 2) &= \\ 0.2 & \end{aligned}$$

a Calculate the covariance and correlation between X and Y .

b Calculate $P(Y = 2)$, $P(Y = 2 | X = -2)$, $P(Y = 2 | X = 2)$, $P(Y = 2 | X = 0)$.

c Calculate $P(X = -2)$, $P(X = -2 | Y = 2)$, $P(X = -2 | Y = 0)$.

d Are X and Y dependent or independent? How would you describe the relationship between X and Y ?

7.12 Write probability statements relating hemline height, stock market prices, and "adventurousness" that correspond to the following description of these relationships:

Another example is that high skirt hemlines tend to occur when the stock market is advancing steadily (a bull market), although no one believes that hemlines cause the stock market to advance, or vice versa. Perhaps both reflect a pervasive feeling of confidence and adventure in our society. If we had an adequate index

of "societal adventuresomeness," hemline height and stock market activity might well be conditionally independent given the index.

Even though we distinguish between continuous and discrete uncertain quantities in reality everything is discrete if only because of limitations inherent in our measuring devices. For example, we can only measure time or distance to a certain level of precision. Why then, do we make the distinction between continuous and discrete uncertain quantities? What value is there in using continuous uncertain quantities in a decision model?

- 7.14** $P(A) = 0.68$, $P(B|A) = 0.30$, and $P(B|\bar{A}) = 0.02$. Find $P(\bar{A})$, $P(A \text{ and } B)$, and $P(\bar{A} \text{ and } B)$. Use these to construct a probability table. Now use the table to find the following:

$$P(\bar{B}|A), P(\bar{B}|\bar{A}), P(B), P(\bar{B}), P(A|B), P(\bar{A}|B), P(A|\bar{B}), P(\bar{A}|\bar{B}).$$

- 7.15** Julie Myers, a graduating senior in accounting, is preparing for an interview with a bi-eight accounting firm. Before the interview, she sets her chances of eventually getting an offer at 50%. Then, on thinking about her friends who have interviewed and gotten offers from this firm, she realizes that of the people who received offers, 95% had good interviews. On the other hand, of those who did not receive offers, 75% said they had good interviews. If Julie Myers has a good interview, what are her chances of receiving an offer?
- 7.16** Find the expected value, variance, and standard deviation of X in the following probability distributions:
- a** $P(X = 1) = 0.05$, $P(X = 2) = 0.45$, $P(X = 3) = 0.30$, $P(X = 4) = 0.20$. **b** $P(X = -20) = 0.13$, $P(X = 0) = 0.58$, $P(X = 100) = 0.29$. **c** $P(X = 0) = 0.368$, $P(X = 1) = 0.632$.
- 7.17** If $P(X = 1) = p$ and $P(X = 0) = 1 - p$, show that $E(X) = p$ and $\text{Var}(X) = p(1 - p)$.
- 7.18** If $P(A|B) = p$, must $P(A|\bar{B}) = 1 - p$? Explain.
- 7.19** Suppose that a company produces three different products. The sales for each product are independent of the sales for the others. The information for these products is given below:

Product	Price (\$)	Expected Unit Sales	Variance of Unit Sales
A	3.50	2000	1000
B	2.00	10,000	6400
C	1.87	8500	1150

- a** What are the expected revenue and variance of the revenue from Product A alone?
- b** What are the company's overall expected revenue and variance of its revenue?

7.20 A company owns two different computers, which are in separate buildings and operated entirely separately. Based on past history, Computer 1 is expected to break down 5.0 times a year, with a variance of 6, and costing \$200 per breakdown. Computer 2 is expected to break down 3.6 times per year, with a variance of 7, and costing \$165 per breakdown. What is the company's expected cost for computer breakdowns and the variance of the

breakdown cost? What assumption must you make to find the variance? Is this a reasonable assumption?

- 7.21 A firm is negotiating with a local club to supply materials for a party. The firm's manager expects to sell 100 large bags of pretzels for \$3 each or 300 for \$2 each; these two outcomes are judged to be equally likely. The expected number of bags sold is $200 = (100 + 300)/2$, and expected price is $\$2.50 = (\$3 + \$2)/2$. The manager then calculates expected revenue as the expected number sold times the expected price: $E(\text{Revenue}) = 200 (\$2.50) = \$500$. What is wrong with the manager's calculation?
- 7.22 Flip the probability tree shown in Figure 7.21.
- 7.23 Figure 7.22 shows part of an influence diagram for a chemical that is considered potentially carcinogenic. How would you describe the relationship between the test results and the field results?
- 7.24 Let CP denote carcinogenic potential, TR test results, and FR field results. Suppose that for Figure 7.22 we have the following probabilities:

$$P(\text{CP High}) = 0.27 \quad P(\text{CP Low}) = 0.73$$

Figure 7.21
A probability tree representing the diagnostic performance of a medical test.

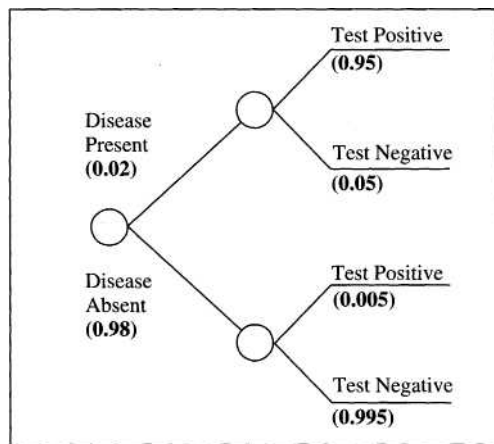
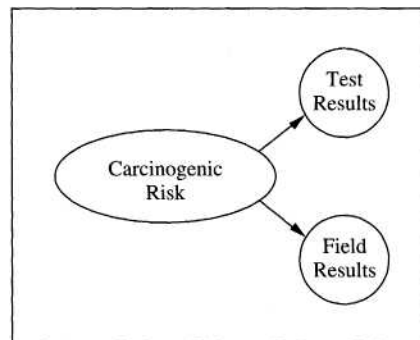


Figure 7.22
An influence diagram for a potentially carcinogenic chemical



$$\begin{aligned} P(\text{TR Positive} \mid \text{CP High}) &= 0.82 & P(\text{TR Positive} \mid \text{CP Low}) &= 0.21 \\ P(\text{TR Negative} \mid \text{CP High}) &= 0.18 & P(\text{TR Negative} \mid \text{CP Low}) &= 0.79 \end{aligned}$$

$$\begin{aligned} P(\text{FR Positive} \mid \text{CP High}) &= 0.95 & P(\text{FR Positive} \mid \text{CP Low}) &= 0.17 \\ P(\text{FR Negative} \mid \text{CP High}) &= 0.05 & P(\text{FR Negative} \mid \text{CP Low}) &= 0.83 \end{aligned}$$

Find the following:

$$\begin{aligned} &P(\text{TR Positive and FR Positive} \mid \text{CP High}) \\ &P(\text{TR Positive and FR Negative} \mid \text{CP High}) \\ &P(\text{TR Negative and FR Negative} \mid \text{CP Low}) \\ &P(\text{TR Negative and FR Positive} \mid \text{CP Low}) \end{aligned}$$

QUESTIONS AND PROBLEMS

- 7.25** Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice and also participated in antinuclear demonstrations. Use your judgment to rank the following statements by their probability, using 1 for the most probable statement and 8 for the least probable:
- Linda is a teacher in an elementary school.
 - Linda works in a bookstore and takes Yoga classes.
 - Linda is active in the feminist movement.
 - Linda is a psychiatric social worker.
 - Linda is a member of the League of Women Voters.
 - Linda is a bank teller.
 - Linda is an insurance salesperson.
 - Linda is a bank teller and is active in the feminist movement.

Source: Amos Tversky and Daniel Kahneman (1982) "Judgments of and by Representativeness." In D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment under Uncertainty: Heuristics and Biases*, pp. 84-98. Cambridge: Cambridge University Press.

- 7.26** The description and statements given in Problem 7.25 often elicit responses that are not consistent with probability requirements. If you are like most people, you ranked statement h (Linda is a bank teller and is active in the feminist movement) as more probable than statement f (Linda is a bank teller).
- Explain why you ranked statements h and f as you did.
 - Statement h is actually a compound event. That is, for h to be true, Linda must be both a bank teller (Outcome A) and active in the feminist movement (Outcome B). Thus, statement h is represented by the Outcome "A and B." Use a Venn diagram to explain why statement h must be less probable than statement f.
 - Suppose that you have presented Problem 7.25 to a friend, who ranks statement h as more probable than statement f. Your friend argues as follows: "Well, it's not very likely that Linda is a bank teller in the first place. But if she is a bank teller, then she

is very likely to be active in the feminist movement. So h would appear to be more likely than f." How is your friend interpreting statement h? Explain why this is not an appropriate interpretation.

7.27 Suppose you are a contestant on the television game show, "Let's Make a Deal." You must choose one of three closed doors, labeled A, B, and C, and you will receive what ever is behind the chosen door. Behind two of the doors is a goat, and behind the third is a new car. Like most people, you have no use for a goat, but the car would be very welcome. Suppose you choose Door A. Now the host opens Door B, revealing a goat, and offers to let you switch from Door A to Door C. Do you switch? Why or why not? [*Hint*: What can you assume about the host's behavior? Would it ever be the case that the host would open a door to reveal the car before asking if you want to switch?]

7.28 On the fictitious television game show, "Marginal Analysis for Everyone," the host subjects contestants to unusual tests of mental skill. On one, a contestant may choose one of two identical envelopes labeled A and B, each of which contains an unknown amount of money. The host reveals, though, that one envelope contains twice as much as the other. After choosing A, the host suggests that the contestant might want to switch.

"Switching is clearly advantageous," intones the host persuasively. "Suppose you have amount x in your Envelope A. Then B must contain either $x/2$ (with probability 0.5) or $2x$ (also with probability 0.5). Thus, the expected value of switching is 1.25 x . In fact now that I think about it, I'll only let you switch if you give me a 10% cut of your winnings. What do you say? You'll still be ahead."

"No deal," replies the contestant. "But I'll be happy to switch for free. In fact, I'll even let you choose which envelope I get. I won't even charge you anything!"

What is wrong with the host's analysis?

Source: This problem was suggested by Ross Shachter.

7.29 Draw an influence diagram for the oil-wildcatting problem. Construct the necessary tables represented by the nodes, including the payoff table, and solve the influence diagram. If possible, use a computer program to do this problem.

7.30 Finding the variance and standard deviation for the payoff from Product A in the GPC case at the end of Chapter 4 (page 145) is somewhat complicated because a pricing decision must be made. How would you calculate the variance and standard deviation for Product A's payoff? (Recall the approach we used in solving decision trees. We were able to "prune" decision branches that were suboptimal.)

7.31 Calculate the variance and standard deviation of the payoffs in the final court decision in the Texaco-Pennzoil case as diagrammed in Figure 4.2.

7.32 In the oil-wildcatting problem, suppose that the company could collect information from a drilling core sample and analyze it to determine whether a dome structure exists at Site 1. A positive result would indicate the presence of a dome, and a negative result would indicate the absence of a dome. The test is not perfect, however. The test is highly accurate for detecting a dome; if there is a dome, then the test shows a positive result 99% of the time. On the other hand, if there is no dome, the probability of a negative result is only 0.85. Thus, $P(+ | \text{Dome}) = 0.99$ and $P(- | \text{No Dome}) = 0.85$. Use these probabilities, the information given in the example, and Bayes' theorem to find the posterior probabilities $P(\text{Dome} | +)$ and $P(\text{Dome} | -)$. If the test gives a positive result, which site

should be selected? Calculate expected values to support your conclusion! If the test result is negative, which site should be chosen? Again, calculate expected values.

- 7.33 In Problem 7.32, calculate the probability that the test is positive and a dome structure exists [P(+ and Dome)]. Now calculate the probability of a positive result, a dome structure, and a dry hole [P(+ and Dome and Dry)]. Finally, calculate P(Dome | + and Dry).
- 7.34 Referring to the oil-wildcatting decision diagrammed in Figure 7.15, suppose that the decision maker has not yet assessed P(Dome) for Site 1. Find the value of P(Dome) for which the two sites have the same EMV. If the decision maker believes that P(Dome) is somewhere between 0.55 and 0.65, what action should be taken?
- 7.35 Again referring to Figure 7.15, suppose the decision maker has not yet assessed P(Dry) for Site 2 or P(Dome) for Site 1. Let P(Dry) = p and P(Dome) = q . Construct a two-way sensitivity analysis graph for this decision problem.
- 7.36 Refer to Exercises 7.23 and 7.24. Calculate P(FR Positive) and P(FR Positive | TR Positive). [Hints: P(FR Positive) is a fairly simple calculation. To find P(FR_Positive | TR Positive), first let B = FR Positive, C = TR Positive, A = CP High, and \bar{A} = CP Low. Now expand P(FR Positive | TR Positive) using the law of total probability in this form:

$$P(B | C) = P(B | A, C) P(A | C) + P(B | \bar{A}, C) P(\bar{A} | C)$$

Now all of the probabilities on the right-hand side can be calculated using the information in the problem.]

Compare P(FR Positive) and P(FR Positive | TR Positive). Would you say that the test results and field results are independent? Why or why not? Discuss the difference between conditional independence and regular independence.

C A S E S T U D I E S

DECISION ANALYSIS MONTHLY

Peter Finch looked at the numbers for the renewals of subscriptions to *Decision Analysis Monthly* magazine. For both May and June he had figures for the percentage of expiring subscriptions that were gift subscriptions, promotional subscriptions, and from previous subscribers. Furthermore, his data showed what proportion of the expiring subscriptions in each category had been renewed (see Table 7.3).

Finch was confused as he considered these numbers. Robert Calloway, who had assembled the data, had told him that the overall proportion of renewals had dropped from May to June. But the figures showed clearly that the proportion renewed had increased in each category. How could the overall proportion possibly have gone down? Peter got a pencil and pad of paper to check Calloway's figures. He had to report to his boss this afternoon and wanted to be able to tell him whether these figures represented good news or bad news.

Question

- 1 Do the data represent good news or bad news regarding renewal trends? Why?

Table 7.3
Subscription data for
Decision Analysis
Monthly.

May Subscription Data	Expiring Subscriptions, %	Proportion Renewed
Gift Subscriptions	70	0.75
Promotional Subscriptions	20	0.50
Previous Subscribers	<u>10</u>	0.10
Total	100	
June Subscription Data	Expiring Subscriptions, %	Proportion Renewed
Gift Subscriptions	45	0.85
Promotional Subscriptions	10	0.60
Previous Subscribers	<u>45</u>	0.20
Total	100	

SCREENING FOR COLORECTAL CANCER

The fecal occult blood test, widely used both in physicians' offices and at home to screen patients for colon and rectal cancer, examines a patient's stool sample for blood, a condition indicating that the cancer may be present. A recent study funded by the National Cancer Institute found that of 15,000 people tested on an annual basis, 10% were found to have blood in their stools. These 10% underwent further testing, including *colonoscopy*, the insertion of an optical-fiber tube through the rectum in order to inspect the colon and rectum visually for direct indications of cancer. Only 2.5% of those having colonoscopy actually had cancer. Additional information in the study suggests that, of the patients who were tested, approximately 5 out of 1000 tested negative (no blood in the stool) but eventually did develop cancer.

Questions

- 1 Create a probability table that shows the relationship between blood in a stool sample and colorectal cancer. Calculate $P(\text{Cancer} \mid \text{Blood})$ and $P(\text{Cancer} \mid \text{No Blood})$.
- 2 The study results have led some medical researchers to agree with the American Cancer Society's long-standing recommendation that all U.S. residents over 50 years of age be tested annually. On the other hand, many researchers claim the costs of such screening, including the cost of follow-up testing on 10% of the population, far exceeds its value. Assume that the test can be performed for as little as \$10 per person, that colonoscopy costs \$750 on average, and that about 60 million people in the United States are over age 50. What is the expected cost (including follow-up colonoscopy) of implementing a policy of screening everyone over age 50? What is the expected number of people who must undergo colonoscopy? What is the expected number of people who must undergo colonoscopy only to find that they do not have cancer after all?

- 3 Over 13 years of follow-up study, 0.6% (6 out of 1000) of those who were screened annually with the fecal occult blood test died from colon cancer anyway. Of those who were not screened, 0.9% (9 out of 1000) died of colon cancer during the same 13 years. Thus, the screening procedure saves approximately 3 lives per 1000 every 13 years. Use this information, along with your calculations from Questions 1 and 2, to determine the expected cost of saving a life by implementing a policy requiring everyone over 50 to be screened every year.
- 4 What is your conclusion? Do you think everyone over 50 should be screened? From your personal point of view, informed now by your calculations above, would the saved lives be worth the money spent and the inconvenience, worry, discomfort, and potential complications of subjecting approximately 6 million people each year to colonoscopy even though relatively few of them actually have detectable and curable cancer?

Source: J. S. Mandel, J. H. Bond, T. R. Church, D. C. Snover, G. M. Braley, L. M. Schuman, and F. Ederer, (1993) "Reducing Mortality from Colorectal Cancer by Screening for Fecal Occult Blood. (Minnesota Colon Cancer Control Study)." *New England Journal of Medicine*, 328 (19), 1365-1371.

AIDS

Acquired immune deficiency syndrome (AIDS) is the most frightening disease of the late twentieth century. The disease attacks and disables the immune system, leaving the body open to other diseases and infection. It is almost always fatal, although years may pass between infection and the development of the disease. As of December 31, 1993, there were 361,509 confirmed AIDS cases and 220,871 deaths due to AIDS in the United States.

Even more frightening is the process by which the disease travels through the population. AIDS is caused by a virus (human T-lymphotropic virus, Type III, or HTLV-III, although more commonly listed as HIV). The virus is transmitted through blood, semen, and vaginal secretions, and may attack virtually anyone who engages in any of several risky behaviors. The extent of the concern about AIDS among public health officials is reflected in the fact that the U.S. Surgeon General's office mailed brochures on AIDS to 107 million households in May and June 1988, the largest single mailing undertaken by the federal government to that time.

When an individual becomes infected, the body produces a special antibody in an effort to counteract the virus. But it can be as long as 12 weeks before these antibodies appear, and it may be years before any signs or symptoms of AIDS infection appear. During this time the individual may not be aware that he or she is infected and thus can spread the disease inadvertently.

Because of the delayed reaction of the virus, there are many more infected individuals than reported AIDS cases. Epidemiologists estimate that about 1 million people in the United States are infected with HIV and thus are potentially infectious to others. Worldwide, the estimate is that over 13 million people are infected. Because of this and because of the way the disease is transmitted, the best way to avoid AIDS simply is to avoid risky behavior. Do not share drug needles. Use a latex condom during intercourse unless you are certain that your partner is not infected.

To help reduce the rate at which AIDS is spread, a number of local (and often controversial) programs have sprung up to provide condoms to sexually active teenagers and clean needles to intravenous drug users.

If 1 million people in the United States *are* HIV-infected, this amounts to an overall rate of approximately 0.38% in a population of 260 million people. The term *seroprevalence* refers to the incidence of HIV in the population. Seroprevalence rates vary dramatically depending on the particular subpopulation. For example, the overall seroprevalence rate for adult men in the United States is about 1%, but for adult women it is 0.13%. Figures 7.23 and 7.24 show seroprevalence rates for male and female military recruits, respectively, in each of the 50 states plus the District of Columbia and Puerto Rico. For higher risk groups the seroprevalence rates are much higher. For example, nationally about 25.5% of gay and bisexual men are infected, but the rate ranges from 3.9% in Providence, Rhode Island, to 47.4% in Atlanta, Georgia. For intravenous drug users, the seroprevalence rate is 7.5% nationwide, reaching as high as 40.4% in New York City.

The best tests available for AIDS detect the antibodies rather than the virus. Two such tests are generally available and widely used. The first is the *enzyme-linked immunosorbent assay* (ELISA). An individual is considered to have a positive result on the ELISA test only if both of two separate trials yield positive results. The performance of such a diagnostic test can be measured by the probabilities associated with correct diagnosis. The probability that an infected individual tests positive, $P(\text{ELISA}^+ \mid \text{Infected})$, is called the *sensitivity* of the test. The probability that an uninfected individual tests negative, $P(\text{ELISA}^- \mid \text{Not Infected})$, is called the *specificity* of the test. A negative result for an infected individual is called a *false-negative*, and

Figure 7.23
HIV seroprevalence among male applicants for military service by state of residence, United States, January 1991 through December 1992.
 Source: Department of Defense.

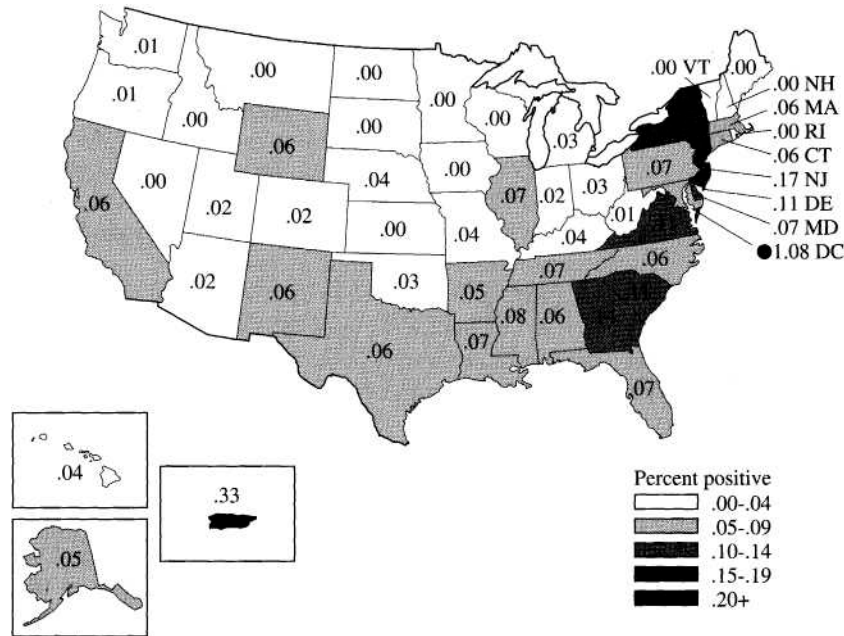
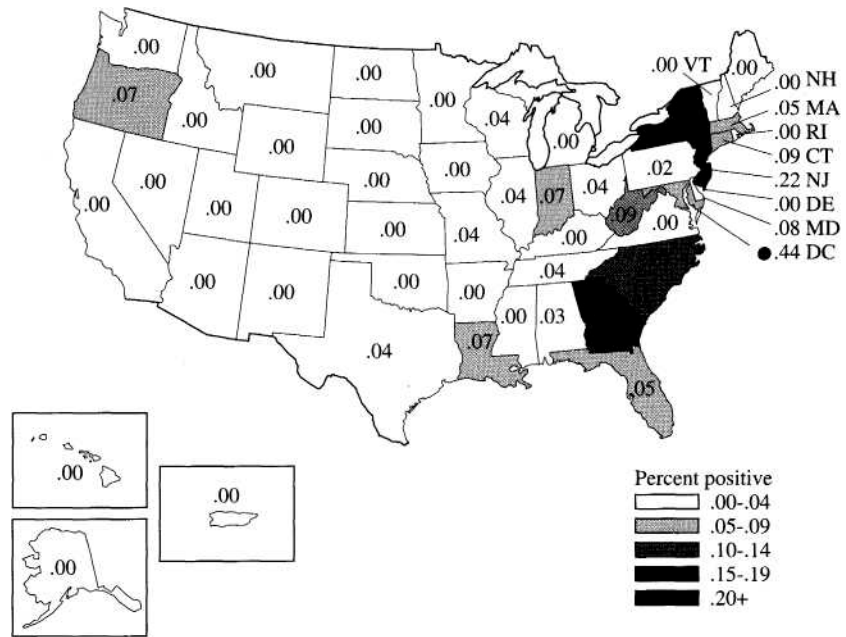


Figure 7.24 HIV seroprevalence among female applicants for military service by state of residence, United States, January 1991 through December 1992. Source: Department of Defense.



a positive result for someone who is not infected is called a *false-positive*. An ideal test would have 100% sensitivity and 100% specificity, giving correct diagnosis all the time with neither false-negatives nor false-positives. In 1990, the Centers for Disease Control (CDC) reported a study indicating that the sensitivity and specificity of the ELISA test are 0.997 and 0.985, respectively.

The *Western blot* test generally is used to confirm or disconfirm a positive ELISA result. For some time the Western blot was considered to be a perfect test, but this may not be true. The Western blot is a labor-intensive test, the results of which require interpretation by a skilled laboratory technician. The same CDC study indicated sensitivity and specificity of the Western blot to be 0.993 and 0.916, respectively. Table 7.4 summarizes the performance characteristics of the two tests.

Questions

- 1 Given that the tests are not perfect, it is worthwhile to calculate the probability of being infected, given the various combinations of results on the tests. Calculate the probability of being infected given a positive ELISA test result, $P(\text{Inf} | \text{ELISA}+)$. Use as your prior probability the $P(\text{Inf}) = 0.0038$, the estimated overall rate of infection in the United States. For $P(\text{ELISA} + | \text{Inf})$ and $P(\text{ELISA} + | \text{Not Inf})$, use the numbers from Table 7.4.
- 2 Create a graph like Figure 7.20 that shows $P(\text{Inf} | \text{ELISA}+)$ as a function of $P(\text{Inf})$, the prior probability or seroprevalence rate. Indicate on your graph the appropriate prior and posterior probabilities for female military recruits in New Jersey, gay men in Providence, Rhode Island, and intravenous drug users in New York City.

Table 7.4
Performance of ELISA
and western blot tests.

Test	Sensitivity $P(+ \text{Inf})$	False-Negative Rate $P(- \text{Inf})$	Specificity $P(- \text{Not Inf})$	False-Positive Rate $P(+ \text{Not Inf})$
ELISA	0.997	0.003	0.985	0.015
Western blot	0.993	0.007	0.916	0.084

- 3 Repeat Questions 1 and 2, but this time find the probability of being infected given a negative ELISA result, $P(\text{Inf} | \text{ELISA}-)$.
- 4 Calculate the probability of being infected given a positive ELISA test result and a positive Western blot, $P(\text{Inf} | \text{ELISA}+, \text{WB}+)$ and the probability of being infected given a positive ELISA and negative Western blot, $P(\text{Inf} | \text{ELISA}+, \text{WB}-)$. [*Hint: All of the information in the case regarding the performance of the Western blot assumes a positive ELISA. The calculations required here can be done with Bayes' theorem, using as the prior probability $P(\text{Inf} | \text{ELISA}+)$, the quantity calculated in Question 1.*]
- 5 Create graphs like those for Questions 2 and 3 that show $P(\text{Inf} | \text{ELISA}+, \text{WB}+)$ and $P(\text{Inf} | \text{ELISA}+, \text{WB}-)$ as functions of the prior probability $P(\text{Inf})$. [*Hint: Note that this prior probability enters the calculations through $P(\text{Inf} | \text{ELISA}+)$.*]
- 6 Some public health officials have called for widespread testing for HIV infection. Certainly there is considerable value to society in identifying HIV carriers, although there are costs inherent in incorrect diagnoses. For example, suppose that you were forced to be tested, and both ELISA and Western blot tests gave a positive result. Imagine that a later tissue culture revealed no HIV exposure. Thus, for some time you would have been falsely labeled as an AIDS carrier. On the other hand, suppose that the tests had been falsely negative. Then you may have engaged in risky behavior under the assumption that you had no infection. Discuss the social trade-offs of costs and benefits that are involved in using imperfect screening tests for AIDS.

Sources: This case study was prepared using several publications available from the Centers for Disease Control National AIDS Clearinghouse, P.O. Box 6003, Rockville, MD 20849. The key publications used were "Surgeon General's Report to the American Public on HIV Infection and AIDS" (1994); "Serologic Testing for HIV-1 Antibody—United States, 1988 and 1989"; *Morbidity and Mortality Weekly Report* (1990), 39, 380-383 (abstracted in *Journal of the American Medical Association*, July 11, 1990, 171-173); and "National HIV Serosurveillance Summary: Results Through 1992," HIV/NCID/11-93/036.

DISCRIMINATION AND THE DEATH PENALTY

Is there a relationship between the race of convicted defendants in murder trials and the imposition of the death penalty on such defendants? This question has been debated extensively, with one side claiming that white defendants are given the death sentence much less frequently than non whites. This case can help you understand one reason for the debate. Table 7.5 shows information regarding 326 cases.

Questions

- 1 On the basis of Table 7.5, estimate $P(\text{Death Penalty} \mid \text{Defendant White})$ and $P(\text{Death Penalty} \mid \text{Defendant Black})$. What is your conclusion on the basis of these calculations?
- 2 Table 7.6 shows the same data disaggregated on the basis of the race of the victim. Use Table 7.6 to estimate the following probabilities:
 $P(\text{Death Penalty} \mid \text{Defendant White, Victim White})$
 $P(\text{Death Penalty} \mid \text{Defendant Black, Victim White})$
 $P(\text{Death Penalty} \mid \text{Defendant White, Victim Black})$
 $P(\text{Death Penalty} \mid \text{Defendant Black, Victim Black})$

Now what is your conclusion?

- 3 Explain the apparent contradiction between your answers to Questions 1 and 2.

Table 7.5
Death-penalty and racial status for 326 convicted murderers.

Race of Defendant	Death Penalty Imposed		Total Defendants
	Yes	No	
White	19	141	160
Black	<u>17</u>	<u>149</u>	<u>166</u>
Total	36	290	326

Table 7.6
Death-penalty and racial status for 326 convicted murderers, disaggregated by victim's race.

Race of Victim	Race of Defendant	Death Penalty Imposed		Total Defendants
		Yes	No	
White	White	19	132	151
	Black	<u>11</u>	<u>52</u>	<u>63</u>
Total		30	184	214
Black	White	0	9	9
	Black	<u>6</u>	<u>97</u>	<u>103</u>
Total		6	<u>106</u>	<u>112</u>
Total		36	290	326

REFERENCES

Probability basics appear in a wide variety of books on probability and statistics, presenting the material at various levels of sophistication. Good lower-level introductions can be found in elementary statistics textbooks such as McClave and Benson (1988), Mendenhall et al. (1989), Sincich (1989), and Wonnacott and Wonnacott (1984). Two excellent resources written at higher levels are Olkin, Gleser, and Derman (1980) and Feller (1968).

All decision-analysis textbooks seem to have at least one example about oil wildcatting. True, this is the quintessential decision-analysis problem, and it includes many sources of uncertainty, concerns about attitudes toward risk, and opportunities for gathering information. But many problems have these characteristics. Probably the real reason for the oil-wildcatting scenario is that, in 1960, C. J. Grayson published one of the first applied dissertations using decision theory, and its area of application was oil drilling. Decision theorists ever since have used oil drilling as an example!

Feller, W. (1968) *An Introduction to Probability Theory and Its Applications*, Vol. 1, 3rd ed. New York: Wiley.

Grayson, C. J. (1960) *Decisions under Uncertainty: Drilling Decisions by Oil and Gas Operators*. Cambridge, MA: Division of Research, Harvard Business School.

McClave, J. T., and P. G. Benson (1988) *Statistics for Business and Economics*, 4th ed. San Francisco: Dellen.

Mendenhall, W., J. Reinmuth, and R. Beaver (1989) *Statistics for Management and Economics*, 6th ed. Boston: PWS-KENT.

Olkin, I., L. J. Gleser, and C. Derman (1980) *Probability Models and Applications*. New York: Macmillan.

Sincich, T. (1989) *Business Statistics by Example*, 3rd ed. San Francisco: Dellen.

Wonnacott, T. H., and R. J. Wonnacott (1984) *Introductory Statistics for Business and Economics*. New York: Wiley.

EPILOGUE John Hinckley Hinckley's defense attorney was not permitted to introduce the CAT scan of Hinckley's brain. In spite of this, the jury's verdict found Hinckley "not guilty by reason of insanity" on all counts, and he was committed to Saint Elizabeth's Hospital in Washington, D.C. The trial caused a substantial commotion among the public, many people viewing the insanity plea and the resulting verdict as a miscarriage of justice. Because of this, some lawmakers initiated efforts to tighten legal loopholes associated with the insanity plea.

AIDS New diagnostic tests for AIDS are under continual development as research on this frightening disease continues. For example, in December 1994, the U.S. Food and Drug administration approved an AIDS diagnostic test that uses saliva instead of blood. This test may be easier to use, and the hope is that more people will be willing to be tested. The ELISA and Western blot tests described in the case study are typical of medical diagnostic tests in general, and the analysis performed shows how to evaluate such tests.

Subjective Probability

All of us are used to making judgments regarding uncertainty, and we make them frequently. Often our statements involve informal evaluations of the uncertainties that surround an event. Statements such as "The weather is likely to be sunny today," "I doubt that the Democrats will win the next presidential election," or "The risk of cancer from exposure to cigarette smoke is small" all involve a personal, subjective assessment of uncertainty at a fundamental level. As we have seen, subjective assessments of uncertainty are an important element of decision analysis. A basic tenet of modern decision analysis is that subjective judgments of uncertainty can be made in terms of probability. In this chapter we will explore how to make such judgments and what they imply.

Although most people can cope with uncertainty informally, perhaps, it is not clear that it is worthwhile to develop a more rigorous approach to measure the uncertainty that we feel. Just how important is it to deal with uncertainty in a careful and systematic way? The following vignettes demonstrate the importance of uncertainty assessments in a variety of public-policy situations.

UNCERTAINTY AND PUBLIC POLICY

Fruit Frost Farmers occasionally must decide whether to protect a crop from potentially damaging frost. The decision must be made on the basis of weather forecasts that often are expressed in terms of probability (U.S. National Weather Service

is responsible for providing these forecasts). Protecting a crop can be costly, but less so than the potential damage. Because of such potential losses, care in assessing probabilities is important.

Earthquake Prediction Geologists are beginning to develop ways to assess the probability of major earthquakes for specific locations. In 1988 the U.S. Geological Survey published a report that estimated a 0.60 probability of a major earthquake (7.5-8 on the Richter scale) occurring in Southern California along the southern portion of the San Andreas Fault within the next 30 years. Such an earthquake could cause catastrophic damage in the Los Angeles metropolitan area.

Environmental Impact Statements Federal and state regulations governing environmental impact statements typically require assessments of the risks associated with proposed projects. These risk assessments often are based on the probabilities of various hazards occurring. For example, in projects involving pesticides and herbicides, the chances of cancer and other health risks are assessed.

Public Policy and Scientific Research Often scientists learn of the possible presence of conditions that may require action by the government. But action sometimes must be taken without absolute certainty that a condition exists. For example, scientists in 1988 reported that the earth had begun to warm up because of the "greenhouse effect," a condition presumably resulting from various kinds of pollution and the destruction of tropical forests. James Hansen of NASA expressed his beliefs in probabilistic terms, saying that he was 99% certain that the greenhouse effect was upon us.

Medical Diagnosis Many physicians in hospital intensive-care units (ICUs) have access to a complex computer system known as APACHE III (Acute Physiology, Age, and Chronic Health Evaluation). Based on information about a patient's medical history, condition, treatment, and lab results, APACHE III evaluates the patient's risk as a probability of dying either in the ICU or later in the hospital.

Some of the above examples include more complicated and more formal probability assessments as well as subjective judgments. For example, the National Weather Service forecasts are based in part on a large-scale computer model of the global atmospheric system. The computer output is just one bit of information used by a forecaster to develop an official forecast that involves his or her subjective judgment regarding the uncertainty in local weather. Some risk assessments are based on cancer studies performed on laboratory animals. The results of such studies must be extrapolated subjectively to real-world conditions to derive potential effects on humans. Because of the high stakes involved in these examples and others, it is important for policy makers to exercise care in assessing the uncertainties they face.

At a reduced scale, personal decisions also involve high stakes and uncertainty. Personal investment decisions and career decisions are two kinds of decisions that typically involve substantial uncertainty. Perhaps even harder to deal with are per-

sonal medical decisions in which the outcomes of possible treatments are not known in advance. If you suffer from chronic chest pain, would you undergo elective surgery in an attempt to eliminate the pain? Because of the risks associated with open heart surgery, this decision must be made under a condition of uncertainty. You would want to think carefully about your chances on the operating table, considering not only statistics regarding the operation but also what you know about your own health and the skills of your surgeon and the medical staff.

Probability: A Subjective Interpretation

Many introductory textbooks present probability in terms of long-run frequency. For example, if a die is thrown many times, it would land with the five on top approximately one-sixth of the time; thus, the probability of a five on a given throw of the die is one-sixth. In many cases, however, it does not make sense to think about probabilities as long-run frequencies. For example, in assessing the probability that the California condor will be extinct by the year 2010 or the probability of a major nuclear power plant failure in the next 10 years, thinking in terms of long-run frequencies or averages is not reasonable because we cannot rerun the "experiment" many times to find out what proportion of the times the condor becomes extinct or a power plant fails. We often hear references to the chance that a catastrophic nuclear holocaust will destroy life on the planet. Let us not even consider the idea of a long-run frequency in this case!

Even when a long-run frequency interpretation might seem appropriate, there are times when an event has occurred, but we remain unsure of the final outcome. For example, consider the following:

- 1 You have flipped a coin that has landed on the floor. Neither you nor anyone else has seen it. What is the probability that it is heads?
- 2 What is the probability that Oregon beat Stanford in their 1970 football game?
- 3 What is the probability that the coin that was flipped at the beginning of that game came up heads?
- 4 What is the probability that Millard Fillmore was President in 1850?

For most of us the answers to these questions are not obvious. In every case the actual event has taken place. But unless you know the answer, you are uncertain.

The point of this discussion is that we can view uncertainty in a way that is different from the traditional long-run frequency approach. In Examples 1 and 3 above, there was a random event (flipping the coin), but the randomness is no longer in the coin. You are uncertain about the outcome because you do not know what the outcome was; the uncertainty is in your mind. In all of the examples, the uncertainty lies in your own brain cells. When we think of uncertainty and probability in this way, we are adopting a subjective interpretation, with a probability representing an individual's *degree of belief* that a particular outcome will occur.

Decision analysis requires numbers for probabilities, not phrases such as "common," "unusual," "toss-up," or "rare." In fact, there is considerable evidence from the cognitive psychologists who study such things that the same phrase has different connotations to different people in different contexts. For example, in one study (Beyth-Marom, 1982), the phrase "there is a non-negligible chance ..." was given specific probability interpretations by individuals that ranged from below 0.36 to above 0.77. Furthermore, it may be the case that we interpret such phrases differently depending on the context. The phrase "a slight chance that it will rain tomorrow" may carry a very different probability interpretation than the phrase "a slight chance that the space shuttle will explode."

The problems with using verbal representations of uncertainty can be seen in the following financial-accounting policy.

ACCOUNTING FOR CONTINGENT LOSSES

When a company prepares its financial statements, in some instances it must disclose information about possible future losses. These losses are called *contingent* because they may or may not occur contingent on the outcome of some future event (e.g., the outcome of a lawsuit or a subsidiary's default on a loan that the parent company guaranteed.) In their *Statement of Financial Accounting Standards No. 5, "Accounting for Contingencies,"* the Financial Accounting Standards Board provides guidelines for different accounting treatments of accounting losses, depending on whether the contingent loss is judged to be "probable," "remote," or "reasonably possible." In defining these verbal terms of uncertainty, "probable" is taken to mean that the future event is likely to occur. Likewise, "remote" means that the chance of the event will occur is slight. Finally, "reasonably possible" means that the chance of the event occurring is somewhere between slight and likely. [Source: Financial Accounting Standards Board (1991). *Original Pronouncements: Accounting Standards. Vol. 1: FASB Statement of Standards.* Homewood, IL: Irwin.]

In this example, verbal terms of uncertainty are defined using other verbal terms; no precise guidance is provided. The wording of the standard moves us from wondering about the meaning of "probable" and "remote" to a concern with "likely to occur" and "slight." How much more straightforward the accountant's job would be if the standard specified the different degrees of risk in terms of quantitative judgments made by a knowledgeable person!

One of the main topics of this chapter is how to assess probabilities — the numbers — that are consistent with one's subjective beliefs. Of the many concepts in decision analysis, the idea of subjective probability is one that seems to give students trouble. Some are uncomfortable assessing a degree of belief, because they think there must be a "correct" answer. There are no correct answers when it comes to sub-

jective judgment; different people have different degrees of belief and hence will assess different probabilities.

If you disagree with a friend about the probability that your favorite team will win a game, do you try to persuade your friend that your probability is better? You might discuss different aspects of the situation, such as which team has the home advantage, which players are injured, and so forth. But even after sharing your information, the two of you still might disagree. Then what? You might place a bet. For many people, betting reflects their personal subjective probabilities. Some people bet on anything even if the outcome is based on some "objectively" random event (flipping a coin, playing cards, and so on). One of the most common bets might be investing — betting — in the stock market. For example, you might be willing to purchase a stock now if you think its value is likely to increase.

We will begin with the assessment of probabilities. We will show how you can view different situations in terms of the bets you might place involving small cash amounts or in terms of hypothetical lotteries. Following the discussion of the assessment of discrete probabilities, we will see how to deal with continuous probability distributions (the fan or crescent shape in the range-of-risk dilemma that was introduced in Chapter 3). Special psychological phenomena are associated with probability assessment, and we will explore the cognitive heuristics that we tend to use in probability assessment. The last two sections discuss procedures for decomposing probability assessments and what it means to be "coherent" in assessing probabilities.

Assessing Discrete Probabilities

There are three basic methods for assessing probabilities. The first is simply to have the decision maker assess the probability directly by asking, "What is your belief regarding the probability that event such and such will occur?" The decision maker may or may not be able to give an answer to a direct question like this and may place little confidence in the answer given.

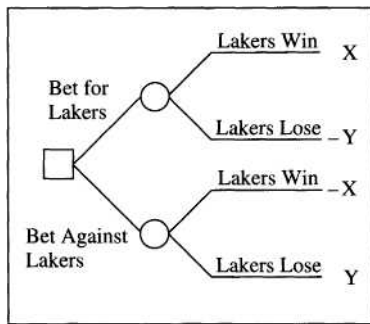
The second method is to ask about the bets that the decision maker would be willing to place. The idea is to find a specific amount to win or lose such that the decision maker is indifferent about which side of the bet to take. If he or she is indifferent about which side to bet, then the expected value of the bet must be the same regardless of which is taken. Given these conditions, we can then solve for the probability.

As an example, suppose that the Los Angeles Lakers are playing the Detroit Pistons in the NBA finals this year. We are interested in finding the decision maker's probability that the Lakers will win the championship. The decision maker is willing to take either of the following two bets:

- | | |
|--------------|---|
| Bet 1 | Win \$X if the Lakers win.
Lose \$Y if the Lakers lose. |
| Bet 2 | Lose \$X if the Lakers win.
Win \$ Y if the Lakers lose. |

Figure 8.1

Decision-tree representation for assessing subjective probability via the betting method. The assessor's problem is to find X and Y so that he or she is indifferent about betting for or against the Lakers.



In these bets, X and Y can be thought of as the amounts that each person puts into the "pot." The winner of the bet takes all of the money therein. Bets 1 and 2 are symmetric in the sense that they are opposite sides of the same bet. Figure 8.1 displays the decision tree that the decision maker faces.

If the decision maker is indifferent between Bets 1 and 2, then in his or her mind their expected values must be equal:

$$\begin{aligned}
 & X P(\text{Lakers Win}) - Y [1 - P(\text{Lakers Win})] \\
 & = -X P(\text{Lakers Win}) + Y [1 - P(\text{Lakers Win})]
 \end{aligned}$$

which implies that

$$2 \{X P(\text{Lakers Win}) - Y [1 - P(\text{Lakers Win})]\} = 0$$

We can divide through by 2 and expand the left-hand side to get X

$$P(\text{Lakers Win}) - Y + Y P(\text{Lakers Win}) = 0$$

Collecting terms gives

$$(X + Y) P(\text{Lakers Win}) - Y = 0$$

which reduces to

$$P(\text{Lakers Win}) = \frac{Y}{X + Y}$$

For example, a friend of yours might be willing to take either side of the following bet:

Win \$2.50 if the Lakers win.

Lose \$3.80 if the Lakers lose.

His subjective probability that the Lakers win, as implied by his betting behavior, is $3.80 / (2.50 + 3.80) = 0.603$.

Finding the bet for which a decision maker would be willing to take either side is fairly straightforward. Begin by offering a bet that is highly favorable to one side or the other, and note which side of the bet she would take. Then offer a bet that favors

the opposite side, and ask which side of this new bet she would prefer. Continue offering bets that first favor one side and then the other, gradually adjusting the payoffs on each round. By adjusting the bet appropriately, making it more or less attractive depending on the response to the previous bet, the indifference point can be found.

The betting approach to assessing probabilities appears straightforward enough, but it does suffer from a number of problems. First, many people simply do not like the idea of betting (even though most investments can be framed as a bet of some kind). For these people, casting the judgment task as a bet can be distracting. Most people also dislike the prospect of losing money; they are *risk-averse*. Thus, the bets that are considered must be for small enough amounts of money that the issue of risk aversion does not arise, but some people may be risk-averse even for very small amounts. Finally, the betting approach also presumes that the individual making the bet cannot make any other bets on the specific event (or even related events). That is, the individual cannot protect himself or herself from losses by "hedging" one bet with another.

To get around the problems with direct assessment or with the betting approach, a third approach adopts a thought-experiment strategy in which the decision maker compares two lotterylike games, each of which can result in a Prize (A or B). For convenience, set it up so that the decision maker prefers A to B. (Prize A might be a fully paid two-week vacation in Hawaii, and Prize B a coupon for a free beer.) We would ask the decision maker to compare the lottery

Win Prize A if the Lakers win.

Win Prize B if the Lakers lose,

with the lottery

Win Prize A with known probability p .

Win Prize B with probability $1 - p$.

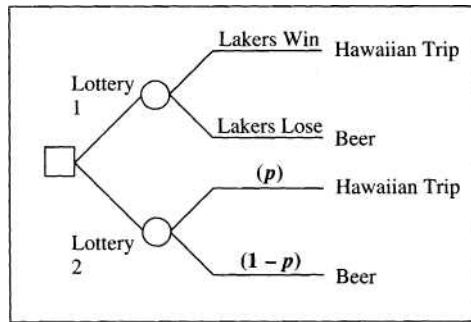
The decision-tree representation is shown in Figure 8.2. The second lottery is called the *reference lottery*, for which the probability mechanism must be well specified. A typical mechanism is drawing a colored ball from an urn in which the proportion of colored balls is known to be p . Another mechanism is to use a "wheel of fortune" with a known area that represents "win"; if the wheel were spun and the pointer landed in the win area, then the decision maker would win Prize A.

Once the mechanism is understood by the decision maker, the trick is to adjust the probability of winning in the reference lottery until the decision maker is indifferent between the two lotteries. Indifference in this case means that the decision maker has no preference between the two lotteries, but slightly changing probability p makes one or the other lottery clearly preferable. If the decision maker is indifferent, then her subjective probability that the Lakers win must be the p that makes her indifferent.

How do we find the p that makes the decision maker indifferent? The basic idea is to start with some p_1 and ask which lottery she prefers. If she prefers the reference lottery, then p_1 must be too high; she perceives that the chance of winning in the reference

Figure 8.2

Decision-tree representation for assessing subjective probability with equivalent-lottery method. The assessor's problem is to find a value of p so that the two lotteries are equivalent.



lottery is higher. In this case, choose p_2 less than p_1 and ask her preference again. Continue adjusting the probability in the reference lottery until the indifference point is found. It is important to begin with extremely wide brackets and to converge on the indifference probability slowly. Going slowly allows the decision maker plenty of time to think hard about the assessment, and she probably will be much happier with the final result than she would be if rushed. This is an important point for finding your own probability assessments. Be patient, and home in on your indifference point gradually.

The wheel of fortune is a particularly useful way to assess probabilities. By changing the setting of the wheel to represent the probability of winning in the reference lottery, it is possible to find the decision maker's indifference point quite easily. Furthermore, the use of the wheel avoids the bias that can occur from using only "even" probabilities (0.1, 0.2, 0.3, and so on). With the wheel, a probability can be any value between 0 and 1. Some decision-analysis programs include wheels as probability-assessment aids. Figures 8.3 and 8.4 show DPL's probability-assessment wheel in the context of the Texaco-Penzoil example from Chapter 4.

The lottery-based approach to probability assessment is not without its own shortcomings. Some people have a difficult time grasping the hypothetical game that they are asked to envision, and as a result they have trouble making assessments. Others dislike the idea of a lottery or carnival-like game. These same people, though, do make trade-offs with their own money whenever they purchase insurance, invest in a small business, or purchase shares of stock in a company. In some cases it may be better to recast the assessment procedure in terms of risks that are similar to the kinds of financial risks an individual might take.

The last step in assessing probabilities is to check for consistency. Many problems will require the decision maker to assess several interrelated probabilities. It is important that these probabilities be consistent among themselves; they should obey the probability laws introduced in Chapter 7. For example, if $P(A)$, $P(B | A)$, and $P(A \text{ and } B)$ were all assessed, then it should be the case that

$$P(A)P(B | A) = P(A \text{ and } B)$$

If a set of assessed probabilities is found to be inconsistent, then the decision maker should reconsider and modify the assessments as necessary to achieve consistency.

Figure 8.3
The Texaco reaction chance node in DPL.




Figure 8.4
DPL's probability-assessment wheel for the Texaco reaction node. The user can change the proportion of the wheel that corresponds to any of the events. Clicking on the "OK" button returns the user to the screen in Figure 8.3 with appropriate probabilities entered on the branches of the chance node.

Assessing Continuous Probabilities

The premise of this chapter is that it always is possible to model a decision maker's uncertainty using probabilities. How would this be done in the case of an uncertain but continuous quantity? We already have learned how to assess individual probabilities; we will apply this technique to assess several cumulative probabilities and then use these to plot a rough CDF. We will discuss two strategies for assessing a subjective CDF.

Let us reexamine the example of the movie star's age that we introduced in Chapter 7 (page 235). As you recall, the problem was to derive a probability distribution representing a probability assessor's uncertainty regarding a particular movie star's age. In that example, several probabilities were found, and these were transformed into cumulative probabilities.

A typical cumulative assessment would be to assess $P(\text{Age} < a)$, where a is a particular value. For example, consider $P(\text{Age} < 46)$. The outcome "Age < 46" is an outcome just like any other, and so a decision maker could assess the probability of this event by using any of the three techniques discussed above. For example, a wheel of fortune might be used as an assessment aid to find the probability p that would make the decision maker indifferent between the two lotteries shown in Figure 8.5.

Using this technique to find a CDF amounts to assessing the cumulative probability for a number of points, plotting them, and drawing a smooth curve through the plotted points. Suppose the following assessments were made:

$$\begin{aligned} P(\text{Age} \leq 29) &= 0.00 \\ P(\text{Age} \leq 40) &= 0.05 \\ P(\text{Age} \leq 44) &= 0.50 \\ P(\text{Age} \leq 50) &= 0.85 \\ P(\text{Age} \leq 65) &= \\ &1.00 \end{aligned}$$

Plotting these cumulative probabilities would result in the graph that we originally drew in Figure 7.9, and which is reproduced here as Figure 8.6.

Figure 8.5
Decision-tree representation for assessing $P(\text{Age} \leq 46)$ in movie star example.

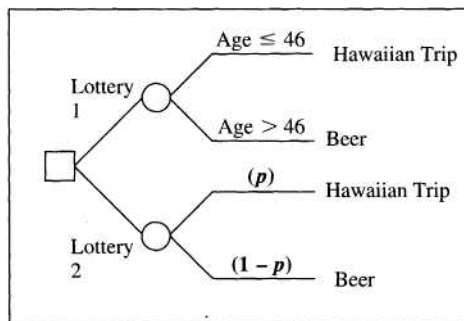



Figure 8.6
Cumulative distribution function for movie star's age.



The strategy that we have used here is to choose a few values from the horizontal axis (some ages) and then to find the cumulative probabilities that correspond to those ages. This is a perfectly reasonable strategy for assessing a CDF. Another strategy builds up the graph the other way around. That is, we pick a few cumulative probabilities from the vertical axis and find the corresponding ages. For example, suppose we pick probability 0.35. Now we want the number of years $a_{0.35}$ such that $P(\text{Age} \leq a_{0.35}) = 0.35$. The number $a_{0.35}$ is called the 0.35 *fractile* of the distribution. In general, the p fractile of a distribution for X is the value x_p such that $P(X \leq x_p) = p$. We can see from Figure 8.6 that the 0.35 fractile of the distribution is approximately 42 years. We know from the assessments that were made that the 0.05 fractile is 40 years, the 0.50 fractile is 44 years, and the 0.85 fractile is 50 years.

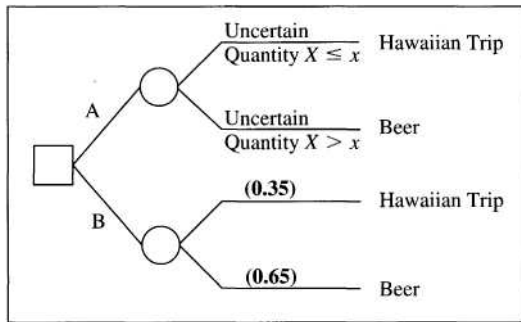
How could you go about assessing a fractile? Figure 8.7 shows a decision tree that represents the process for assessing the 0.35 fractile. In Lottery B, or the reference lottery, the probability of winning is fixed at 0.35. The assessment task is to adjust the number x in Lottery A until indifference between the lotteries is achieved. Indifference would mean that the probability of winning in Lottery A must be 0.35. Hence, there must be a 0.35 chance that X is less than or equal to the assessed x . By definition, then, x must be the 0.35 fractile of the distribution.

It is important to recognize the difference between Figures 8.5 and 8.7. In Figure 8.5 we adjusted p in the reference lottery to find indifference. To assess the 0.35 fractile in Figure 8.7, we fix the probability in the reference lottery at 0.35, and we adjust x in the upper lottery.

The term *fractile* is a general one, but other similar terms are useful for referring to specific fractiles. The idea of a *median* may be familiar. If we can find an amount such that the uncertain quantity is as likely to be above as below that amount, then we have found the median. The median is defined as the 0.50 fractile; the median for the movie star's age is 44 years. We also can speak of *quartiles*. The first quartile is an amount such that $P(X \leq \text{first quartile}) = 0.25$, or the 0.25 fractile. In our example, the first quartile appears to be around 42 years.

Figure 8.7

Decision tree for assessing the 0.35 fractile of a continuous distribution for X . The decision maker's task is to find x in Lottery A that results in indifference between the two lotteries.



Likewise, the third quartile is defined as the 0.75 fractile. The third quartile of our example is approximately 48 years because $P(\text{Age} \leq 48)$ is approximately 0.75. The second quartile is, of course, the median. Fractiles also can be expressed as *percentiles*. For example, the 90th percentile is defined as the 0.90 fractile of a distribution.

As mentioned above, we can exploit this idea of fractiles so as to assess a continuous distribution. In general, the strategy will be to select specific cumulative probabilities and assess the corresponding fractiles. The first step might be to find the uncertain quantity's extreme values. How small or large could this quantity be? Because it often is difficult, and sometimes even misleading, to think in terms of probabilities of 0 or 1, we take the 0.05 and 0.95 fractiles (or the fifth and 95th percentiles). In our movie star example, the 0.05 fractile is a value $a_{0.05}$ such that there is only a 5% chance that the movie star's age would be less than or equal to $a_{0.05}$. Likewise, the 0.95 fractile is the value $a_{0.95}$ such that there is a 95% chance that the age would be less than or equal to $a_{0.95}$. Informally, we might think of these as the smallest and largest values that the uncertain quantity could reasonably assume. Anything beyond these values would be quite surprising.

After assessing the extreme points, the median might be assessed. For example, Figure 8.8 shows the decision tree that corresponds to the assessment for the median in our example. The task would be for the decision maker to find an age a that makes the two lotteries equivalent. The value of a that leaves the decision maker indifferent is the median of the distribution and can be plotted as a point on the decision maker's subjective CDF.

Next, assess the first and third quartiles. These assessments can be made using a lottery setup similar to that in Figure 8.7. Another way to think of the quartiles is that they "split" the probability intervals above and below the median. For example, the first quartile is a value x such that the uncertain quantity is just as likely to fall below x as between x and the median.

To do this, the decision maker must find a point such that the uncertain quantity is equally likely to fall above or below this point. Having assessed the extreme points, the median, and the quartiles, we have five points on the cumulative distribu-

Figure 8.8
 Decision tree for assessing the median of the distribution for the movie star's age. The assessment task is to adjust the number of years a in Lottery A to achieve indifference.

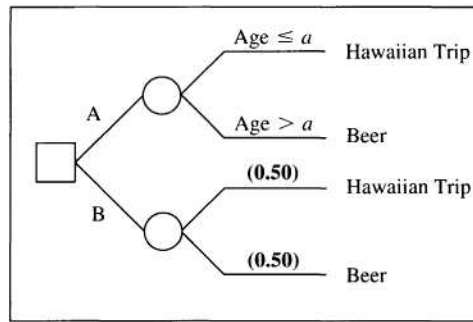



Figure 8.9
 A subjectively assessed CDF for pretzel demand.



tion function. These points can be plotted on a graph, and a smooth curve drawn through the points.

As an example, suppose you have developed a new soft pretzel that you are thinking about marketing through sidewalk kiosks. You are interested in assessing the annual demand for the pretzels as a continuous quantity. You might make the following assessments:

- 0.05 fractile for demand = 5000.
- 0.95 fractile for demand = 45,000.
- Demand is just as likely to be above 23,000 as below or equal to 23,000.
- There is a 0.25 chance that demand will be below 16,000.
- There is a 0.75 chance that demand will be below 31,000.

The last three assessments establish the median to be 23,000, the first quartile to be 16,000, and the third quartile to be 31,000. Plotting the points, we obtain the graph in Figure 8.9. A smooth curve drawn through the five points represents your subjective cumulative probability distribution of demand for the new pretzels.

Once we assess a continuous distribution, how can we use it? Apparently, our

motivation for assessing it in the first place was that we faced uncertainty in the form of a range-of-risk dilemma. In our example, we may be deciding whether to go ahead and market the pretzels. We need the probability distribution to fill out our decision tree, calculate expected values, and find an optimum choice. But how will we calculate an expected value for this subjectively assessed distribution? There are several possibilities, some of which we will explore in later chapters. For example, in the next chapter, we will see how to fit a theoretical distribution to an assessed one; we then can use the mathematical properties of the theoretical distribution. Advanced simulation and numerical integration techniques also are possible. At this point, however, we will content ourselves with some simple and useful approximation techniques.

The easiest way to use a continuous distribution in a decision tree or influence diagram is to approximate it with a discrete distribution. The basic idea is to find a few representative points in the distribution and then to assign those points specific probability values. A particularly simple approach, from Keefer and Bodily (1983), is called the *extended Pearson-Tukey method*. This three-point approximation uses the median and the 0.05 and 0.95 fractiles. Thus, one primary advantage of this method is that we can use assessments that already have been made. In assigning probabilities, the median gets probability 0.63, and the 0.05 and 0.95 fractiles each have probability 0.185. (These probabilities do not appear to admit any obvious interpretation. They are used because the resulting approximation is reasonably accurate for a wide variety of distributions.) For the pretzel-demand example, we can create the three-point discrete approximation as in Figure 8.10. The extended Pearson-Tukey method works best for approximating symmetric distributions. Given its simplicity, however, it also works surprisingly well for asymmetric distributions.

A slightly more complex approximation technique is to find *bracket medians*. Suppose that we consider an interval in which an uncertain quantity could fall: $a \leq X \leq b$. The bracket median is a value m^* between a and b such that $P(a \leq X \leq m^*) = P(m^* \leq X \leq b)$. Figure 8.11 shows how a bracket median relates to the underlying CDF. The bracket median divides the probability of the original interval in half and is associated with a cumulative probability halfway between the cumulative probabilities for a and b .

To use bracket medians, the typical approach is to break the subjective probability distribution into several equally likely intervals and then to assess the bracket median for each interval. In practice, three, four, or five intervals are typically used; the more intervals used, the better the approximation. With five intervals, the assessor would use the extreme points and the 0.20, 0.40, 0.60, and 0.80 fractiles. These would correspond to cumulative probabilities as follows:

$$\begin{aligned} P(X \leq x_{0.0}) &= 0.00 \\ P(X \leq x_{0.2}) &= 0.20 \\ P(X \leq x_{0.4}) &= 0.40 \\ P(X \leq x_{0.6}) &= 0.60 \\ P(X \leq x_{0.8}) &= 0.80 \\ P(X \leq x_{1.0}) &= 1.00 \end{aligned}$$

Figure 8.10

Replacing a continuous distribution with a three-branch discrete uncertainty node in a decision tree.

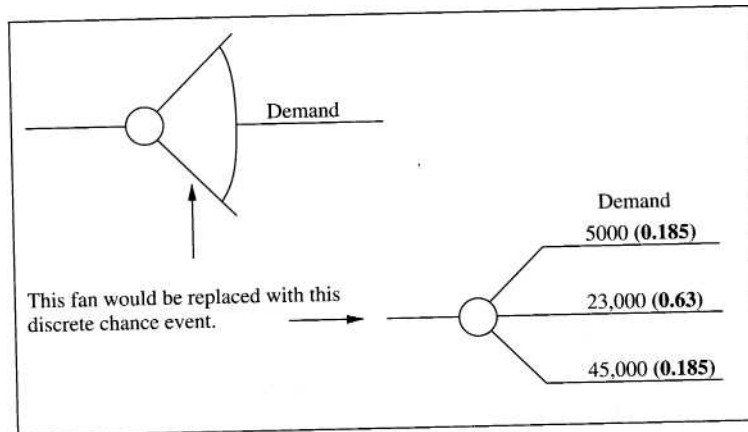



Figure 8.11

Finding the bracket median for the interval between a and b . The cumulative probabilities p and q correspond to a and b , respectively. Bracket median m^* is associated with a cumulative probability that is halfway between p and q .




Because these six points define five intervals, we need five bracket medians ($m_1 - m_5$) with the five cumulative probabilities

$$\begin{aligned}
 P(X \leq m_1) &= 0.10 \\
 P(X \leq m_2) &= 0.30 \\
 P(X \leq m_3) &= 0.50 \\
 P(X \leq m_4) &= 0.70 \\
 P(X \leq m_5) &= 0.90
 \end{aligned}$$

These bracket medians can be assessed either by using the general assessment technique for assessing fractiles above or by assessing the value m_i so that X is just as likely to be between the lower interval bound and m_i as between m_i and the upper bound.

Figure 8.12 shows the determination of five bracket medians for the pretzel-demand distribution. These bracket medians then can be used as "representative points" for their respective brackets. Thus, the discrete approximation to the original distribution would have

Figure 8.12
 Finding bracket medians for the pretzel-demand distribution.

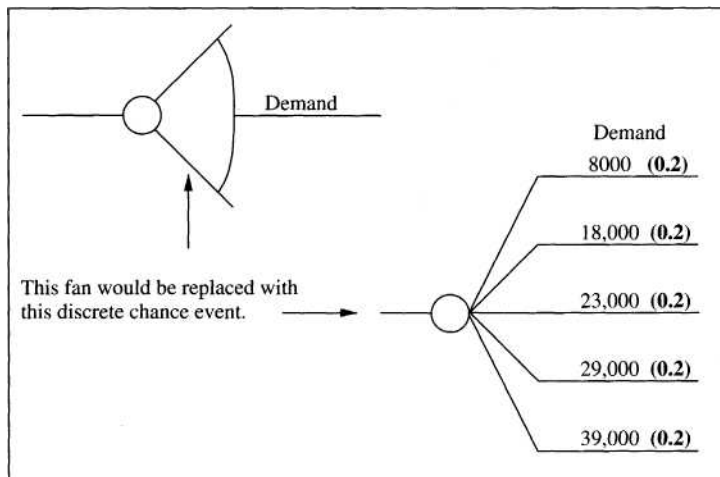


$$\begin{aligned}
 P(X = m_1 = 8) &= 0.20 \\
 P(X = m_2 = 18) &= 0.20 \\
 P(X = m_3 = 23) &= 0.20 \\
 P(X = m_4 = 29) &= 0.20 \\
 P(X = m_5 = 39) &= 0.20
 \end{aligned}$$

as shown in Figure 8.13.

Both methods that we have described for creating a discrete approximation are straightforward and work reasonably well. The extended Pearson-Tukey method has the advantage of requiring no additional assessments. The advantages of the bracket-median approach are its ability to approximate virtually any kind of distribution and the intuitive nature of the bracket-median assessments.

Figure 8.13
 Replacing a continuous distribution with bracket medians in a decision tree.



Which method should you use? Both work well, but a general strategy exists that dovetails perfectly with the sensitivity-analysis approach we described in Chapter 5. Suppose that you have constructed a tornado diagram for a decision problem and have identified several variables that are candidates for probabilistic modeling. A simple first step would be to use the extended Pearson-Tukey method with the base value as the median, and the upper and lower values as the 0.95 and 0.05 fractiles, respectively. In many cases, this may be an adequate model of the uncertainty. If not, then it may be worthwhile to use bracket medians to construct a more complete model of the uncertainty. (Additional information on discrete approximations is given in the reference section at the end of the chapter.)

Pitfalls: Heuristics and Biases

The methods presented above make probability assessment sound easy. As you probably realize, however, thinking in terms of probabilities is not easy. It takes considerable practice before one is comfortable making probability assessments. Even then, we tend to use rather primitive cognitive techniques to make our probability assessments. Tversky and Kahneman (1974) have labeled these techniques *heuristics*. In general, heuristics can be thought of as rules of thumb for accomplishing tasks. For example, an inventory-control heuristic might be to keep 10% of total yearly demand for any given item on hand. When placing an order for an item that sells some 200 units per year, one then would check the stock and order enough to have 20 units on hand. Heuristics tend to be simple, are easy to perform, and usually do not give optimal answers.

Heuristics for assessing probabilities operate in basically the same way. They are easy and intuitive ways to deal with uncertain situations, but they tend to result in probability assessments that are biased in different ways depending on the heuristic used. In this section we will look at the various heuristics and the biases they can create. Before we begin, however, consider the case of Tom W.

TOM W.

"Tom W. is of high intelligence, although lacking in true creativity. He has a need for order and clarity, and for neat and tidy systems in which every detail finds its appropriate place. His writing is rather dull and mechanical, occasionally enlivened by somewhat corny puns and by flashes of imagination of the sci-fi type. He has a strong drive for competence. He seems to have little feel and little sympathy for other people and does not enjoy interacting with others. Self-centered, he nonetheless has a deep moral sense."

The preceding personality sketch of Tom W. was written during his senior year in high school by a psychologist on the basis of projective tests. Tom W. is now a graduate student. Please rank the following nine fields of graduate specialization in order of the likelihood that Tom W. is now a graduate student in each of these fields:

- a Business administration
- b Computer science
- c Engineering
- d Humanities and education
- e Law
- f Library science
- g Medicine
- h Physical and life sciences
- i Social science and social work

Write down your rankings before you read on. [*Source: Kahneman and Tversky (1973).*]

Representativeness

If you are like most people, you wrote down your ranks on the basis of how similar the description of Tom W. is to your preconceived notions of the kinds of people in the nine different fields of study. Specifically, Tom W.'s description makes him appear to be a "nerd," and so most people think that he has a relatively high chance of being in engineering or computer science. But judging the probability of membership in a group on the basis of similarity ignores important information. There are many more graduate students in humanities and education and in social science and social work than in computer science or engineering. Information relating to the incidence or *base rate* of occurrence in the different fields is ignored, however, when we make probability judgments on the basis of similarity.

Making such judgments on similarity is one example of a kind of heuristic that Kahneman and Tversky (1973) call *representativeness*. In its most fundamental form, the representativeness heuristic is used to judge the probability that someone or something belongs to a particular category. Using the representativeness heuristic means that the judgment is made by comparing the information known about the person or thing with the stereotypical member of the category. The closer the similarity between the two, the higher the judged probability of membership in the category.

The representativeness heuristic surfaces in many different situations and can lead to a variety of different biases. For example, as in the Tom W. problem, people can be insensitive to base rates or prior probabilities. If one were to consider base rates carefully in the Tom W. problem, humanities and education may well be the most likely category. Insensitivity to sample size is another possible result of the representativeness heuristic. Sometimes termed the *law of small numbers*, people (even

scientists!) sometimes draw conclusions from highly representative small samples even though small samples are subject to considerably more statistical error than are large samples. Other situations in which representativeness operates include relying on old and unreliable information to make predictions, and making equally precise predictions regardless of the inherent uncertainty in a situation.

Misunderstanding of random processes is another phenomenon attributed to the representativeness heuristic. One of the most important aspects of this situation for managers relates to changes over time and misunderstanding the extent to which a process can be controlled. For example, Kahneman and Tversky relate their experience with a flight instructor in Israel. The instructor had been in the habit of praising students for good landings and scolding them for poor ones. He observed that after receiving praise for a good landing, a pilot's subsequent landing tended to be worse. Conversely, after a pilot received a scolding, his next landing tended to be better. The instructor concluded that scolding was effective feedback and that praise was not. In fact, this phenomenon is more easily explained by what is known as the statistical phenomenon as *regression toward the mean*. If performance or measurements are random, then extreme cases will tend to be followed by less extreme ones. Landing a jet is not an easy task, and the pilot must deal with many different problems and conditions on each landing. It is perfectly reasonable to assume that performance for any pilot will vary from one landing to the next. Regression toward the mean suggests that a good landing probably will be followed by one that is not as good, and that a poor one will most likely be followed by one that is better.

Availability

Another heuristic that people tend to use to make probability judgments is termed *availability*. According to this heuristic, we judge the probability that an event will occur according to the ease with which we can retrieve similar events from memory. As with the representativeness heuristic, availability comes into play in several ways. External events and influences, for example, can have a substantial effect on the availability of similar incidents. Seeing a traffic accident can increase one's estimate of the chance of being in an accident. Being present at a house fire can have more effect on the retrievability of fire incidents than reading about the fire in the newspaper. Furthermore, differential attention by the news media to different kinds of incidents can result in availability bias. Suppose the local newspaper plays up deaths resulting from homicide but plays down traffic deaths. To some extent, the unbalanced reporting can affect readers' judgments of the relative incidence of homicides and traffic fatalities, thus affecting the community's overall perception.

Bias from availability arises in other ways as well. For example, some situations are simply easier to imagine than others. In other cases, it may be difficult to recall contexts in which a particular event occurs. Another situation involves *illusory correlation*. If a pair of events is perceived as happening together frequently, this perception can lead to an incorrect judgment regarding the strength of the relationship between the two events.

Anchoring and Adjusting

This heuristic refers to the notion that in making estimates we often choose an initial anchor and then adjust that anchor based on our knowledge of the specific event in question. An excellent example is sales forecasting. Many people make such forecasts by considering the sales figures for the most recent period and then adjusting those values based on new circumstances. The problem is that the adjustment usually is insufficient.

The anchor-and-adjust heuristic affects the assessment of probability distributions for continuous uncertain quantities more than it affects discrete assessments. We tend to begin by assessing an anchor, say, the mean or median. Then extreme points or fractiles are assessed by adjusting away from the anchor. Because of the tendency to underadjust, most subjectively assessed probability distributions are too narrow, inadequately reflecting the inherent variability in the uncertain quantity. One of the consequences is that we tend to be overconfident, having underestimated the probability of extreme outcomes, often by a substantial amount (Capen 1976). The technique described above for assessing a subjective CDF using the median and quartiles is subject to this kind of overconfidence from anchoring and adjusting. This is one reason why we assess the 0.05 and 0.95 fractiles instead of the 0.00 and 1.00 fractiles. Assessing the 0.05 and 0.95 fractiles essentially admits that there is a remote possibility that the uncertain quantity could fall beyond these assessed points. Also, we assess the 0.05 and 0.95 fractiles before the median to reduce the tendency to anchor on a central value.

Motivational Bias

The cognitive biases described above relate to the ways in which we as human beings process information. But we also must be aware of motivational biases. Incentives often exist that lead people to report probabilities or forecasts that do not entirely reflect their true beliefs. For example, a salesperson asked for a sales forecast may be inclined to forecast low so that he will look good (and perhaps receive a bonus) when he sells more than the amount forecasted. Occasionally incentives can be quite subtle or even operate at a subconscious level. For example, some evidence suggests that weather forecasters, in assessing the probability of precipitation, persistently err on the high side; they tend to overstate the probability of rain. Perhaps they would rather people were prepared for bad weather (and were pleasantly surprised by sunshine) instead of expecting good weather and being unpleasantly surprised. Even though forecasters generally are good probability assessors and strive for accurate forecasts, their assessments may indeed be slightly affected by such implicit incentives.

Heuristics and Biases: Implications

This discussion of heuristics and biases in probability judgments sounds quite pessimistic. If people really are subject to such deficiencies in assessing probabilities, is there any hope? There is indeed. First, some evidence suggests that individuals can

learn to become good at assessing probabilities. As mentioned, weather forecasters are good probability assessors; in general, they provide accurate probabilities. For example, on those occasions when a forecaster says the probability of rain is 0.20, rain actually occurs very nearly 20% (or slightly less) of the time. Weather forecasters have three advantages; they have a lot of specialized knowledge about the weather, they make many forecasts, and they receive immediate feedback regarding the outcome. All of these appear to be important in improving probability-assessment performance.

Second, awareness of the heuristics and biases may help individuals make better probability assessments. If nothing else, knowing about some of the effects, you now may be able to recognize them when they occur. For example, you may be able to recognize regression toward the mean, or you may be sensitive to availability effects that result from unbalanced reporting in the news media. Moreover, when you obtain information from other individuals, you should realize that their judgments are subject to these same problems.

Third, the techniques we have discussed for assessing probabilities involve thinking about lotteries and chances in a structured way. These contexts are quite different from the way that most people think about uncertainty. By thinking hard about probabilities using these methods, it may be possible to avoid some heuristic reasoning and attendant biases. At the very least, thinking about lotteries provides a new perspective in the assessment process.

Finally, some problems simply cannot be addressed well in the form in which they are presented. In many cases it is worthwhile to decompose a chance event into other events. The result is that more assessments must be made, although they may be easier. In the next section we will see how decomposition may improve the assessment process.

Decomposition and Probability Assessment

In many cases it is possible to break a probability assessment into smaller and more manageable chunks. This process is known as *decomposition*. There are at least three different scenarios in which decomposition of a probability assessment may be appropriate. In this section, we will discuss these different scenarios.

In the simplest case, decomposition involves thinking about how the event of interest is related to other events. A simple example might involve assessing the probability that a given stock price increases. Instead of considering only the stock itself, we might think about its relationship to the market as a whole. We could assess the probability that the market goes up (as measured by the Dow Jones average, say), and then assess the conditional probabilities that the stock price increase given that the market increases and given that the market does not increase. Finding the probability that the stock price increases is then a matter of using the law of total probability:

$$P(\text{Stock Price Up}) = P(\text{Stock Price Up} \mid \text{Market Up}) P(\text{Market Up}) \\ + P(\text{Stock Price Up} \mid \text{Market Not Up}) P(\text{Market Not Up})$$

The reason for performing the assessment in this way is that it may be more comfortable to assess the conditional probabilities and the probability about the market rather than to assess $P(\text{Stock Price Up})$ directly. In terms of an influence diagram or a probability tree, we are adding a chance node that is relevant to the assessment of the probabilities in which we are interested. Figure 8.14 shows the decomposition of the stock-price assessment.

In the second scenario, it is a matter of thinking about what kinds of uncertain outcomes could eventually lead to the outcome in question. For example, if your car will not start, there are many possible reasons why it will not. The decomposition strategy would be to think about the chances that different possible things could go wrong and the chance that the car will not start given each of these specific underlying problems or some combination of them.

For a more complicated example that we can model, suppose that you are an engineer in a nuclear power plant. Your boss calls you into his office and explains that the Nuclear Regulatory Commission has requested safety information. One item that the commission has requested is an assessment of the probability of an accident resulting in the release of radioactive material into the environment. Your boss knows that you have had a course in decision analysis, and so you are given the job of assessing this probability.

How would you go about this task? Of course, one way is to sit down with a wheel of fortune and think about lotteries. Eventually you would be able to arrive at a probability assessment. Chances are that as you thought about the problem, however, you would realize that many different kinds of situations could lead to an accident. Thus, instead of trying to assess the probability directly, you might construct an influence diagram that includes some of the outcomes which could lead to an accident. Figure 8.15 shows the simple influence diagram that you might draw.

Figure 8.14
Decomposing the probability assessment for stock-price movement.

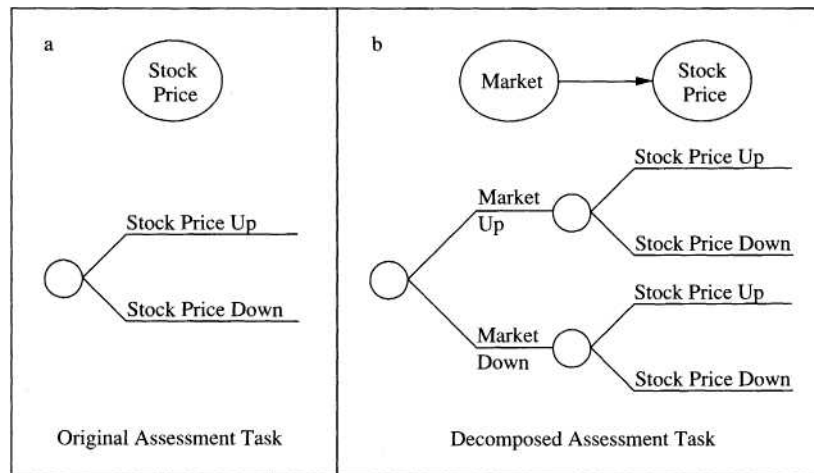



Figure 8.15
Simple influence diagram for assessing the probability of a nuclear power plant accident.



The intuition behind Figure 8.15 is that an accident could result from a failure of the cooling system or of the control system. The cooling system could either spring a leak itself, thus spilling radioactive material, or its pumps could fail, allowing the reactor core to overheat, and thus resulting in a possible accident. The control system also is critical. If the control system fails, it may become impossible to maintain safe operation of the reactor, and thus possibly result in an accident. Furthermore, the control system and the cooling system do not operate independently; they both depend on the electrical power system within the plant. Thus, failure of the electrical system would make failure of both the cooling and the control systems more likely. (Many other relationships also are possible. An influence diagram to assess the probability of an accident in an actual nuclear power plant would be considerably more complicated than Figure 8.15.)

For each of the four chance nodes in Figure 8.15, we have two possible outcomes. Because failures of both the cooling and control systems are relevant to the assessment of an accident, we have four conditional probabilities to assess. Let A denote the outcome of an accident, L the outcome of a cooling system failure, N the outcome of a control system failure, and E the outcome of an electrical system failure. The four conditional probabilities we must assess for Outcome A are $P(A | L, N)$, $P(A | L^-, N)$, $P(A | L, N^-)$ and $P(A | L^-, N^-)$. For the cooling system node, probabilities $P(L | E)$ and $P(L | E^-)$ must be assessed. Likewise, for the control system node $P(N | E)$ and $P(N | E^-)$ must be assessed. Finally, $P(E)$ must be assessed for the electrical system node.

There are nine assessments in all. Again, the reason for decomposing the assessment task into multiple assessments is that you may be more comfortable with the assessments that are required in the decomposed version. For example, you may be

able to conclude that $P(A | L^-, N^-) = 0$, or that if neither the cooling system nor the control system fails, then the probability of an accident is zero.

Assembling the probabilities in this case again is a matter of using the law of total probability, although it must be used more than once. Start out by using the law of total probability to expand $P(A)$:

$$P(A) = P(A | L, N)P(L, N) + P(A | L^-, N)P(L^-, N) \\ + P(A | L, N^-)P(L, N^-) + P(A | L^-, N^-)P(L^-, N^-)$$

Now the problem is to find $P(L, N)$, $P(L^-, N)$, $P(L, N^-)$, and $P(L^-, N^-)$. Each in turn can be expanded using the law of total probability. For example, consider $P(L, N)$:

$$P(L, N) = P(L, N | E)P(E) + P(L, N | E^-)P(E^-)$$

Now we must find $P(L, N | E)$ and $P(L, N | E^-)$. From the influence diagram, the only connection between the cooling system (L) and the control system (N) is through the electrical system. Thus, cooling and control failures are conditionally independent given the state of the power system. From the definition of conditional independence in Chapter 7, we can write

$$P(L, N | E) = P(L | E)P(N | E)$$

and

$$P(L, N | E^-) = P(L | E^-)P(N | E^-)$$

Thus, by expanding out the probabilities, it is possible to build up the probability $P(A)$ from the nine assessments and their complements.

The third scenario is related to the second. In this case, however, it is not a matter of different possible underlying causes but a matter of thinking through all of the different events that must happen before the outcome in question occurs. For example, in assessing the probability of an explosion at an oil refinery, an engineer would have to consider the chances that perhaps some critical pipe would fail, that all of the different safety measures also would fail at the same time, and that no one would notice the problem before the explosion occurred. Thus, many different individual outcomes would have to occur before the explosion. In contrast, the nuclear power plant example involved alternative paths that could lead to a failure. Of course, the second and third scenarios can be combined. That is, there may be alternative paths to a failure, each requiring that certain individual outcomes occur. This kind of analysis often is called *fault-tree analysis* because it is possible to build a tree showing the relationship of prior outcomes to the outcome in question, which often is the failure of some complicated system.

As you may have noticed in the nuclear power plant example, the probability manipulations can become somewhat complicated. Fortunately, in complicated assessment problems, computer programs can perform the probability manipulations for us. Using such a program allows us to focus on thinking hard about the assessments rather than on the mathematics.

Figure 8.16

Influence diagram for a decision analysis of alternative sites for a nuclear-waste repository. Source: Merkhofer (1987b), p. 105. Reprinted by permission of Kluwer Academic Publishers.

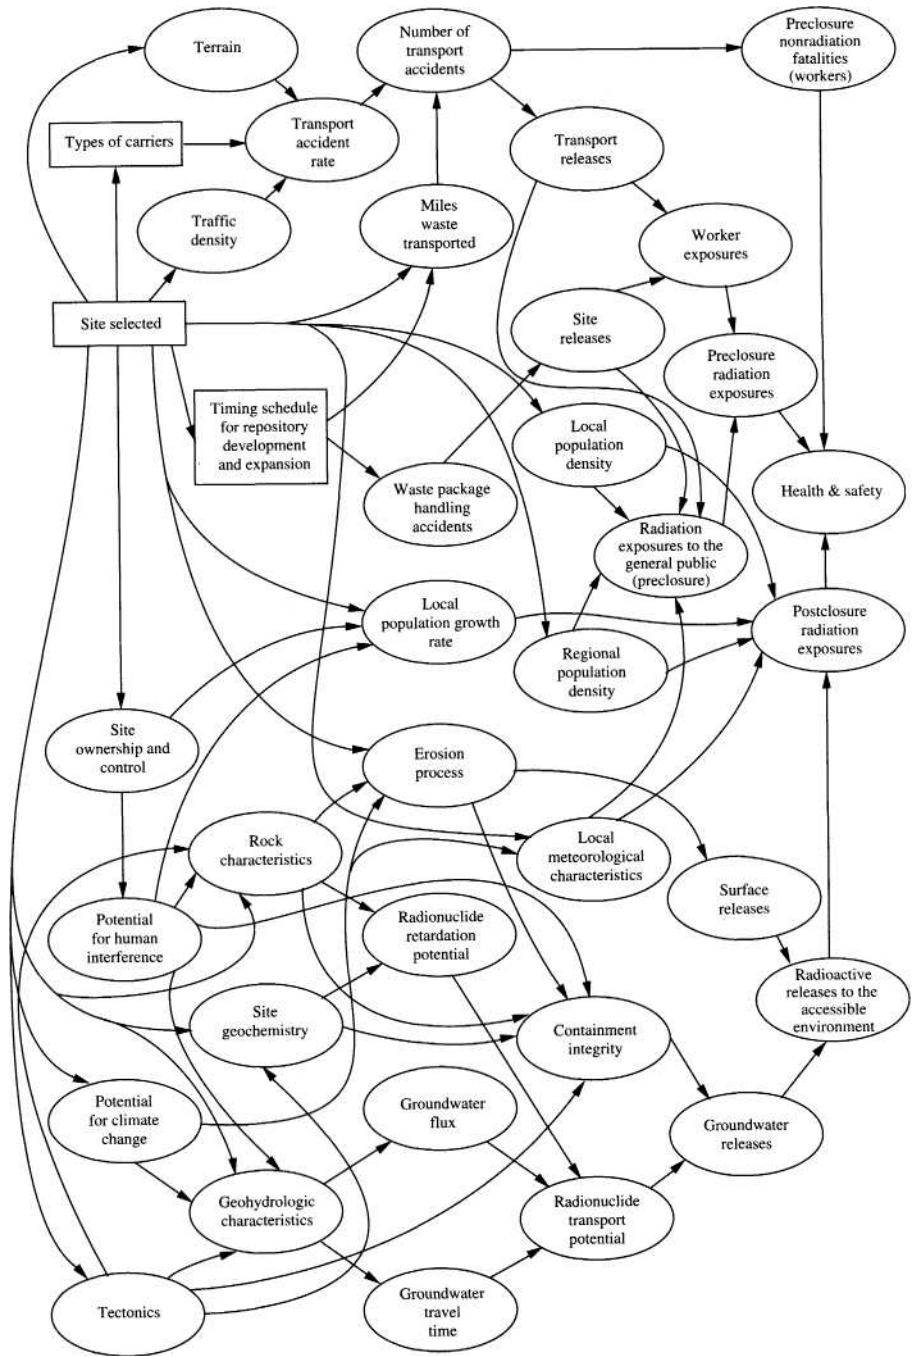
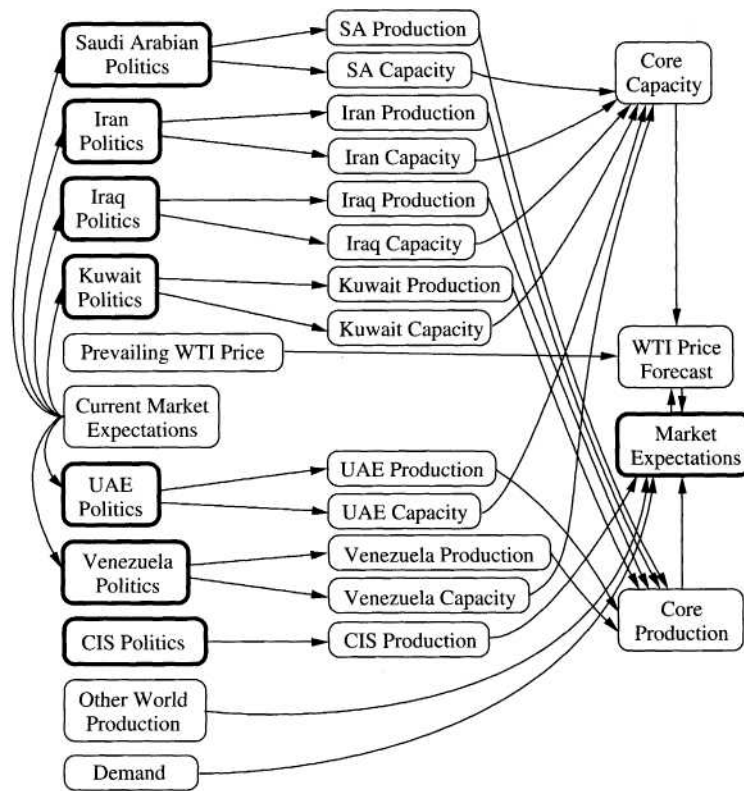


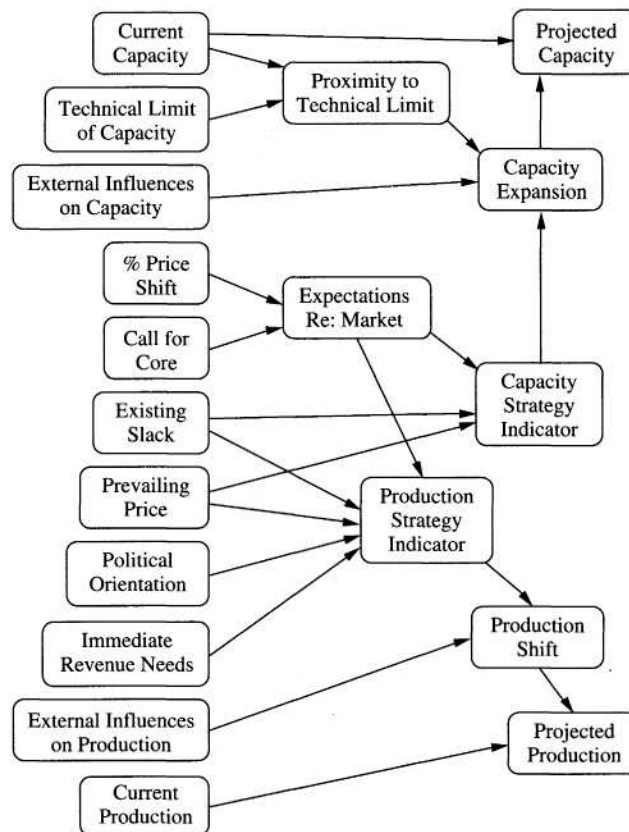
Figure 8.17
Influence diagram (or belief network) for forecasting crude oil prices. Rounded rectangles are used instead of ovals to represent chance nodes. Nodes with bold outlines represent entire submodels. *Source:* Abramson and Finizza (1995).



As with many decision-analysis techniques, there may be more than one way to decompose a probability assessment. The whole reason to use decomposition is to make the assessment process easier. The best decomposition to use is the one that is easiest to think about and that gives the clearest view of the uncertainty in the decision problem.

As indicated above, decomposition in decision-analysis assessment of probabilities is important for many reasons. Perhaps most important, though, is that it permits the development of large and complex models of uncertainty. Two examples are given in Figures 8.16, 8.17, and 8.18. The influence diagram in Figure 8.16 was constructed as part of an analysis of alternative sites for a nuclear-waste repository in the United States (Merkhofer 1987b). Figures 8.17 and 8.18 show parts of the influence diagram for a probabilistic model for forecasting crude oil prices (Abramson and Finizza, 1995). In Figures 8.17 and 8.18 the authors have used rounded rectangles to represent chance nodes. In addition, each of the nodes with bold outlines in Figure 8.17 actually represents an entire submodel. For example, Figure 8.18 shows the generic structure for each of the "Politics" nodes. In all, this model of the oil market includes over 150 chance nodes.

Figure 8.18
 Details of the
 "Politics" submodels
 in Figure 8.17.
 Source: Abramson
 and Finizza (1995).



Experts and Probability Assessment: Pulling It All Together

Our discussion thus far has taken a conventional decision-analysis approach; the decision maker makes the probability assessments that are required in the decision model. In practice, though, decisions are often quite complex, and the decision maker must rely on experts to provide information — in the form of probability assessments — regarding crucial uncertainties. For example, all of the probability assessments required in the influence diagrams in Figures 8.16, 8.17, and 8.18 were provided by experts regarding nuclear-waste technology and the world oil market, respectively.

In such complex problems, expert risk assessment plays a major role in the decision-making process. As such, the process by which the expert information was acquired must stand up to professional scrutiny, and thus policy makers who acquire and use expert information must be able to document the assessment process. In general, this is no different than standard scientific principles of data collection; scientists who run experiments or conduct surveys are expected to adhere to standards

that ensure the scientific validity of their data. Adherence to these standards is necessary if the conclusions drawn on the basis of the data are to have any validity.

The assessment of expert judgments must also adhere to standards. Not surprisingly, though, the standards for experts are quite different from those for data collection; expert judgments are not as well behaved as data collected in a laboratory. There are important parallels between the two situations, however. For example, the definition of an expert in any given situation is not always without controversy. Thus, the policy maker must be able to document and justify the expert-selection process, just as the data-collecting scientist must be able to document and justify the process by which specific data points were selected. Also, as we have mentioned, experts can be subject to numerous biases. Thus, policy makers must be able to show that the environment in which the judgments were made did as much as possible to reduce or avoid these biases. The counterpart in data collection is that the scientist must be able to show that measurements were taken without bias. If judgments from multiple experts are combined to obtain a single probability distribution, then issues of relative expertise and redundancy among the experts must be taken into account. The corresponding situation in data collection occurs when a scientist combines multiple data sets or uses results from multiple studies to draw conclusions.

Over the past 20 years, as the use of expert information has grown in importance, procedures have been developed for acquiring expert probability assessments. In general, the approach requires the creation of a *protocol* for expert assessment that satisfies the need for professional scrutiny. Thorough discussions of protocol development are found in Merkhofer (1987a) and Morgan and Henrion (1990). Although procedures vary (for example, Morgan and Henrion describe three different approaches), every assessment protocol should include the following steps.

1 Background The first step is to identify those variables for which expert assessment is needed. Although this sounds obvious, it is an important first step. Relevant scientific literature should be searched to determine the extent of scientific knowledge. The objectives of stakeholders should be examined to be sure that information is being obtained about pertinent concerns. For example, if stakeholders in an environmental-management situation care about habitat for endangered species, ecosystem viability, extraction of natural resources for economic purposes, recreation opportunities, and saving wilderness for future generations, then information on all five of these issues must be acquired. Some may require expert assessment, while others may be better addressed by conventional scientific studies. In many cases, experts may be required to make issue-specific probabilistic judgments based on state-of-the-art (but more general) scientific knowledge.

2 Identification and Recruitment of Experts Identification of appropriate experts can range from straightforward to very difficult. Often an organization can find in-house expertise. In some cases, experts must be recruited externally. Recommendations by peers (for example, through professional associations) can help to justify the selection of specific experts.

3 Motivating Experts Experts often are leery of the probability-assessment process. Typically they are scientists themselves and prefer to rely on the process of science to generate knowledge. Their opinions may or may not be "correct," and hence they hesitate to express those opinions. The fact remains, though, that a decision must be made with the limited information available, and the choice was made in Step 1 that expert opinion is the appropriate way to obtain that information. Thus, it is important to establish rapport with the experts and to engender their enthusiasm for the project.

4 Structuring and Decomposition This step might be called *knowledge exploration*. This step identifies specific variables for which judgments are needed and explores the experts' understanding of causal and statistical relationships among the relevant variables. The objective is to develop a general model (expressed, for example, as an influence diagram) that reflects the experts' thinking about the relationships among the variables. The resulting model may be an elaborate decomposition of the original problem, showing which probability distributions must be assessed conditional on other variables. The model thus gives an indication of the order in which the probability assessments must be made.

5 Probability-Assessment Training Because many experts do not have specific training in probability assessment, it is important to explain the principles of assessment, to provide information on the inherent biases in the process and ways to counteract those biases, and to give the experts an opportunity to practice making probability assessments.

6 Probability Elicitation and Verification In this step the experts make the required probability assessments, typically under the guidance of an individual trained in the probability-elicitation process. The expert's assessments are checked to be sure they are consistent (probabilities sum to 1, conditional probabilities are consistent with marginal and joint probability assessments, and so on). As part of this process, an expert may provide detailed chains of reasoning for the assessments. Doing so can help to establish a clear rationale for specific aspects of the assessed distributions (e.g., a variable's extreme values or particular dependence relationships). At the same time, encouraging a thorough examination of the expert's knowledge base can help to counteract the biases associated with the psychological heuristics of availability, anchoring, and representativeness. Thus, as output this step produces the required probability assessments and a documentation of the reasoning behind the assessments.

7 Aggregation of Experts' Probability Distributions If multiple experts have assessed probability distributions, then it may be necessary to aggregate their assessed probability distributions into a single distribution for the use of the decision maker. Another reason for aggregation is that a single distribution may be needed in a probabilistic model of a larger system. In general, two approaches are possible. One is to ask the experts themselves to generate a consensus distribution. Doing so may require considerable sharing of information, clarification of individual definitions, and possibly compromise on the parts of individuals. Unfortunately, the necessary interactions can

also lead to biases in the consensus opinion. For this reason, several methods have been proposed to control the interaction among the individual experts.

If the experts are unable to reach a consensus, either because of irreconcilable differences or because convening the group is logistically inconvenient, then one can aggregate their individual distributions using a mathematical formula. Simply averaging the distributions is a straightforward (and intuitively appealing) aggregation approach, but it ignores the relative expertise among the experts as well as the extent to which their information is redundant or dependent.

The seven steps described above give a feel for the process of obtaining expert probability assessments. In a full-blown risk assessment, this process can involve dozens of people and may take several months to complete. The following example describes a risk assessment in which expert climatologists provided probability assessments regarding possible future climate changes at the site of the proposed nuclear-waste repository in Nevada.

CLIMATE CHANCE AT YUCCA MOUNTAIN, NEVADA


The U.S. government has proposed construction of a long-term nuclear-waste repository at Yucca Mountain, about 100 miles northwest of Las Vegas, Nevada. Spent fuel rods from nuclear reactors around the nation would be sealed in large steel casks, shipped to Yucca Mountain, and stored in a large cavern carved out of solid rock about 300 meters below the surface. When full, the repository will be sealed, but the nuclear waste will remain dangerous for millennia. Thus, requirements for licensing the facility include showing that the repository can safely contain the radioactive material for 10,000 years.

One of several risks that the repository faces is that the local climate is expected to change over time as the global climate changes. In broad-brush terms, one might expect that increased human activity will lead to some warming but that natural climate cycles will lead to a global cooling or even another ice age. Climatologists have studied general climate trends over long periods of time and have produced global climate-change forecasts. What does the future hold for Yucca Mountain, Nevada, though? One of the attractions of Yucca Mountain is the dry climate; it experiences an annual average precipitation of about 15 centimeters. If the future climate changes enough, groundwater could enter the repository. How much future precipitation is likely at Yucca Mountain over the next 10,000 years?


To address this question, the Center for Nuclear Waste Repository Analyses (CNWRA) undertook a project to obtain the opinions of expert climatologists. The seven steps described above helped to provide a framework for the project, during which the experts assessed subjective probability distributions for a variety of different climatological variables (average annual rainfall, average temperature, amount of cloud cover, and others) at several different points in time over the next 10,000 years.

Although the exceptionally long forecast horizon suggests that forecasting will be difficult if not impossible, the experts found this to be an interesting project, and they were able to develop carefully thought-out rationales for different scenarios and assess probability distributions that reflected their reasoning. In general, the experts agreed that the Yucca Mountain area would warm up slightly in the short term due to atmospheric buildup of greenhouse gases such as carbon dioxide and water vapor, followed by slight cooling as the earth enters a "mini ice age," but they disagreed regarding the extent and persistence of the global-warming effect. This disagreement shows up, for example, in Figure 8.19a, which shows the medians of the experts' CDFs for average temperature traced over time. Some experts think the effect of global warming will be short-lived, followed quickly by global cooling, while others believe it will last longer and be more pronounced. This disagreement is also seen in their assessed CDFs. Figures 8.19b and c, for example, show the experts' CDFs for change in average temperature (°C) and change in average annual precipitation (mm) 3000 years in the future.

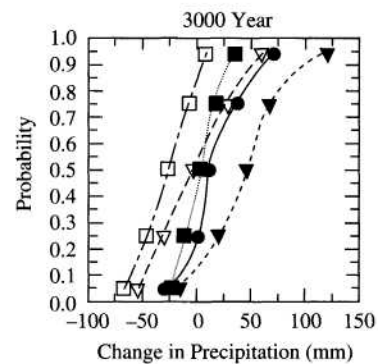
Figure 8.19
 Some assessments of climate change at Yucca Mountain, (a) Medians of CDFs for average temperature over time, (b) CDFs for change in average temperature 3000 years in future, (c) CDFs for change in average annual precipitation 3000 years in future. Source: DeWispelare et al. (1993). *Expert elicitation of future climate in the Yucca Mountain vicinity*, Technical Report CNWRA 93-016. Prepared for Nuclear Regulatory Commission. San Antonio, TX: Southwest Research Institute.



(a)



(b)



(c)

Legend: ● Expert 1 ∇ Expert 2 ▼ Expert 3 □ Expert 4 ■ Expert 5

The Yucca Mountain climate forecast is just one example of a probabilistic risk analysis that assessed expert opinions in the form of probability distributions. Other examples of such analyses involve sales forecasting, research-and-development decisions, analysis of health risks due to ozone, policy making regarding hazardous waste, power plant design analysis, and many others. Risk analyses are performed by both government agencies and private corporations and are becoming more important over time as decision makers, both public and private, face important decisions with substantial uncertainty and imperfect information.

Coherence and the Dutch Book (Optional)

Subjective probabilities must obey the same postulates and laws as so-called objective probabilities. That is, probabilities must be between 0 and 1, probabilities for mutually exclusive and collectively exhaustive sets of events must add to 1, and the probability of any event occurring from among a set of mutually exclusive events must be the sum of the probabilities of the individual events. All of the other properties of probabilities follow from these postulates, and subjective probabilities thus also must obey all of the probability "laws" and formulas that we studied in Chapter 7.

If an individual's subjectively assessed probabilities obey the probability postulates, the person is said to be *coherent*. What benefit is there in being coherent? The mathematician de Finetti proved a famous theorem, the Dutch Book Theorem, which says that if a person is not coherent, then it is possible to set up a Dutch book against him or her. A Dutch book is a series of bets that guarantees your opponent will lose and you will win.

For example, suppose that it is nearly the end of the NBA season, and that the Los Angeles Lakers and the Boston Celtics are playing in the finals. A friend of yours says that the probability is 0.4 that the Lakers will win and 0.5 that the Celtics will win. You look at this statement and point out that his probabilities only add to 0.9. Is there some other possible outcome? He sullenly replies that those are his probabilities, and that there is nothing wrong with them.

Let us find what is wrong. If those really are his probabilities, then he should be willing to agree to the following bets:

- Bet 1** He wins \$40 if the Lakers lose.
 You win \$60 if the Lakers win.
- Bet 2** You win \$50 if the Celtics win.
 He wins \$50 if the Celtics lose.

Note that, according to his stated probabilities, his expected value for each bet is 0. In Bet 1 he has a 0.6 chance of winning \$40 and a 0.4 chance of losing \$60. His expected value is $0.6(40) - 0.4(60) = 0$. In Bet 2, his expected value is $0.5(50) - 0.5(50) = 0$. What can happen? If the Lakers win (and hence the Celtics lose), then he pays you \$60 in Bet 1 and you pay him \$50 in Bet 2; he has a net loss of \$10. If the Lakers

lose and the Celtics win, you pay him \$40 in Bet 1 and he pays you \$50 in Bet 2. Again he has a net loss of \$10. He is bound to pay you \$10 no matter what happens! He wants to know how you figured this out, so we will show him. First, the bets are determined in the following way. Let p represent his stated probability that the Lakers win and q his probability that the Celtics win. To be coherent, $p + q$ should equal 1. To make up the two bets with zero expected value to him, consider his sides of the bets:

Bet 1 He wins the amount pX with probability $1 - p$
 (that is, if the Lakers lose/Celtics win). You win (he loses)
 the amount $(1 - p)X$ with probability p .

Bet 2 You win the amount $(1 - q)Y$ with probability q
 (that is, if the Celtics win/Lakers lose). He
 wins the amount qY with probability $1 - q$.

As before, we can think of Y and X as the total stakes that are involved in each bet. For example, in Bet 1 you put pX into the pot and he puts in $(1 - p)X$, and whoever wins the bet gets all of the money in the pot. It is easy to verify that his expected value for each bet is zero; for example, in Bet 1 his expected value is

$$(1 - p)(pX) - p[(1 - p)X] = 0$$

Now, think about the equations that give his net gain or loss. If the Lakers win, his net position is L :

The two terms on the left-hand side of the equation are his loss from Bet 1 and his gain from Bet 2. Likewise, if the Celtics win, his position is K :

$$pX - (1 - q)Y = K$$

Now, here's the trick. He has supplied p and q . *You have decided to set L and K each equal to $-\$10$ so that he is guaranteed to pay you \$10.* Now it is just a matter of solving for X and Y to find the bets necessary to guarantee that he loses \$ 10. There are two equations and two unknowns (X and Y). What could be easier?

"Aha!" he exclaims, "In my linear algebra course, I learned about solving linear equations simultaneously. No matter what L and K are, you should be able to find an X and Y to defeat me. Unless . . . Oh, I see!" What does he see? Solve the second equation for X to get

$$X = \frac{K + (1 - q)Y}{p}$$

Now substitute this into the first equation and solve for Y . Rearranging yields

$$Y = \frac{L - K = \frac{K}{p}}{1 - \frac{1 - q}{p}}$$

This equation always can be solved for Y unless the denominator $[1 - (1 - q)/p]$ on the right-hand side happens to equal 0. That will occur only when $p = 1 - q$, or $p + q = 1$. This was exactly the condition he violated in the first place; his probabilities did not sum to 1. If his probabilities had satisfied this condition, you would not be able to find an X and Y that would solve the two equations, and the Dutch book would have been impossible.

This example points out how incoherence can be exploited. In fact, if a person's probabilities do not conform to the probability laws, it is always possible to do this kind of thing, no matter which probability law is violated. The contribution of de Finetti was to prove that it is only possible not to be exploited if subjective probabilities obey the probability postulates and laws. The practical importance of this is less that you can set up Dutch books against incoherent probability assessors — because no one in his or her right mind would agree to such a series of bets — but to provide insight into why subjective probability should work the same way as "long-run frequency" probability. Because no one in his or her right mind would make decisions in a way that could be exploited, coherence is a reasonable condition to guide our assessments of probabilities for decision-making purposes. When we assess probabilities for use in a decision tree or influence diagram, those probabilities should obey all the normal probability properties.

The idea of coherence has an important implication for assessment when a decision analysis involves assessing several probabilities. Once all probabilities have been assessed, it is important to check for coherence. If the assessments are not coherent, the decision maker should be made aware of this and given the chance to adjust the assessments until they are coherent.

SUMMARY

Many of the decision problems that we face involve uncertain future events. We may have some feeling for such uncertainties, and we can build models of them using subjective probability-assessment techniques. The basic approach to assessing a probability involves either setting up a bet or comparing lotteries. We also have considered assessment methods for continuous uncertain quantities and found that it is straightforward to assess continuous distributions in terms of cumulative distribution functions. A reasonable and practical way to incorporate a continuous distribution into a decision tree is to use a discrete approximation. We discussed bracket medians and the Pearson-Tukey three-point approximation.

Our discussion also touched on the pitfalls of probability assessment. Individuals tend to use cognitive heuristics to judge probabilities. Heuristics such as representativeness, availability, and anchoring and adjustment can lead to bias in probability assessment. Certain ideas for improving probability assessments were discussed, including decomposition of the assessment task. We also presented a protocol-based approach to obtaining expert probability assessments and showed how this approach was used in a risk assessment of a nuclear-waste repository.

Finally, we discussed the idea of coherence, that is, that subjective probabilities must obey the same probability laws which "long-run frequency" probabilities do.

Being coherent in probability assessment means being sure that one cannot be exploited through a Dutch book.

EXERCISES

- 8.1** Explain in your own words the idea of subjective probability.
- 8.2** An accounting friend of yours has gone to great lengths to construct a statistical model of bankruptcy. Using the model, the probability that a firm will file for bankruptcy within a year is calculated on the basis of financial ratios. On hearing your explanation of subjective probability, your friend says that subjective probability may be all right for decision analysis, but his model gives objective probabilities on the basis of real data. Explain to him how his model is, to a great extent, based on subjective judgments. Comment also on the subjective judgments that a bank officer would have to make in using your friend's model.
- 8.3** Explain in your own words the difference between assessing the probability for a discrete event and assessing a probability distribution for a continuous unknown quantity.
- 8.4** For each of the following phrases write down the probability number that you feel is represented by the phrase. After doing so, check your answers to be sure they are consistent. [For example, is your answer to e less than your answer to j?]
- a "There is a better than even chance that..."
 - b "A possibility exists that..."
 - c "... has a high likelihood of occurring."
 - d "The probability is very high that..."
 - e "It is very unlikely that..."
 - f "There is a slight chance..."
 - g "The chances are better than even..."
 - h "There is no probability, no serious probability, that..."
 - i "... is probable."
 - j "... is unlikely."
 - k "There is a good chance that..."
 - l "... is quite unlikely."
 - m "... is improbable."
 - n "... has a high probability."
 - o "There is a chance that..."
 - p "... is very improbable."
 - q "... is likely."
 - r "Probably..."
- 8.5** Suppose your father asked you to assess the probability that you would pass your decision-analysis course. How might you decompose this probability assessment? Draw an influence diagram to represent the decomposition.

QUESTIONS AND PROBLEMS

- 8.6** Assess your probability that the following outcomes will occur. Use the equivalent-lottery method as discussed in the chapter. If possible, use a wheel of fortune with an adjustable win area, or a computer program that simulates such a wheel. What issues did you account for in making each assessment?
- It will rain tomorrow in New York City.
 - You will have been offered a job before you graduate.
 - The women's track team at your college will win the NCAA championship this year.
 - The price of crude oil will be more than \$30 per barrel on January 1, 2010.
 - The Dow Jones industrial average will go up tomorrow.
 - Any other uncertain outcome that interests you.
- 8.7** Consider the following two outcomes:
- You will get an A in your most difficult course.
 - You will get an A or a B in your easiest course.
- Can you assess the probability of these outcomes occurring? What is different about assessing probabilities regarding your own performance as compared to assessing probabilities for outcomes like those in Problem 8.6?
- 8.8** Describe a decomposition strategy that would be useful for assessing the probabilities in Problem 8.7.
- 8.9** Many people deal with uncertainty by assessing odds. For example, in horse racing different horses' odds of winning are assessed. Odds of " a to b for Outcome E " means that $P(E) = a/(a + b)$. Odds of " c to d against Outcome E " means that $P(E) = c/(c + d)$. For the outcomes in Problem 8.6, assess the odds for that outcome occurring. Convert your assessed odds to probabilities. Do they agree with the probability assessments that you made in Problem 8.6?
- 8.10** It is said that Napoleon assessed probabilities at the Battle of Waterloo in 1815. His hopes for victory depended on keeping the English and Prussian armies separated. Believing that they had not joined forces on the morning of the fateful battle, he indicated his belief that he had a 90% chance of defeating the English; $P(\text{Napoleon Wins}) = 0.90$. When told later that elements of the Prussian force had joined the English, Napoleon revised his opinion downward on the basis of this information, but his posterior probability was still 60%; $P(\text{Napoleon Wins} \mid \text{Prussian and English Join Forces}) = 0.60$.
- Suppose Napoleon were using Bayes' theorem to revise his information. To do so, he would have had to make some judgments about $P(\text{Prussian and English Join Forces} \mid \text{Napoleon Wins})$ and $P(\text{Prussian and English Join Forces} \mid \text{Napoleon Loses})$. In particular, he would have had to judge the ratio of these two probabilities. Based on the prior and posterior probabilities given above, what is that ratio?
- 8.11** Should you drop your decision-analysis course? Suppose you faced the following problem: If you drop the course, the anticipated salary in your best job offer will depend on your current GPA:

$$\text{Anticipated Salary} \mid \text{Drop} = (\$4000 \times \text{Current GPA}) + \$16,000$$


If you take the course, the anticipated salary in your best job offer will depend on both your current GPA and your overall score (on a scale of 0 to 100) in the course:

$$\begin{aligned} \text{Anticipated Salary} \mid \text{Do Not Drop} &= 0.6(\$4000 \times \text{Current GPA}) \\ &+ 0.4(\$170 \times \text{Course Score}) + \\ &\$16,000 \end{aligned}$$

The problem is that you do not know how well you will do in the course. You can, however, assess a distribution for your score. Assuming that 90–100 is an A, 80–89 is a B, 70–79 a C, 60–69 a D, and 0–59 an F, assess a continuous probability distribution for your numerical score in the course. Use that distribution to decide whether or not to drop the course. Figure 8.20 shows your decision tree.

- 8.12** Assess these fractiles for the following uncertain quantities: 0.05 fractile, 0.25 fractile (first quartile), 0.50 (median), 0.75 fractile (third quartile), and 0.95 fractile. Plot your assessments to create graphs of your subjective CDFs.
- The closing Dow Jones industrial average (DJIA) on the last Friday of the current month.
 - The closing DJIA on the last Friday of next year.
 - The exchange rate, in Japanese yen per dollar, at the end of next Monday.
 - The official high temperature at O'Hare International Airport tomorrow.
 - The number of fatalities from airline accidents in the United States next year.
 - The number of casualties from nuclear power plant accidents in the United States over the next 10 years.
 - The value of the next jackpot won in the California state lottery.
- 8.13** For each of the following 10 items, assess the 0.05 and 0.95 fractiles. That is, choose upper and lower estimates such that you are 90% sure that the actual value falls between your estimates. Your challenge is to be neither too narrow (i.e., overconfident) nor too wide (underconfident).

Figure 8.20
Decision tree for Question 8.11. Should you drop your decision-analysis course?



	0.05 Fractile (Low)	0.95 Fractile (High)
1. Martin Luther King's age at death		
2. Length of the Nile River		
3. Number of countries that are members of OPEC		
4. Number of books in the Old Testament		
5. Diameter of the moon		
6. Weight of an empty Boeing 747		
7. Year of Wolfgang Amadeus Mozart's birth		
8. Gestation period of an Asian elephant		
9. Air distance from London to Tokyo		
10. Depth of deepest known point in the oceans		

If you have done a good job, most of your intervals contain the actual value (given in the back of the book — no peeking!). If three or more of your intervals missed the actual value, though, you have demonstrated overconfidence. Given your results in this question, would you adjust your assessments in Problem 8.12? [Source: Adapted from J. E. Russo, and P. J. H. Schoemaker (1989) *Decision Traps: The Ten Barriers to Brilliant Decision Making and How to Overcome Them*. New York: Fireside.]

8.14 Forecasters often provide only point forecasts, which are their best guesses as to an upcoming event. For example, an economic forecaster might predict that U.S. gross national product (GNP) will increase at a 3% annual rate over the next three months. Occasionally a forecaster also will provide an estimate of the degree of confidence in the point forecast to indicate how sure (or unsure) the forecaster is.

In what sense can your answers to Problem 8.12 be interpreted as forecasts? What advantages do subjective probability distributions have over typical point forecasts? What disadvantages? How could a decision maker use probabilistic forecasts such as those in Problem 8.12?

8.15 Choose a course that you are currently taking in which the final exam is worth 100 points. Treating your score on the exam as if it were a continuous uncertain quantity, assess the subjective probability distribution for your score. After you have finished, check your assessed distribution for consistency by:

- a** choosing any two intervals you have judged to have equal probability content, and
- b** determining whether you would be willing to place small even-odds bets that your score would fall in one of the two intervals. (The bet would be called off if the score fell elsewhere.)
- c** After assessing the continuous distribution, construct a three-point approximation to this distribution with the extended Pearson-Tukey method. Use the approximation to estimate your expected exam score.

- d. Now construct a five-point approximation with bracket medians. Use this approximation to estimate your expected exam score. How does your answer compare with the estimate from part c?
- 8.16** Compare the discrete-approximation methods by doing the following:
- a Use the extended Pearson-Tukey method to create three-point discrete approximations for the continuous distributions assessed in Problem 8.12. Use the approximations to estimate the expected values of the uncertain quantities.
 - b Repeat part a, but construct five-point discrete approximations using bracket medians. Compare your estimated expected values from the two methods.
- 8.17** Assess the probability that the New York Mets will win the World Series (WS) next year. Call this probability p . Now assess the following probabilities: $P(\text{Win WS} \mid \text{Win Pennant})$ and $P(\text{Win Pennant})$. Use these to calculate $q = P(\text{Win WS}) = P(\text{Win WS} \mid \text{Win Pennant}) P(\text{Win Pennant}) = P(\text{Win WS and Win Pennant})$. (To play in the World Series, the team must first win the pennant.)
- a Does $p = q$? Which of the two assessments are you more confident about? Would you adjust your assessments to make $p = q$? Why or why not?
 - b If you do not like the Mets, do this problem for your favorite major league baseball team.
- 8.18** Assess the probability that you will be hospitalized for more than one day during the upcoming year. In assessing this probability, decompose the assessment based on whether or not you are in an automobile accident. With this decomposition, you must assess $P(\text{Hospitalized} \mid \text{Accident})$, $P(\text{Hospitalized} \mid \text{No Accident})$, and $P(\text{Accident})$. In what other ways might this assessment be decomposed?
- 8.19** Choose a firm in which you are particularly interested, perhaps one where you might like to work. Go to the library and read about this firm, about its industry, and about the relationship of this firm and the industry to the economy as a whole. After your reading, assess a subjective CDF for the firm's revenue over the next fiscal year. Discuss the assessment process. In particular, what did you learn in your research that had an impact on the assessment? What kinds of decomposition helped you in your assessment process?
- 8.20** After observing a long run of red on a roulette wheel, many gamblers believe that black is bound to occur. Such a belief often is called the *gambler's fallacy* because the roulette wheel has no memory. Which probability-assessment heuristic is at work in the gambler's fallacy? Explain.
- 8.21** Look again at Problems 7.25 and 7.26. These problems involve assessments of the relative likelihood of different statements. When an individual ranks "Linda is a bank teller and is active in the feminist movement" as more probable than the statement "Linda is a bank teller," which of the three probability-assessment heuristics may be at work? Explain.
- 8.22** Suppose that you are a solid B student. Your grades in your courses have always been B's. In your statistics course, you get a D on the midterm. When your parents express concern, what statistical phenomenon might you invoke to persuade them not to worry about the D?
- 8.23** When we assess our subjective probabilities, we are building a model of the uncertainty we face. If we face a range-of-risk problem, and we begin to assess a continuous distribution subjectively, then clearly it is possible to perform many assessments, which would make the sketched CDF smoother and smoother. How do we know when to stop making

assessments and when to keep going? When is the model of our uncertainty adequate for solving the problem, and when is it inadequate?

- 8.24** Most of us have a hard time assessing probabilities with a lot of precision. For instance, in assessing the probability of rain tomorrow, even carefully considering the lotteries and trying to adjust a wheel of fortune to find the indifference point, many people would eventually say something like this: "If you set $p = 0.2$, I'd take Lottery A, and if $p = 0.3$, I'd take Lottery B. My indifference point must be somewhere in between these two numbers, but I am not sure where."

How could you deal with this kind of imprecision in a decision analysis? Illustrate how your approach would work using the umbrella problem (Figure 4.9). (*Hint*: The question is *not* how to get more precise assessments. Rather, given that the decision maker refuses to make precise assessments, and you are stuck with imprecise assessments, what kinds of decision-analysis techniques could you apply to help the individual make a decision?)

- 8.25** It is not necessary to have someone else set up a series of bets against you in order for incoherence to take its toll. It is conceivable that one could inadvertently get oneself into a no-win situation through inattention to certain details and the resulting incoherence, as this problem shows.

Suppose that an executive of a venture-capital investment firm is trying to decide how to allocate his funds among three different projects, each of which requires a \$100,000 investment. The projects are such that one of the three will definitely succeed, but it is not possible for more than one to succeed. Looking at each project as an investment, the anticipated payoff is good, but not wonderful. If a project succeeds, the payoff will be a net gain of \$150,000. Of course, if the project fails, he loses all of the money invested in that project. Because he feels as though he knows nothing about whether a project will succeed or fail, he assigns a probability of 0.5 that each project will succeed, and he decides to invest in each project.

- a According to his assessed probabilities, what is the expected profit for each project?
 - b What are the possible outcomes of the three investments, and how much will he make in each case?
 - c Do you think he invested wisely? Can you explain why he is in such a predicament?
 - d If you feel you know nothing about some event, is it reasonable to assess equal probabilities for the outcomes? Give an example where this might be reasonable and another where it might not.
- 8.26** Could you ever set up a Dutch book against a bookie who places bets for a living? Work through the following problem and think about this question.

Consider a baseball season when the Chicago Cubs and the New York Yankees meet in the World Series. A friend of yours is a Cubs fan, and you are trying to find out how confident he is about the Cubs' chances of winning. He tells you that he would bet on the Cubs at odds of 3:2 *or better* and on the Yankees at odds of 1:2 *or better*. [Odds of $a:b$ for an event means the probability of the event is $a/(a + b)$.] This means that he would be happy with either of the following bets:

- | | |
|--------------|--|
| Bet 1 | He wins \$20 if the Cubs win.
He loses \$30 if the Yankees win. |
| Bet 2 | He loses \$20 if the Cubs win.
He wins \$40 if the Yankees win. |

- a Explain why he would be happy to accept any modification of these bets as long as the amount he would win increases and the amount he would lose stays the same or is lower. Explain in terms of his expected bet value.
- b Show mathematically that $0.60 \leq P(\text{Cubs Win}) \leq 0.67$. (*Hint*: If he is willing to accept either bet, what does that imply about his expected value for each bet?)
- c Would it be possible to set up a Dutch book against this individual? If your answer is yes, what bets would you place, and how much would you be sure to win? If your answer is no, explain why not. (Be careful! Make sure he is willing to accept the bets you propose.)
- d Most individuals have a hard time assessing subjective probabilities with a very high degree of precision. (See Problem 8.23.) Does this problem shed any light on the issue?

8.27 Ellsberg Paradox A barrel contains a mixture of 90 red, blue, and yellow balls. Thirty of the balls are red, and the remaining 60 are a mixture of blue or yellow, but the proportion of blue and yellow is unknown. A single ball will be taken randomly from the barrel.

- a Suppose you are offered the choice between gambles A and B:

A: Win \$ 1000 if a red ball is chosen.

B: Win \$1000 if a blue ball is chosen.

Would you prefer A or B? Why? **b** Now suppose that you are offered a choice between gambles C and D:

C: Win \$1000 if either a red or a yellow ball is chosen. **D:**

Win \$1000 if either a blue or a yellow ball is chosen.

Would you prefer C or D? Why?

- c Many people prefer A in the first choice and D in the second. Do you think this is inconsistent? Explain.

C A S E S T U D I E S

ASSESSING CANCER RISK — FROM MOUSE TO MAN

Cancer is a poorly understood and frightening disease. Its causes are essentially unknown, and the biological process that creates cancerous cells from healthy tissue remains a mystery. Given this state of ignorance, much research has been conducted into how external conditions relate to cancer. For example, we know that smoking leads to substantially higher incidence of lung cancer in humans. Cancer appears spontaneously, however, and its onset seems to be inherently probabilistic, as shown by the fact that some people smoke all their lives without developing lung cancer.

Some commentators claim that the use of new and untested chemicals is leading to a cancer epidemic. As evidence they point to the increase in cancer deaths over the years. Indeed, cancer deaths have increased, but people generally have longer life spans, and more elderly people now are at risk for the disease. When cancer rates are

adjusted for the increased life span, cancer rates have not increased substantially. In fact, when data are examined this way, some cancer rates (liver, stomach, and uterine cancer) are less common now than they were 50 years ago (1986 data of the American Cancer Society). Nevertheless, the public fears cancer greatly. The Delaney Amendment to the Food, Drug, and Cosmetics Act of 1954 outlaws residues in processed foods of chemicals that poses any risk of cancer to animals or humans. One of the results of this fear has been an emphasis in public policy on assessing cancer risks from a variety of chemicals.

Scientifically speaking, the best way to determine cancer risk to humans would be to expose one group of people to the chemical while keeping others away from it. But such experiments would not be ethical. Thus, scientists generally rely on experiments performed on animals, usually mice or rats. The laboratory animals in the experimental group are exposed to high doses of the substance being tested. High doses are required because low doses probably would not have a statistically noticeable effect on the relatively small experimental group. After the animals die, cancers are identified by autopsy. This kind of experiment is called a *bioassay*. Typical cancer bioassays involve 600 animals, require two to three years to complete, and cost several hundred thousand dollars.

When bioassays are used to make cancer risk assessments, two important extrapolations are made. First, there is the extrapolation from high doses to low doses. Second, it is necessary to extrapolate from effects on test species to effects on humans. On the basis of data from laboratory experiments and these extrapolations, assessments are made regarding the incidence of cancer when humans are exposed to the substance.

Questions

- 1 Clearly, the extrapolations that are made are based on subjective judgments. Because cancer is viewed as being an inherently probabilistic phenomenon, it is reasonable to view these judgments as probability assessments. What kinds of assessments do you think are necessary to make these extrapolations? What issues must be taken into account? What kind of scientific evidence would help in making the necessary assessments?
- 2 It can be argued that most cancer risk assessments are weak evidence of potential danger or lack thereof. To be specific, a chemical manufacturer and a regulator might argue different sides of the same study. The manufacturer might claim that the study does not conclusively show that the substance is dangerous, while the regulator might claim that the study does not conclusively demonstrate safety. Situations like this often arise, and decisions must be made with imperfect information. What kind of strategy would you adopt for making these decisions? What trade-offs does your strategy involve?
- 3 In the case of risk assessment, as with many fields of scientific inquiry, some experiments are better than others for many reasons. For example, some experiments may be more carefully designed or use larger samples. In short, some sources of information are more "credible."

For a simple, hypothetical example, suppose that you ask three "experts" whether a given coin is fair. All three report that the coin is fair; for each one the best estimate of $P(\text{Heads})$ is 0.50. You learn, however, that the first expert flipped the coin 10,000 times and observed heads on 5000 occasions. The second flipped the coin 20 times and observed 10 heads. The third expert did not flip the coin at all, but gave it a thorough physical examination, finding it to be perfectly balanced, as nearly as he could measure.

How should differences in credibility of information be accounted for in assessing probabilities? Would you give the same weight to information from the three examiners of the coin? In the case of putting together information on cancer risk from multiple experiments and expert sources, how might you deal with the information sources' differential credibility?

Source: D. A. Freedman and H. Zeisel (1988) "From Mouse-to-Man: The Quantitative Assessment of Cancer Risks." *Statistical Science*, 3, 3-56. Includes discussion.

BREAST IMPLANTS

The controversy surrounding breast implants is a good example of scientific uncertainty. Yanked from the market by the Food and Drug Administration in 1991 because of some evidence of dangerous side effects, breast implants received a reprieve in 1994 when the *New England Journal of Medicine* published an article that found no evidence of danger. An editorial in the *San Jose Mercury News* opined regarding the controversy:

The wisdom of letting science settle this question [regarding breast implants] will win the ready assent of everyone. "Scientific certainty" has such a reassuring ring to it.

But the implant case is a sterling example of the limits of science. Regarding long-term harm, science often can't provide definitive answers within the time a decision needs to be made.

And many policy decisions are out of the realm of science. What level of risk is reasonable for a woman who's lost a breast to cancer? What about reasonable risk for a woman who wants a fuller figure? Who decides what's reasonable? Is cosmetic breast enhancement a misuse of medicine to reinforce outdated notions of female beauty? [*Source:* "Uncertain Science," *San Jose Mercury News*, June 17, 1994, p. 10B.]

Questions

- 1 What kinds of information should a jury consider when deciding whether a plaintiff's claims of damages are reasonable? Anecdotes? The number of plaintiffs filing similar lawsuits? Scientific studies? What does a juror need to know in order to evaluate the quality of a scientific study? Do you think the average juror (or judge for that matter!) in the United States has the ability to critically evaluate the quality of scientific studies?

Discuss the questions asked in the last paragraph of the quote above. Do these questions relate to uncertainty (scientific or otherwise) or to values? Given the imperfect information available about breast implants, what role should the manufacturers play in deciding what risks are appropriate for which women? What role should government agencies play? What role should individual consumers play?

THE SPACE SHUTTLE *CHALLENGER*

On January 28, 1986, the space shuttle *Challenger* lifted off from an ice-covered launch pad. Only 72 seconds into the flight, the shuttle exploded, killing all seven astronauts aboard. The United States and the rest of the world saw the accident firsthand as films from NASA were shown repeatedly by the television networks.

Before long the cause of the accident became known. The shuttle's main engines are fueled by liquid hydrogen and oxygen stored in a large tank carried on the shuttle's belly. Two auxiliary rockets that use solid fuel are mounted alongside the main fuel tank and provide additional thrust to accelerate the shuttle away from the launch pad. These boosters use their fuel rapidly and are jettisoned soon after launch.

The solid rocket boosters are manufactured in sections by Morton Thiokol, Inc. (MTI), in Utah. The sections are shipped individually to Kennedy Space Center (KSC) in Florida where they are assembled. The joints between sections of the rocket are sealed by a pair of large rubber O-rings, whose purpose is to contain the hot gases and pressure inside the rocket. In the case of the *Challenger*, one of the joint seals failed. Hot gases blew past the O-rings and eventually burned through the large belly tank, igniting the highly explosive fuel inside. The resulting explosion destroyed the spacecraft.

Before long it also became known that the launch itself was not without controversy. MTI engineers had been aware of the problems with the O-rings for some time, having observed eroded O-rings in the boosters used on previous flights. A special task force was formed in 1985 to try to solve the problem, but ran into organizational problems. One memo regarding the task force began, "Help! The seal task force is constantly being delayed by every possible means." The problem came to a head when, on the evening before the launch, MTI engineers recommended not launching the shuttle because of the anticipated cold temperatures on the launch pad. After a teleconference involving officials at KSC and the Marshall Space Flight Center (MSFC) in Alabama, management officials at MTI reversed their engineers' recommendation and approved the launch.

Questions

- 1 To a great extent, the engineers were concerned about the performance of the O-ring under anticipated cold weather conditions. The coldest previous flight had been 53°F, and, knowing of the existing problems with the seals, the engineers hesitated to recommend a launch under colder conditions. Technically, the problem

was that an O-ring stiffens as it gets colder, thus requiring a longer time to seal a joint. The real problem, however, was that the engineers did not know much about the performance of the O-rings at cold temperatures. Robert K. Lund, vice president of engineering for MTI, testified to the presidential commission investigating the accident, "[W]e just don't know how much further we can go below the 51 or 53 degrees or whatever it was. So we were concerned with the unknown... . They [officials at MSFC] said they didn't accept that rationale" (*Report of the Presidential Commission on the Space Shuttle Challenger Accident*, p. 94).

The MTI staff felt as if it were in the position of having to prove that the shuttle was unsafe to fly instead of the other way around. Roger Boisjoly, an MTI engineer, testified, "This was a meeting where the determination was to launch, and it was up to us to prove beyond a shadow of a doubt that it was not safe to do so. This is in total reverse to what the position usually is in a preflight conversation or a flight readiness review. It is usually exactly opposite that" (*Report*, p. 93).

NASA solicited information regarding ice on the launch pad from Rockwell International, the shuttle's manufacturer. Rockwell officials told NASA that the ice was an unknown condition. Robert Glaysher, a vice president at Rockwell, testified that he had specifically said to NASA, "Rockwell could not 100 percent assure that it is safe to fly" (*Report*, p. 115). In this case, the presidential commission also found that "NASA appeared to be requiring a contractor to prove that it was not safe to launch, rather than proving it was safe" (*Report*, p. 118).

The issue is how to deal with unknown information. What do you think the policy should be regarding situations in which little or no information is available? Discuss the problems faced by both MTI and NASA. What incentives and pressures might they have faced?

- 2 Professor Richard Feynman, Nobel Laureate in physics, was a member of the commission. He issued his own statement, published as an appendix to the report, taking NASA to task for a variety of blunders. Some of his complaints revolved around assessments of the probability of failure.

Failure of the solid rocket boosters. A study of 2900 flights of solid-fuel rockets revealed 121 failures, or approximately 1 in 25. Because of improved technology and special care in the selection of parts and in inspection, Feynman is willing to credit a failure rate of better than 1 in 100 but not as good as 1 in 1000. But in a risk analysis prepared for the Department of Energy (DOE) that related to DOE radioactive material aboard the shuttle, NASA officials used a figure of 1 in 100,000. Feynman writes:

If the real probability is not so small [as 1 in 100,000], flights would show troubles, near failures, and possibly actual failures with a reasonable number of trials, and standard statistical methods could give a reasonable estimate. In fact, previous NASA experience had shown, on occasion, just such difficulties, near accidents, and accidents, all giving warning that the probability of flight failure was not so very small. (*Report*, p. F-1)

Failure of the liquid fuel engine. In another section of his report, Feynman discussed disparate assessments of the probability of failure of the liquid fuel engine. His own calculations suggested a failure rate of approximately 1 in 500. Engineers at Rocketdyne, the engine manufacturer, estimated the probability to be approximately

1 in 10,000. NASA officials estimated 1 in 100,000. An independent consultant for NASA suggested that a failure rate of 1 or 2 per 100 would be a reasonable estimate.

How is it that these probability estimates could vary so widely? How should a decision maker deal with probability estimates that are so different?

- 3 To arrive at their overall reliability estimates, NASA officials may have decomposed the assessment, estimated the reliability of many different individual components, and then aggregated their assessments. Suppose that, because of an optimistic viewpoint, each probability assessment had been slightly overoptimistic (that is, a low assessed probability of failure). What effect might this have on the overall reliability estimate?
- 4 In an editorial in *Space World* magazine, editor Tony Reichhardt commented on the accident:

One person's safety is another's paranoia. How safe is safe? What is acceptable risk? It's no small question, in life or in the space program. It's entirely understandable that astronauts would come down hard on NASA policies that appear to be reckless with their lives. But unless I'm misreading the testimony [before the commission], at the end of the teleconference that night of January 27, most of the participating engineers believed that it was safe to go ahead and launch. A few argued that it was not safe *enough*. There was an element of risk in the decision, and in many others made prior to *Challenger's* launch, and seven people were killed.

Whether this risk can be eliminated is a question of monumental importance to the space program. Those who have put the blame squarely on NASA launch managers need to think hard about the answer. If no Shuttle takes off until everyone at every level of responsibility is in complete agreement, then it may never be launched again. No single person can be absolutely sure that the whole system will work. On this vehicle, or on some other spacecraft next year or 30 years from now — even if we ease the financial and scheduling pressures — something will go wrong again. [Source: T. Reichhardt (1986) "Acceptable Risk," *Space World*, April, p. 3.]

Comment on Reichhardt's statement. What is an acceptable risk? Does it matter whether we are talking about risks to the general public from cancer or risks to astronauts in the space program? Would your answer change if you were an astronaut? A NASA official? A manufacturer of potentially carcinogenic chemicals? A cancer researcher? How should a policy maker take into account the variety of opinions regarding what constitutes an acceptable risk?

Source: Information for this case was taken from many sources, but by far the most important was the report by the Rogers Commission (1986) *Report of the Presidential Commission on the Space Shuttle Challenger Accident*. Washington, DC: U.S. Government Printing Office.

REFERENCES

The subjective interpretation of probability is one of the distinguishing characteristics of decision theory and decision analysis. This interpretation was presented first by Savage (1954) and has been debated by probabilists and statisticians ever since. Winkler (1972) provides an excellent introduction to subjective probability and Bayesian statistics (so

called because of the reliance on Bayes' theorem for inference), as well as extensive references to the literature.

Although we argued that verbal expressions of uncertainty are by nature less precise than numerical probabilities, Wallsten, Budescu, and Zwick (1993) show how to *calibrate* a set of verbal probability labels so that the meaning of each label is precisely understood.

Spetzler and Stael von Holstein (1975) is the standard reference on probability assessment. Winkler (1972) also covers this topic. Wallsten and Budescu (1983) review the field from a psychological perspective.

The construction of an appropriate discrete distribution is an interesting problem that has occupied a number of decision-analysis researchers. Many different approaches exist; the extended Pearson-Tukey and bracket median approaches are two of the more straightforward ones that perform well. Other good ones include the *extended Swanson-Megill* method, in which the distribution's median gets probability 0.4 and the 0.10 and 0.90 fractiles get probability 0.3 each. More complicated (and more precise) methods create discrete distributions with mathematical characteristics that match those of the continuous distribution (e.g., Miller and Rice 1983, Smith 1993); such approaches are difficult to use "by hand" but are easily implemented in computer software. Studies that report on the relative performance of different discrete-approximation methods include Keefer and Bodily (1983), Keefer (1994), and Smith (1993).

The literature on heuristics and biases is extensive. One reference that covers the topic at an introductory level is Tversky and Kahneman (1974). Hogarth (1987) provides a unique overview of this material in the context of decision analysis. Kahneman, Slovic, and Tversky (1982) have collected key research papers in the area. With respect to the issue of overconfidence in particular, Capen (1976) reports an experiment that demonstrates the extent of overconfidence in subjective probability assessments and possible ways to cope with this phenomenon.

Probability decomposition is a topic that has not been heavily researched, although there are many applications. For example, see Bunn (1984) and von Winterfeldt and Edwards (1986) for discussions and examples of fault trees. A recent research paper by Ravinder, Kleinmuntz, and Dyer (1988) discusses when decomposition is worth doing. Fault-tree analysis is widely used in reliability engineering and risk analysis. It is described in Gottfried (1974), Bunn (1984), and Merkhofer (1987b). Covello and Merkhofer (1993) provide additional references regarding fault trees.

As mentioned, the concept of coherence was first introduced by de Finetti (1937). Thinking in terms of coherence as did de Finetti contrasts with the axiomatic approach taken by Savage (1954). The idea of coherence has been developed considerably; examples of ways in which various decision rule or probability assessments are incoherent are given by Bunn (1984), French (1986), and Lindley (1985).

Finally, a topic that arose in Problems 8.23-8.26 deserves mention. This is the matter of vagueness or ambiguity in probability assessments. Problem 8.26 is the classic paradox of Ellsberg (1961), essentially showing that people shy away from risky prospects with vague probabilities. Ongoing research is attempting to understand this phenomenon both descriptively and prescriptively. Einhorn and Hogarth (1985) present a psychological model of the process. Frisch and Baron (1988) define vagueness in terms of the lack of information that could have an impact on a probability assessment. The setup in Problem 8.25 is a simplified version of Nau's (1990) model in which a decision maker places constraints on the size of the bets that would be acceptable at different odds levels.

- Abramson, B., and A. J. Finizza (1995) "Probabilistic Forecasts from Probabilistic Models Case Study: The Oil Market." *International Journal of Forecasting*, 11, 63-72.
- Beyth-Marom, R. (1982) "How Probable Is Probable? A Numerical Translation of Verbal Probability Expressions." *Journal of Forecasting*, 1, 257-269.
- Bunn, D. (1984) *Applied Decision Analysis*. New York: McGraw-Hill.
- Capen, E. C. (1976) "The Difficulty of Assessing Uncertainty." *Journal of Petroleum Technology*, August, 843-850. Reprinted in R. Howard and J. Matheson (eds.) (1983) *The Principles and Applications of Decision Analysis*. Menlo Park, CA: Strategic Decisions Group.
- Covello, V. T., and M. W. Merkhofer (1993) *Risk Assessment Methods: Approaches for Assessing Health and Environmental Risks*. New York: Plenum.
- de Finetti, B. (1937) "La Prevision: Ses Lois Logiques, Ses Sources Subjectives." *Annales de l'Institut Henri Poincaré*, 7, 1-68. Translated by H. E. Kyburg in H. E. Kyburg, Jr., and H. E. Smokier (eds.) (1964) *Studies in Subjective Probability*. New York: Wiley.
- Einhorn, H., and R. M. Hogarth (1985) "Ambiguity and Uncertainty in Probabilistic Inference." *Psychological Review*, 92, 433-461.
- Ellsberg, D. (1961) "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics*, 75, 643-669.
- French, S. (1986) *Decision Theory: An Introduction to the Mathematics of Rationality*. London: Wiley.
- Frisch, D., and J. Baron (1988) "Ambiguity and Rationality." *Journal of Behavioral Decision Making*, 1, 149-157.
- Gottfried, P. (1974) "Qualitative Risk Analysis: FTA and FMEA." *Professional Safety*, October, 48-52.
- Hogarth, R. M. (1987) *Judgment and Choice*, 2nd ed. New York: Wiley.
- Kahneman, D., P. Slovic, and A. Tversky (eds.) (1982) *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- Kahneman, D., and A. Tversky (1973) "On the Psychology of Prediction." *Psychological Review*, 80, 237-251.
- Keefer, D. L. (1994) "Certainty Equivalents for Three-Point Discrete-Distribution Approximations." *Management Science*, 40, 760-773.
- Keefer, D., and S. E. Bodily (1983) "Three-Point Approximations for Continuous Random Variables." *Management Science*, 29, 595-609.
- Lindley, D.V. (1985) *Making Decisions*, 2nd ed. New York: Wiley.
- Merkhofer, M. W. (1987a) "Quantifying Judgmental Uncertainty: Methodology, Experiences, and Insights." *IEEE Transactions on Systems, Man, and Cybernetics*, 17, 741-752.
- Merkhofer, M. W. (1987b) *Decision Science and Social Risk Management*. Dordrecht, Holland: Reidel.
- Miller, A. C., and T. R. Rice (1983) "Discrete Approximations of Probability Distributions." *Management Science*, 29, 352-362.

- Morgan, M. G., and M. Henrion (1990) *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge: Cambridge University Press.
- Nau, R. (1990) "Indeterminate Probabilities on Finite Sets." *Annals of Statistics*, 20, 1737-1787.
- Ravinder, H. V., D. Kleinmuntz, and J. S. Dyer (1988) "The Reliability of Subjective Probabilities Obtained Through Decomposition." *Management Science*, 34, 186-199.
- Savage, L. J. (1954) *The Foundations of Statistics*. New York: Wiley.
- Smith, J. E. (1993) "Moment Methods for Decision Analysis." *Management Science*, 39, 340-358.
- Spetzler, C. S., and C. A. Stael von Holstein (1975) "Probability Encoding in Decision Analysis." *Management Science*, 22, 340-352.
- Tversky, A., and D. Kahneman (1974) "Judgments under Uncertainty: Heuristics and Biases." *Science*, 185, 1124-1131.
- von Winterfeldt, D., and W. Edwards (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press.
- Wallsten, T. S., and D. V. Budescu (1983) "Encoding Subjective Probabilities: A Psychological and Psychometric Review." *Management Science*, 29, 151-173.
- Wallsten, T. S., D. V. Budescu, and R. Zwick (1993) "Comparing the Calibration and Coherence of Numerical and Verbal Probability Judgments." *Management Science*, 39, 176-190.
- Winkler, R. L. (1972) *Introduction to Bayesian Inference and Decision*. New York: Holt.

EPILOGUE

Stanford beat Oregon in their 1970 football contest. The score was 33-10. Jim Plunkett was Stanford's quarterback that season and won the Heisman trophy.

Millard Fillmore was Zachary Taylor's vice president. Taylor died in 1850 while in office. Fillmore succeeded him and was president from 1850 through 1853.

Theoretical Probability Models

The last chapter dealt with the subjective assessment of probabilities as a method of modeling uncertainty in a decision problem. Using subjective probabilities often is all that is necessary, although they occasionally are difficult to come up with, and the nature of the uncertainty in a decision situation can be somewhat complicated. In these cases, we need another approach.

An alternative source for probabilities is to use theoretical probability models and their associated distributions. We can consider the characteristics of the system from which the uncertain event of interest arises, and, if the characteristics correspond to the assumptions that give rise to a standard distribution, we may use the distribution to generate the probabilities. It is important to realize, however, that in this situation a substantial subjective judgment is being made: that the physical system can be represented adequately using the model chosen. In this sense, such probability models would be just as "subjective" as a directly assessed probability distribution.

One approach is to assess a subjective probability distribution and then find a standard distribution that provides a close "fit" to those subjective probabilities. A decision maker might do this simply to make the probability and expected-value calculations easier. Of course, in this case the underlying subjective judgment is that the theoretical distribution adequately fits the assessed judgments.

How important are theoretical probability distributions for decision making? Consider the following applications of theoretical models.

THEORETICAL MODELS APPLIED

Educational Testing Most major educational and intelligence tests generate distributions of scores that can be well represented by the normal distribution, or the familiar bell-shaped curve. Many colleges and universities admit only individuals whose scores are above a stated criterion that corresponds to a specific percentile of the normal distribution of scores.

Market Research In many market research studies, a fundamental issue is whether a potential customer prefers one product to another. The uncertainty involved in these problems often can be modeled using the binomial or closely related distributions.

Quality Control How many defects are acceptable in a finished product? In some products, the occurrence of defects, such as bubbles in glass or blemishes in cloth, can be modeled quite nicely with a Poisson process. Once the uncertainty is modeled, alternative quality-control strategies can be analyzed for their relative costs and benefits.

Predicting Election Outcomes How do the major television networks manage to extrapolate the results of exit polls to predict the outcome of elections? Again, the binomial distribution forms the basis for making probability statements about who wins an election based on a sample of results.

Capacity Planning Do you sometimes feel that you spend too much time standing in lines waiting for service? Service providers are on the other side; their problem is how to provide adequate service when the arrival of customers is uncertain. In many cases, the number of customers arriving within a period of time can be modeled using the Poisson distribution. Moreover, this distribution can be extended to the placement of orders at a manufacturing facility, the breakdown of equipment, and other similar processes.

Environmental Risk Analysis In modeling the level of pollutants in the environment, scientists often use the lognormal distribution, a variant of the normal distribution. With uncertainty about pollutant levels modeled in this way, it is possible to analyze pollution-control policies so as to understand the relative effects of different policies.

You might imagine that many different theoretical probability models exist and have been used in a wide variety of applications. We will be scratching only the surface here, introducing a few of the more common distributions. For discrete probability situations, we will discuss the binomial and Poisson distributions, both of which are commonly encountered and easily used. Continuous distributions that will be discussed are the exponential, normal, and beta distributions. The references at the end of this chapter will direct you to sources on other distributions.

The Binomial Distribution

Perhaps the easiest place to begin is with the *binomial* distribution. Suppose, for example, that you were in a race for mayor of your hometown, and you wanted to find out how you were doing with the voters. You might take a sample, count the number of individuals who indicated a preference for you, and then, based on this information, judge your chances of winning the election. In this situation, each voter interviewed can be either for you or not. This is the kind of situation in which the binomial distribution can play a large part. Of course, it is not limited to the analysis of voter preferences but also can be used in analyses of quality control where an item may or may not be defective, in market research, and in many other situations.

The binomial distribution arises from a situation that has the following characteristics:

- 1 *Dichotomous outcomes.* Uncertain events occur in a sequence, each one having one of two possible outcomes — success/failure, heads/tails, yes/no, true/false, on/off, and so on.
- 2 *Constant probability.* Each event, or *trial* has the same probability of success. Call that probability p .
- 3 *Independence.* The outcome of each trial is independent of the outcomes of the other trials. That is, the probability of success does not depend on the preceding outcomes. (As we noted in Problem 8.19 ignoring this sometimes is called the *gambler's fallacy*: Just because a roulette wheel has come up red five consecutive times does not mean that black is "due" to occur. Given independence from one game to the next, the probability of black still is 0.5, regardless of the previous outcomes.)

Now suppose we look at a sequence of n trials; for example, four tosses of a loaded coin that has $P(\text{Heads}) = 0.8$. How many successes (heads) could there be? Let R denote the uncertain quantity or random variable that is the number of successes in the sequence of trials. Clearly, there can be no more than n successes: $0 < R < n$. In our example, $0 < R < 4$, or there will be somewhere between zero and four heads.

What is the probability of four heads? Let H_i denote the event that a head occurs on the i th toss. Then four heads can happen only if the following sequence occurs: H_1, H_2, H_3, H_4 . What is $P(H_1, H_2, H_3, H_4)$?

$$\begin{aligned}
 P(R = 4 \mid n = 4, p = 0.8) &= P(H_1, H_2, H_3, H_4) \\
 &= P(H_1) P(H_2 \mid H_1) P(H_3 \mid H_2, H_1) P(H_4 \mid H_3, H_2, H_1) \\
 &\quad \text{(by using conditional probabilities)} \\
 &= P(H_1) P(H_2)P(H_3) P(H_4) \\
 &\quad \text{(by the independence property)} = p^4 \\
 &\quad \text{(by the constant probability property)} = 0.8^4 \\
 &\quad [P(\text{Heads}) = 0.8 \text{ for the loaded coin}] = 0.41
 \end{aligned}$$

What is the probability of three heads? In this case, three heads can occur in any of four ways:

- H₁, H₂, H₃, T
- H₁, H₂, T, H₄
- H₁, T, H₃, H₄
- T, H₂, H₃, H₄

The probability for any one of these sequences is $p^3(1 - p)$. Because these four sequences are mutually exclusive, the probability of three heads is simply the sum of the individual probabilities, or $P(\text{Three Heads in Four Tosses}) = P(R = 3 | n = 4, p = 0.8) = 4p^3(1 - p)^1 = 4(0.8^3)(0.2)$.

In general, the probability of obtaining r successes in n trials is given by

$$P_B(R = r | n, p) = \frac{n!}{r!(n - r)!} p^r (1 - p)^{n-r} \quad (9.1)$$

where the subscript B indicates that this is a binomial probability. The term with the factorials is called the *combinatorial term*. It gives the number of ways that a sequence of n trials can have r successes. The second term is just the probability associated with a particular sequence with r successes in n trials. The expected number of successes is simply $E(R) = np$. (If I have n trials, I expect proportion p of them to be successes.) The variance of the number of successes is $\text{Var}(R) = np(1 - p)$.

Binomial probabilities are not difficult to calculate using Formula (9.1), but you do not have to calculate them yourself. Individual probabilities [$P_B(r$ Successes)] and cumulative terms [$P_B(r$ or Fewer Successes)] can be found in the tables in Appendixes A and B at the end of this book. For example, $P_B(R = 2 | n = 9, p = 0.12)$ is found in the following way, as illustrated in Figure 9.1. First, we find the value $n = 9$ along the left margin, and then the value for $r = 2$. Now read across that row until you find the column headed by $p = 0.12$. The probability you read is 0.212. Thus, $P_B(R = 2 | n = 9, p = 0.12) = 0.212$.

Figure 9.1

Using the binomial table to find $P_B(R = 2 | n = 9, p = 0.12)$.

n	r	p	.10	.12	.14
.
.
.
9	0		.387	.316	.257
	1		.387	.388	.377
	2		.172	.212	.245
	3		.045	.067	.093
.
.

The tables can be used to find binomial probabilities in many different forms. For example, what is $P_B(R > 3 \mid n = 15, p = 0.2)$? To find this probability, all that is necessary is to consider the complement:

$$\begin{aligned} P_B(R > 3 \mid n = 15, p = 0.2) &= 1 - P_B(R \leq 3 \mid n = 15, p = 0.2) \\ &= 1 - 0.648 \quad (\text{from Appendix B}) \\ &= 0.352 \end{aligned}$$

You may have noticed that the values for p only go up to 0.50. What if you need to find a binomial probability with a p larger than 0.50? Let us try it. What is $P_B(R \leq 5 \mid n = 8, p = 0.7)$? To do this, we will look at the flip side of the situation. Getting five or fewer successes in eight trials is equivalent to getting three or more failures in eight trials, where $P(\text{Failure}) = 1 - P(\text{Success})$. Thus, we are arguing that

$$P_B(R \leq 5 \mid n = 8, p = 0.7) = P_B(R' \geq 3 \mid n = 8, p = 0.3)$$

where $R' = n - R$, or the number of failures. Now, to find $P_B(R' \geq 3 \mid n = 8, p = 0.3)$, use the idea of a complement:

$$\begin{aligned} P_B(R' > 3 \mid n = 8, p = 0.3) &= 1 - P_B(R' \leq 2 \mid n = 8, p = 0.3) \\ &= 1 - 0.552 = \\ &0.448 \end{aligned}$$

An Example: Melissa Bailey It is the beginning of winter term at the College of Business at Eastern State. Melissa Bailey says she is eager to study, but she has plans for a ski trip each weekend for 12 weeks. She will go only if there is good weather. The probability of good weather on any weekend during the winter is approximately 0.65. What is the probability that she will be gone for eight or more weekends?

Are all of the requirements satisfied for the binomial distribution? Weekend weather is good or bad, satisfying the dichotomous outcomes property. We will assume that the probabilities are the same from one weekend to the next, and moreover it seems reasonable to assume that the weather on one weekend during the winter is independent of the weather of previous weekends. Given these assumptions, the binomial distribution is an appropriate model for the uncertainty in this problem. Keep in mind that we are building a model of the uncertainty. Although our assumptions may not be exactly true, the binomial distribution should provide a good approximation.

To solve the problem, we must find $P_B(R \geq 8 \mid n = 12, p = 0.65)$. Of course, it is always possible to calculate this probability directly using the formula; however, let us see how to use the tables. The first problem is that the tables only give probabilities for p less than or equal to 0.50. To obtain the probability we want, we look at the problem in terms of bad weather rather than good. Being gone 8 or more weekends (good weather) out of 12 is the same as staying home on 4 or fewer weekends (bad weather). Find $P_B(R' \leq 4 \mid n = 12, p = 0.35)$. Using Appendix B, this is 0.583. Thus, there is more than a 50% chance that Melissa will be home on four or fewer weekends and gone on eight or more weekends.

Another Example: Soft Pretzels Having just completed your degree in business, you are eager to try your skills as an entrepreneur by marketing a new pretzel that you have developed. You estimate that you should be able to sell them at a competitive price of 50 cents each. The potential market is estimated to be 100,000 pretzels per year. Unfortunately, because of a competing product, you know you will not be able to sell that many. After careful research and thought, you conclude that the following model of the situation captures the relevant aspects of the problem: Your new pretzel might be a hit, in which case it will capture 30% of the market in the first year. On the other hand, it may be a flop, in which case the market share will be only 10%. You judge these outcomes to be equally likely.

Being naturally cautious, you decide that it is worthwhile to bake a few pretzels and test market them. You bake 20, and in a taste test against the competing product, 5 out of 20 people preferred your pretzel. Given these new data, what do you think the chances are that your new pretzel is a hit? The following analysis is one way that you might analyze the situation.

The question we are asking is this: What is $P(\text{New Pretzel a Hit} \mid 5 \text{ of } 20 \text{ Preferred New Pretzel})$? How can we get a handle on this probability? A problem like this that involves finding a probability given some new evidence almost certainly requires an application of Bayes' theorem. Let us use some notation to make our life simpler. Let "Hit" and "Flop" denote the outcomes that the new pretzel is a hit or a flop, respectively. Let R be the number of tasters (out of 20) who preferred the new pretzel. Now we can write down Bayes' theorem using this notation:

$$\begin{aligned} & P(\text{New Pretzel a Hit} \mid 5 \text{ of } 20 \text{ Preferred New Pretzel}) \\ &= P(\text{Hit} \mid R = 5) \\ &= \frac{P(R = 5 \mid \text{Hit}) P(\text{Hit})}{P(R = 5 \mid \text{Hit}) P(\text{Hit}) + P(R = 5 \mid \text{Flop}) P(\text{Flop})} \end{aligned}$$

Next we must fill in the appropriate probabilities on the right-hand side of the Bayes' theorem equation. The probabilities $P(\text{Hit})$ and $P(\text{Flop})$ are easy. Based on the judgment (stated above) that these two outcomes are considered to be equally likely, we can say that $P(\text{Hit}) = P(\text{Flop}) = 0.50$.

What about $P(R = 5 \mid \text{Hit})$ and $P(R = 5 \mid \text{Flop})$? These are a bit trickier. Consider $P(R = 5 \mid \text{Hit})$. This is the binomial probability that 5 out of 20 people prefer your pretzel, given that 30% ($p = 0.3$) of the entire population of pretzel customers would prefer yours. That is, if your pretzel is a hit, you will capture 30% of the market. How does this idea of 30% across the entire population relate to the chance of 5 out of a sample of 20 preferring your pretzel? Provided that we can view the 20 people in our sample as "randomly selected," we can apply the binomial distribution. ("Randomly selected" means that each member of the population had the same chance of being chosen. That is, each taster has the same chance of preferring your pretzel, thus satisfying the independence and constant probabilities for the binomial distribution.) We have $p = 0.30$ (Hit) and $n = 20$ (the sample size), and so $P(R = 5 \mid \text{Hit}) = P_B(R = 5 \mid n = 20, p = 0.30)$. We can use Formula (9.1) or the table to find $P_B(R = 5 \mid n = 20, p = 0.30) = 0.179$.

The same argument can be made regarding $P(R = 5 \mid \text{Flop})$. Now the condition is that the pretzels are a flop. This means that only 10% of the population prefer your pretzel over the other. Thus, we have $p = 0.10$ in this case. This gives us $P(R = 5 \mid \text{Flop}) = P_B(R = 5 \mid n = 20, p = 0.10)$. From the table, this probability is 0.032.

We now have everything we need to do the calculations that are required by Bayes' theorem:

$$\begin{aligned} P(\text{New Pretzel a Hit} \mid 5 \text{ of } 20 \text{ Preferred New Pretzel}) &= P(\text{Hit} \mid R = 5) \\ &= \frac{P(R = 5 \mid \text{Hit}) P(\text{Hit})}{P(R = 5 \mid \text{Hit}) P(\text{Hit}) + P(R = 5 \mid \text{Flop}) P(\text{Flop})} \\ &= \frac{0.179(0.50)}{0.179(0.50) + 0.032(0.50)} \\ &= 0.848 \end{aligned}$$

Thus, this evidence (5 out of 20 people preferring your pretzel) is good news. Your posterior probability that your pretzel will be a hit is almost 85%. Of course, we did the analysis on the basis of prior probabilities being $P(\text{Hit}) = P(\text{Flop}) = 0.50$. If you had assessed different prior probabilities, your answer would be different, although in any case your probability that the pretzel is a hit increases with the evidence.

The Poisson Distribution

While the binomial distribution is particularly good for representing successes in several trials, the *Poisson* distribution is good for representing occurrences of a particular event over time or space. Suppose, for example, that you are interested in the number of customers who arrive at a bank in one hour. Clearly this is an uncertain quantity; there could be none, one, two, three, and so on. The Poisson distribution also may be appropriate for modeling the uncertainty surrounding the number of machine breakdowns in a factory over some period of time. Other Poisson applications include modeling the uncertain number of blemishes in a bolt of fabric or the number of chocolate chips in a chocolate chip cookie. The Poisson distribution requires the following:

- 1 Events can happen at any of a large number of places within the unit of measurement (hour, square yard, and so on), and preferably along a continuum.
- 2 At any specific point, the probability of an event is small. This simply means that the events do not happen too frequently. For example, we would be interested in a steady flow of customers to a bank, not a run on the bank.
- 3 Events happen independently of other events. In other words, the probability of an event at any one point is the same regardless of the time (or location) of other events.

- 4 The average number of events over a unit of measure (time or space) is constant no matter how far or how long the process has gone on.

Let X represent the uncertain number of events in a unit of time or space. Under the conditions given above, the probability that $X = k$ events is given by

$$P_P(X = k) = \frac{e^{-m} m^k}{k!} \tag{9.2}$$

where the subscript P indicates this is a Poisson probability, e is the constant 2.718 ... (the base of the natural logarithms), and m is a parameter that characterizes the distribution. In particular, m turns out to be both the expected number of events and the variance of the number of events. In symbols, $E(X) = m$ and $\text{Var}(X) = m$.

It is easy to calculate Poisson probabilities using Formula (9.2) and a good calculator. For example,

$$\begin{aligned} P_P(X = 2 \mid m = 1.5) &= \frac{e^{-1.5}(1.5)^2}{2!} \\ &= 0.251 \end{aligned}$$

Again, however, tables (Appendixes C and D) are available that give both individual Poisson probabilities as well as cumulative probabilities (the probability of k or fewer events) for different values of m . Using these tables is much like using the binomial tables. Find the value for m across the top and the value for k along the left side. Figure 9.2 illustrates the use of the Poisson table in Appendix C to confirm our answer above.

An Example: Blemishes in Fabric As a simple example, suppose that you are interested in estimating the number of blemishes in 200 yards of cloth. Based on earlier experience with the cloth manufacturer, you estimate that a blemish occurs (on average) every 27 yards. At a rate of 1 blemish per 27 yards, this amounts to an approximate 7.4 blemishes in the 200 yards of cloth.

Is the Poisson distribution appropriate? Condition 1 is satisfied — we are looking at a continuous 200 yards of fabric. Condition 2 also is satisfied; apparently there are

Figure 9.2
Using the Poisson table to find $P_P(X = 2 \mid m = 1.5)$.

k	m	1.4	1.5	1.6
0		.247	.223	.202
1		.345	.335	.323
2		.242	.251	.258
3		.113	.126	.138
.		.	.	.
.		.	.	.

only a few blemishes in the 200 yards. Conditions 3 and 4 both should be satisfied unless blemishes are created by some machine malfunction that results in many blemishes occurring together. Thus, the Poisson distribution appears to provide an appropriate model of the uncertainty in this problem.

The expected value of 7.4 suggests that we could use a Poisson distribution with $m = 7.4$. Probabilities can be calculated using Formula (9.2). For example, the probability of nine blemishes to the cloth is

$$\begin{aligned} P_p(X = 9 \mid m = 7.4) &= \frac{e^{-7.4}(7.4)^9}{9!} \\ &= 0.112 \end{aligned}$$

We also can confirm this answer by looking in Appendix C, although information from Appendix D may be more useful. For example, the probability of 4 or fewer blemishes is $P_p(X \leq 4 \mid m = 7.4) = 0.140$. We also can find the probability of more than 10 blemishes: $P_p(X > 10 \mid m = 7.4) = 1 - P_p(X \leq 10 \mid m = 7.4) = 1.000 - 0.871 = 0.129$. Although it is theoretically possible for there to be a very large number of blemishes, we can see that the probability of more than 18 is extremely low. In fact, it is less than 0.0005 (the probabilities in Appendixes C and D are rounded to three decimal places).

Soft Pretzels, Continued Let us continue with the problem of the soft pretzels. You introduced the pretzels, and they are doing quite well. You have been distributing the product through several stores as well as through one street vendor. This vendor has been able to sell an average of 20 pretzels per hour. He had tried a different location earlier but had to move; sales there were only some 8 pretzels per hour, which was not enough to support his business. Now you are ready to try a second vendor in an altogether different location. The new location could be a "good" one, meaning an average of 20 pretzels per hour, "bad" with an average of only 10 per hour, or "dismal" with an average of 6 per hour. You have carefully assessed the probabilities of good, bad, or dismal using the assessment techniques from Chapter 8. Your judgment is that the new location is likely (probability 0.7) to be a good one. On the other hand, it could be (probability 0.2) a bad location, and it is just possible (probability 0.1) that the sales rate will be dismal.

After having the new stand open for a week, or long enough to establish a presence in the neighborhood, you decide to run a test. In 30 minutes you sell 7 pretzels. Now what are your probabilities regarding the quality of your new location? That is, what are $P(\text{Good} \mid X = 7)$, $P(\text{Bad} \mid X = 7)$, and $P(\text{Dismal} \mid X = 7)$?

As in the binomial case, we are interested in finding posterior probabilities given some new evidence. We will use Bayes' theorem to solve the problem:

$$\begin{aligned} P(\text{Good} \mid X = 7) &= \frac{P(X = 7 \mid \text{Good}) P(\text{Good})}{P(X = 7 \mid \text{Good}) P(\text{Good}) + P(X = 7 \mid \text{Bad}) P(\text{Bad}) + P(X = 7 \mid \text{Dismal}) P(\text{Dismal})} \end{aligned}$$

We have our prior probabilities, $P(\text{Good}) = 0.7$, $P(\text{Bad}) = 0.2$, and $P(\text{Dismal}) = 0.1$. What about the probabilities $P(X = 7 \mid \text{Good})$, $P(X = 7 \mid \text{Bad})$ and $P(X = 7 \mid \text{Dismal})$?

First, note that we are talking about a 30-minute period. Thus, the expected number of sales is either 10, 5, or 3 per half hour, depending on whether the location is good, bad, or dismal. If the conditions for the Poisson distribution hold, then $P(X = 7 | \text{Good})$ is the Poisson probability of 7 occurrences when $m = 10$, $P(X = 7 | m = 10)$. From Appendix C, this probability is 0.090. Likewise, $P(X = 7 | \text{Bad})$ is the Poisson probability of 7 occurrences when $m = 5$, or 0.104. Finally, $P(X = 7 | \text{Dismal})$ is 0.022. Now we can plug these values into Bayes' theorem:

$$\begin{aligned} P(\text{Good} | X = 7) &= \frac{P(X = 7 | \text{Good}) P(\text{Good})}{P(X = 7 | \text{Good}) P(\text{Good}) + P(X = 7 | \text{Bad}) P(\text{Bad}) + P(X = 7 | \text{Dismal}) P(\text{Dismal})} \\ &= \frac{0.090(0.7)}{0.090(0.7) + 0.104(0.2) + 0.022(0.1)} \\ &= 0.733 \end{aligned}$$

Likewise, we can calculate $P(\text{Bad} | X = 7) = 0.242$, and $P(\text{Dismal} | X = 7) = 0.025$. The posterior probabilities $P(\text{Good} | X = 7)$ and $P(\text{Bad} | X = 7)$ are not much different from the corresponding prior probabilities, so you might conclude that this information did not tell you much. On the basis of the information, however, the probability of the new location being dismal is now quite small.

The Exponential Distribution

The Poisson and binomial distributions are examples of discrete probability distributions because the outcome can take on only specific "discrete" values. What about continuous uncertain quantities? For example, tomorrow's high temperature could be any value between, say, 0° and 100°F . The per-barrel price of crude oil at the end of the year 2000 might be anywhere between \$10 and \$100. As discussed in Chapter 7, it is more natural in these cases to speak of the probability that the uncertain quantity falls within some interval. In the weather example we might look at historical weather patterns and determine that on 50% of days during May in Columbus, Ohio, the daily high temperature has been 70° or lower. On the basis of this, we could assess $P(\text{High Temperature in Columbus on May 28th} \leq 60^\circ) = 0.50$.

In this section, we will look briefly at the *exponential* distribution for a continuous random variable. In fact, the exponential distribution is closely related to the Poisson. If in the Poisson we were considering the number of arrivals within a specified period of time, then the uncertain time between arrivals (T) has an exponential distribution. The two go hand in hand; if the four conditions listed for the Poisson hold, then the time (or space) between events follows an exponential distribution.

For the binomial and Poisson distributions above, we were able to express in Formulas (9.1) and (9.2) the probability of a specific value for the random variable.

The corresponding expression for a continuous random variable is the *density function*. This function shows the relative likelihood for the different values that the uncertain quantity can take. For the exponential, the density function is

$$f_E(t | m) = me^{-mt} \quad (93)$$

where m is the same average rate that we used in the Poisson and t represents the possible values for the uncertain quantity T . An exponential density function with $m = 2$ is illustrated in Figure 9.3. Recall from Chapter 7 that areas under the density function correspond to probabilities. Thus, the area from 0 to a in Figure 9.3 represents $P_E(T \leq a | m = 2)$, where the subscript E indicates that this is a probability from an exponential distribution.

The exponential distribution turns out to be an easy distribution to work with. Probabilities are calculated by using the following formulas:

$$P_E(T \leq a | m) = 1 - e^{-am} \quad P_E(T > a | m) = 1 - P(T \leq a | m) = e^{-am}$$

$$P_E(b < T \leq a | m) = P(T \leq a | m) - P(T \leq b | m) = e^{-bm} - e^{-am}$$

For example, we can calculate that $P_E(T > 15 \text{ Min} | m = 2 \text{ Arrivals per Hr}) = P_E(T > 0.25 \text{ Hr} | m = 2 \text{ Arrivals per Hr})$

$$= e^{-0.25(2)}$$

$$= 0.607$$

Moreover, we can check this particular probability by using the Poisson table! The outcome "The time until the next arrival is greater than 15 minutes" is equivalent to the outcome "There are no arrivals in the next 15 minutes." Thus,

$$P_E(T > 15 \text{ Min} | m = 2 \text{ Arrivals per Hr})$$


$$= P_p(X = 0 \text{ Arrivals in 15 Min} | m = 2 \text{ Arrivals per Hr})$$

$$= P_p(X = 0 \text{ Arrivals in 15 Min} | m = 0.5 \text{ Arrivals per 15 Min})$$

$$= 0.607$$

from Appendix C.

Figure 9.3
Exponential density function with parameter $m = 2$. The shaded area represents $P_E(T \leq a | m = 2)$.



The expected value of an exponential random variable is $E(T) = 1/m$, and the variance is $\text{Var}(T) = 1/m^2$. Remember, these are units of time. For example, if $m = 6$ events per hour, then the expected time until the next event is $1/6$ of an hour, or 10 minutes.

Soft Pretzels, Again The problem at hand is whether a pretzel can be prepared during the time between customers. If it takes 3.5 minutes to cook a pretzel, what is the probability that the next customer will arrive after it is finished, $P(T > 3.5 \text{ Min})$?

This is a difficult calculation because we do not know exactly what the rate m would be for an exponential distribution. But we can expand $P(T > 3.5 \text{ Min})$ by using the total probability formula:

$$\begin{aligned} P(T > 3.5 \text{ Min}) &= P_E(T > 3.5 \text{ Min} \mid m = 20/\text{Hr}) P(m = 20) \\ &\quad + P_E(T > 3.5 \text{ Min} \mid m = 10/\text{Hr}) P(m = 10) + \\ &\quad P_E(T > 3.5 \text{ Min} \mid m = 6/\text{Hr}) P(m = 6) \\ &= P_E(T > 0.0583 \text{ Hr} \mid m = 20/\text{Hr}) P(m = 20) + \\ &\quad P_E(T > 0.0583 \text{ Hr} \mid m = 10/\text{Hr}) P(m = 10) + \\ &\quad P_E(T > 0.0583 \text{ Hr} \mid m = 6/\text{Hr}) P(m = 6) \\ &= e^{-0.583(20)} P(m = 20) + e^{-0.0583(10)} P(m = 10) \\ &\quad + e^{-0.0583(6)} P(m = 6) \\ &= 0.3114 P(m = 20) + 0.5580 P(m = 10) + 0.7047 P(m = 6) \end{aligned}$$

Now we can substitute in the posterior probabilities that we calculated above, $P(\text{Good} \mid X = 7) = 0.733$, $P(\text{Bad} \mid X = 7) = 0.242$, and $P(\text{Dismal} \mid X = 7) = 0.025$:

$$\begin{aligned} P(T > 3.5 \text{ Min}) &= 0.3114(0.733) + 0.5580(0.242) + 0.7047(0.025) = \\ &0.3809 \end{aligned}$$

Thus, the probability is 0.3809 that the time between arrivals is greater than 3.5 minutes. Put another way, because $P(T < 3.5 \text{ min}) = 1 - 0.3809 = 0.6191$, we can see that the majority of customers will arrive before the next pretzel pops out of the oven.

The Normal Distribution

Another particularly useful continuous distribution is the *normal* distribution, which is the familiar bell-shaped curve. The normal distribution is particularly good for modeling situations in which the uncertain quantity is subject to many different sources of uncertainty or error. For example, in measuring something, errors may be introduced by a wide range of environmental conditions, equipment malfunctions, human error, and so on. Many measured biological phenomena (height, weight, length) often follow a bell-shaped curve that can be represented well with a normal distribution.

We will let Y represent an uncertain quantity that follows a normal distribution. If this is the case, the density function for Y is

$$f_N(y | \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(y-\mu)^2/2\sigma^2} \quad (9.4)$$

where μ and σ are parameters of the distribution and y represents the possible values that Y can take. In fact, it turns out that $E(Y) = \mu$ and $\text{Var}(Y) = \sigma^2$. Figure 9.4 illustrates a normal density function. As with the exponential, the area under the density function represents the probability that the random variable falls into the corresponding interval. For example, the shaded area in Figure 9.4 represents $P_N(a \leq Y \leq b | \mu, \sigma)$. Strictly speaking, a normal random variable can take values anywhere between plus and minus infinity. But the probabilities associated with values more than three or four standard deviations from the mean are negligible, so we often use the normal to represent values that have a restricted range (for example, weight or height, which can only be positive) as long as the extreme points are several standard deviations from the mean.


In the case of the normal distribution, there are no simple formulas for finding probabilities as there are for the exponential. A simple rule of thumb exists for the normal, however. The probability is approximately 0.68 that a normal random variable is within one standard deviation of the mean μ , and the probability is approximately 0.95 that it is within two standard deviations of the mean. In symbols:

$$P_N(\mu - \sigma \leq Y \leq \mu + \sigma) \approx 0.68$$

$$P_N(\mu - 2\sigma \leq Y \leq \mu + 2\sigma) \approx 0.95$$

Figure 9.4

Normal density function.
The shaded area represents
 $P_N(a \leq Y \leq b | \mu, \sigma)$.



These are handy approximations for normal random variables and sometimes are called the *Empirical Rule* because these probabilities are found to hold (approximately) in many real-world situations.

A table has been provided in Appendix E so that you can find cumulative normal probabilities. But there is a catch! The table is for a *standard normal* distribution, one that has $\mu = 0$ and $\sigma = 1$. To use this table, we must know how to convert from a normal with any mean and standard deviation to a standard normal. The conversion works by subtracting the mean and dividing by the standard deviation. In symbols,

$$\begin{aligned} P_N(Y \leq a \mid \mu, \sigma) &= P_N\left(Z \leq \frac{a - \mu}{\sigma} \mid \mu = 0, \sigma = 1\right) \\ &= P_N\left(Z \leq \frac{a - \mu}{\sigma}\right) \end{aligned}$$

In most cases, we will leave off the values $\mu = 0$, $\sigma = 1$ when we talk about the standard normal random variable Z .

For example, if Y has a normal distribution with mean 10 and variance 400, then the probability that Y is less than or equal to 35 is

$$\begin{aligned} P_N(Y \leq 35 \mid \mu = 10, \sigma^2 = 400) &= P_N\left(Z \leq \frac{35 - 10}{20}\right) \\ &= P_N(Z \leq 1.25) \end{aligned}$$

Find this probability by looking in Appendix E as illustrated in Figure 9.5. This table has cumulative probabilities listed for values of z from -3.50 up to 3.49 . Find $z = 1.25$, and read off $P(Z \leq 1.25) = 0.8944$.

Figure 9.5
Using the normal distribution table to find $P_N(Z \leq 1.25)$.

....	z	$P(Z \leq z)$
....
....	1.20	.8849
....	1.21	.8869
....	1.22	.8888
....	1.23	.8907
....	1.24	.8925
....	1.25	.8944
....	1.26	.8962
....	1.27	.8980
....	1.28	.8997
....	1.29	.9015
....

An Investment Example Consider three alternative investments, A, B, and C. Investing in C yields a sure return of \$40. Investment A has an uncertain return (X), which is modeled in this case by a normal distribution with mean \$50 and standard deviation \$10. Investment B's return (represented by Y) also is modeled with a normal distribution having a mean of \$59 and a standard deviation of \$20. On the basis of expected values, B is the obvious choice because \$59 is greater than \$50 or \$40. The decision tree is shown in Figure 9.6.

Although it is obvious that the expected payoff for B is greater than the sure \$40 for C, we might be interested in the probability that B's payoff will be less than \$40. To find this probability, we must convert the value of 40 in our problem to a *standardized* value. Subtract the mean and then divide by the standard deviation to get a standardized random variable Z ;

$$P_N(B \leq 40 \mid \mu = 59, \sigma = 20) = P_N\left(Z \leq \frac{40 - 59}{20}\right) = P_N(Z \leq -0.95)$$

From Appendix E we find that

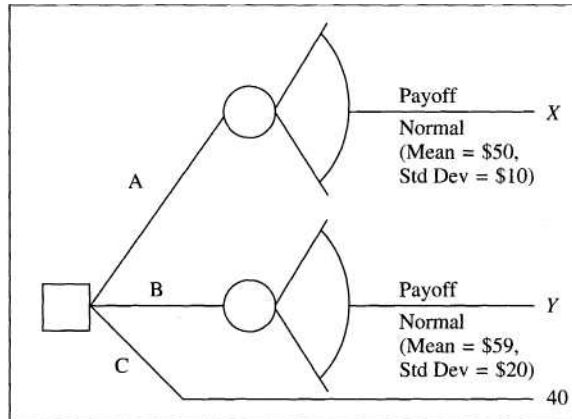
$$P(Z < -0.95) = 0.1711$$

Doing the same thing for Investment A, we want $P_N(A \leq \$40 \mid \mu = 50, \sigma = 10)$:

$$P_N(A \leq \$40 \mid \mu = 50, \sigma = 10) = P\left(Z \leq \frac{40 - 50}{10}\right) = P(Z \leq -1.0) = 0.1587$$

Thus, even though Investment A has a lower expected value than does B, A has a smaller probability of having a return of less than \$40. Why? The larger variance for B means that the distribution for B is spread out more than is the distribution for A.

Figure 9.6
Decision tree for three alternative investments.



We also might be interested in the probability of a particularly large return, say,

$$\begin{aligned} P_N(B > 78 \mid \mu = 59, \sigma = 20) &= P_N\left(Z > \frac{78 - 59}{20}\right) \\ &= P_N(Z > 0.95) \end{aligned}$$

Because Appendix E shows cumulative probabilities $P(Z \leq z)$, we must do extra work in this case. We can find the probability of the complement $P(Z < 0.95)$ and subtract this from 1:

$$\begin{aligned} P_N(Z > 0.95) &= 1 - P_N(Z \leq 0.95) \\ &= 1 - 0.8289 = \\ &0.1711 \end{aligned}$$


The probability $P(Z > 0.95)$ turns out to be the same as $P(Z \leq -0.95)$. This example points out the symmetry of the normal distribution. For any z , $P(Z > z) = P(Z \leq -z)$. To get our probability $P(Z > 0.95)$, we could have found it directly in the table by looking up $P(Z \leq -0.95)$. Figure 9.7 shows this property graphically.

Example: Quality Control Suppose that you are the manager for a manufacturing plant that produces disk drives for personal computers. One of your machines produces a part that is used in the final assembly. The width of this part is important to the disk drive's operation; if it falls below 3.995 or above 4.005 millimeters (mm), the disk drive will not work properly and must be repaired at a cost of \$10.40.

The machine can be set to produce parts with a width of 4 mm, but it is not perfectly accurate. In fact, the actual width of a part is normally distributed with mean 4 mm and a variance that depends on the speed of the machine. If the machine is run at a slower speed, the width of the produced parts has a standard deviation of 0.0019 mm. At the higher speed, however, the machine is less precise, producing parts with a standard deviation of 0.0026 mm.

Of course, the higher speed means that more parts can be produced in less time, thus reducing the overall cost of the disk drive. In fact, it turns out that the cost of the

Figure 9.7
The symmetry of the normal distribution.



disk drive when the machine is set at high speed is \$20.45. At low speed, the cost is \$20.75.

The question that you as plant manager face is whether it would be better to run the machine at high or low speed. Is the extra expense of lower speed more than offset by the increased precision and hence the lower defect rate? You would like to choose the strategy with the lower expected cost. Your decision tree is shown in Figure 9.8.

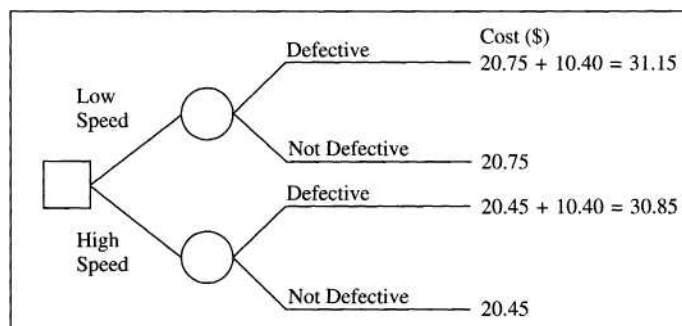
To decide, we need to know the probability of a defective unit under both machine settings. Because the width of the part follows a normal distribution in each case, we can get these probabilities by calculating z values and using Appendix E:

$$\begin{aligned}
 P(\text{Defective} \mid \text{Low Speed}) &= 1 - P(\text{Not Defective} \mid \text{Low Speed}) \\
 &= 1 - P_N(3.995 \leq Y \leq 4.005 \mid \mu = 4, \sigma = 0.0019) \\
 &= 1 - P\left(\frac{3.995 - 4}{0.0019} \leq Z \leq \frac{4.005 - 4}{0.0019}\right) \\
 &= 1 - P(-2.63 \leq Z \leq 2.63) \\
 &= 1 - [P(Z \leq 2.63) - P(Z \leq -2.63)] \\
 &= 1 - (0.9957 - 0.0043) \\
 &= 1 - 0.9914 \\
 &= 0.0086
 \end{aligned}$$

Likewise, we can calculate $P(\text{Defective} \mid \text{High Speed})$:

$$\begin{aligned}
 P(\text{Defective} \mid \text{High Speed}) &= 1 - P(\text{Not Defective} \mid \text{High Speed}) \\
 &= 1 - P_N(3.995 \leq Y \leq 4.005 \mid \mu = 4, \sigma = 0.0026) \\
 &= 1 - P\left(\frac{3.995 - 4}{0.0026} \leq Z \leq \frac{4.005 - 4}{0.0026}\right) \\
 &= 1 - P(-1.92 \leq Z \leq 1.92) \\
 &= 1 - [P(Z \leq 1.92) - P(Z \leq -1.92)] \\
 &= 1 - (0.9726 - 0.0274) \\
 &= 1 - 0.9452 \\
 &= 0.0548
 \end{aligned}$$

Figure 9.8
Decision tree for a
quality-control
problem.



With these probabilities, it now is possible to calculate the expected cost for each alternative:

$$\begin{aligned} E(\text{Cost} \mid \text{Low Speed}) &= 0.9914(\$20.75) + 0.0086(\$31.15) \\ &= \$20.84 \\ E(\text{Cost} \mid \text{High Speed}) &= \\ &0.9452(\$20.45) + 0.0548(\$30.85) \\ &= \$21.02 \end{aligned}$$

Thus, in this case the increased cost from the slower speed is more than offset by the increased precision and lower defect rate, and you would definitely choose the slower speed.

The Beta Distribution

Suppose you are interested in the proportion of voters who will vote for the Republican candidate in the next presidential election. If you are uncertain about this proportion, you may want to encode your uncertainty as a continuous probability distribution. Because the proportion can take only values between 0 and 1, neither the exponential nor the normal distribution can adequately reflect the uncertainty you face. The *beta* distribution, however, may be appropriate. Let Q denote an uncertain quantity that can take any value between 0 and 1. Then the beta density function is

$$f_{\beta}(q \mid r, n) = \frac{(n-1)!}{(r-1)!(n-r-1)!} q^{r-1}(1-q)^{n-r-1} \quad (9.5)$$

As before, q represents the possible values between 0 and 1 that Q can take.

The numbers r and n are parameters that determine the shape of the density function. If n is large, the distribution is fairly "tight," whereas if n is small, the distribution is more "spread out" (Figure 9.9). If $r = n/2$, the density function is symmetric around 0.5. If this is not the case, however, the distribution is skewed to the right or left depending on whether $r < n/2$ or $r > n/2$, as in Figure 9.10. As usual with density functions, the area under the curve represents probability. Thus, in Figure 9.10 the shaded area represents $P_{\beta}(0.2 \leq Q \leq 0.4 \mid n = 6, r = 4)$. A table in Appendix F provides cumulative probabilities for a wide variety of different beta distributions. We will demonstrate this table's use shortly.

Formula (9.5) for the density function looks much like the binomial distribution [Formula (9.1)]. Keep in mind that the beta distribution is a distribution for Q , a continuous random variable, whereas the binomial distribution is a distribution for R , the number of successes in n trials. Thus, we are considering two entirely different uncertain quantities. The two distributions, however, are closely related.

The expected value of a beta random variable is $E(Q) = r/n$, and the variance is $\text{Var}(Q) = r(n-r)/[n^2(n+1)]$. Looking at the formula for the expected value,

Figure 9.9
Some symmetric beta distributions.





Figure 9.10
Some asymmetric beta distributions.



r/n , the relationship between the beta and the binomial becomes apparent. Loosely speaking, r and n still can be interpreted as r successes in n trials. For example, if you had observed 4 successes in 10 trials from a binomial distribution with unknown proportion Q , your best guess for Q would be $4/10$, or 0.40 . Moreover, you might concede that Q might not be exactly 0.40 , although it should be close to 0.40 .

One way to use a beta distribution to model your subjective beliefs about an uncertain Q is to imagine a sample that would be roughly equivalent to your information. In the case of the proportion of voters who will vote Republican, you may have a feeling that the proportion is "around 0.30 ," and perhaps you feel as if all of your information (reading newspapers, talking to friends) is roughly equivalent to having polled a random sample of 20 people. These beliefs about Q might be represented by a beta distribution with $r = 6$ and $n = 20$.

Another way to fit a beta distribution is to use the cumulative probability table for the beta distribution (Appendix F). This table works a lot like the others. First, find the appropriate r and n along the left margin; for example, look at a distribution with $r = 6$ and $n = 20$. Read across the row to find the cumulative probabilities for values of q from 0 to 1 (column headings). For example, reading across, you find the number ".11" in the $q = 0.18$ column (Figure 9.11). This means that for a beta distribution with $n = 20$ and $r = 6$, $P_{\beta}(Q \leq 0.18 | n = 20, r = 6) = 0.11$, or there is an 11% chance that Q is less than 0.18. Likewise, you can see that $P_{\beta}(Q \leq 0.28 | n = 20, r = 6) = 0.45$, and $P_{\beta}(Q \leq 0.30 | n = 20, r = 6) = 0.53$. From these two values, we can deduce that the median for this distribution is approximately 0.29. That is, $P_{\beta}(Q < 0.29 | n = 20, r = 6)$ is approximately 0.50. We also can see that $P_{\beta}(Q \leq 0.48 | n = 20, r = 6) = 0.95$. This statement means there is a 95% chance that Q is less than 0.48 in this distribution, and hence only a 5% chance that Q is greater than 0.48.

To use the table to fit a beta distribution to your beliefs, you might assess the median and upper and lower quartiles of your subjective distribution for Q using the techniques discussed in Chapter 8 for assessing a continuous subjective probability distribution. Then look in the table for a combination of r and n that gives a beta distribution with the same (or nearly the same) median and quartiles. For example, suppose you assess the median of your subjective distribution to be 0.40, the lower quartile to be 0.29, and the upper quartile to be 0.52. That is,

$$P(Q \leq 0.29) = 0.25$$

$$P(Q \leq 0.40) = 0.50$$

$$P(Q \leq 0.52) = 0.75$$

Looking in the table, a beta distribution with $r = 4$ and $n = 10$ has approximately the same median and quartiles:

$$P_{\beta}(Q \leq 0.29 | n = 10, r = 4) = 0.25$$

$$P_{\beta}(Q \leq 0.40 | n = 10, r = 4) = 0.52$$

$$P_{\beta}(Q \leq 0.52 | n = 10, r = 4) = 0.78$$

Given the closeness between this beta distribution and your subjective assessments, you might find it useful to represent your subjective beliefs with the distribution.

Figure 9.11

Using the beta distribution table to find $P_{\beta}(Q \leq 0.18 | n = 20, r = 6) = 0.11$.

n	r	Values for Q		
	 0.16	0.18	0.20
.
20	2	.83	.88	.92
	4	.36	.45	.54
	607	.11	.16
	8	.01	.01	.02
.
.

Suppose, however, that you had assessed

$$P(Q \leq 0.36) = 0.25 \quad P(Q \leq 0.40) = 0.50 \quad P(Q \leq 0.52) = 0.75$$

No beta distributions have characteristics that are very close to these. The one with $r = 2$ and $n = 5$ has (roughly) the correct median and upper quartile: $P_\beta(Q \leq 0.40 | r = 2, n = 5) = 0.52$ and $P_\beta(Q \leq 0.52 | r = 2, n = 5) = 0.72$. On the lower end, however, the approximation is poor: $P_\beta(Q \leq 0.36 | r = 2, n = 5) = 0.45$. In this case, we would have three choices. We could be satisfied with a not-so-great approximation, use our subjective assessment directly (perhaps with a discrete approximation), or find another way to approximate our beliefs.

Soft Pretzels One Last Time Let us return to the issue of how much market share your pretzel will capture. Denote the market share by Q ; it is an uncertain quantity that must be between 0 and 1. You might want to model your beliefs about Q with a beta distribution. You know that Q is not likely to be close to 1. Suppose also that you think that the median is approximately 0.20 and the upper quartile is 0.38. To represent these assessments, you might choose a beta distribution with $r = 1$ and $n = 4$. From the table, you can see that this distribution for Q has a median around 0.21 and an upper quartile around 0.37:

$$P_\beta(Q \leq 0.20 | r = 1, n = 4) = 0.49 \quad P_\beta(Q \leq 0.38 | r = 1, n = 4) = 0.76$$

The expected value of Q is $r/n = 1/4 = 0.25$.

Should you become a pretzel entrepreneur? Recall that the decision was to price the pretzels at 50 cents each and that the total market is estimated to be 100,000 pretzels. You figure that the total fixed cost of the project amounts to \$8000 for marketing, finance costs, and overhead. The variable cost of each pretzel is 10 cents.

You only have one year to prove yourself, and then you will go on to graduate school no matter what happens. It also is apparent that this venture alone will not provide enough income to live on, but embarking on it will not interfere with other plans for gainful employment. Your rich uncle has agreed to help you financially, and you have savings of your own. The only question is whether the net contribution of the project will be positive or negative.

Write the equation for net contribution as the difference between total revenue and total costs:

$$\begin{aligned} \text{Net Contribution} &= 100,000(0.50)Q - 100,000(0.10)Q - 8000 = \\ &= 100,000(0.40)Q - 8000 \end{aligned}$$

The marginal contribution per pretzel sold is 40 cents, or 0.40. The expected net contribution is

$$\begin{aligned} E(\text{Net Contribution}) &= 40,000E(Q) - 8000 = \\ &= 40,000(0.25) - 8000 = \\ &= 10,000 - 8000 = 2000 \end{aligned}$$

On the basis of expected value, becoming a pretzel entrepreneur is a good idea! But the project also is a pretty risky proposition. What is the probability, for instance, that the pretzels could result in a loss? To answer this, we can find the specific value for Q (call it q^*) that would make the return equal to zero:

$$0 = 40,000q^* - 8000$$

$$q^* = \frac{8000}{40,000} = 0.20$$

If Q turns out to be less than 0.20, you would have been better off with the savings account. We assessed the median of Q to be approximately 0.20. This means that there is just about a 50% chance that the pretzel project will result in a loss. Are you willing to go ahead with the pretzels?

Probability Distributions and Decision-Analysis Software

It should come as no surprise that decision-analysis software typically includes facilities for working with theoretical distributions like those we have described in this chapter. For example, DPL provides access to over 20 different continuous and discrete distributions. DATA incorporates 15 different continuous distributions. Both programs use discrete-approximation techniques to represent the distributions. In particular, DATA has a built-in way to construct bracket medians. DPL builds the discrete approximation by means of a powerful and sophisticated "moment-matching" method that creates a discrete distribution with mathematical characteristics (expected value, variance, and so on) which match the theoretical distribution.

SUMMARY

In this chapter we examined ways to use theoretical probability distributions in decision-analysis problems. The distributions we considered were the binomial, Poisson, exponential, normal, and beta. These are only a few of the simpler ones; many other theoretical distributions are available. Each is appropriate in different situations, and a decision analyst must develop expertise in recognizing which distribution provides the best model of the uncertainty in a decision situation.

EXERCISES

- 9.1 In the binomial example with Melissa Bailey (page 318), what is the probability that she will be gone six or more weekends?
- 9.2 Suppose you are interested in an investment with an uncertain return. You think that the return could be modeled as a normal random variable with mean \$2000 and standard deviation \$1500. What is the probability that the investment will end up with a loss? What is the probability that the return will be greater than \$4000?

9.3 If there are, on average, 3.6 chocolate chips per cookie, what is the probability of finding no chocolate chips in a given cookie? Fewer than 5 chocolate chips? More than 10 chips?

9.4 Refer to the discussion of the pretzel problem in the section on the beta distribution (pages 334-335). Find the probability that the net contribution of the pretzel project would be greater than \$4500.

9.5 Table look-up and calculation practice

a Binomial distribution. Find the following probabilities:

$$\begin{array}{ll}
 P_B(R = 5 \mid n = 10, p = 0.22) & P_B(R = 10 \mid n = 15, p = 0.63) \\
 P_B(R \leq 1 \mid n = 7, p = 0.04) & P_B(R \leq 3 \mid n = 10, p = 0.35) \\
 P_B(7 < R \leq 10 \mid n = 12, p = 0.50) & P_B(R > 1 \mid n = 19, p = 0.06) \\
 P_B(R = 4 \mid n = 6, p = 0.42) & P_B(R < 2 \mid n = 4, p = 0.67)
 \end{array}$$

b Poisson distribution. Find the following probabilities:

$$\begin{array}{ll}
 P_P(X = 3 \mid m = 2.0) & P_P(X > 17 \mid m = 15) \\
 P_P(X > 4 \mid m = 5.2) & P_P(X < 7 \mid m = 10.0) \\
 P_P(X \leq 1 \mid m = 3.9) & P_P(3 \leq X < 7 \mid m = 1.5) \\
 P_P(X = 4 \mid m = 1.75) & P_P(X \leq 2 \mid m = 3.56)
 \end{array}$$

c Exponential distribution. Find the following probabilities:

$$\begin{array}{ll}
 P_E(T \leq 5 \mid m = 1) & P_E(T \geq 4 \mid m = 0.25) \\
 P_E(0.25 < T < 1 \mid m = 2) & P_E(T \geq 3.2 \mid m = 2) \\
 P_E(T = 3.2 \mid m = 2) & P_E(T \geq k \mid m = 1/k)
 \end{array}$$

d Normal distribution. Find the following probabilities:

$$\begin{array}{ll}
 P_N(Y < 12 \mid \mu = 10, \sigma = 4) & P_N(Y > 50 \mid \mu = 100, \sigma = 30) \\
 P_N(100 < Y < 124 \mid \mu = 133, \sigma = 15) & P_N(Y > 20 \mid \mu = 15, \sigma = 4) \\
 P_N(Y > 0 \mid \mu = -10, \sigma = 4) & P_N(-44 < Y < 12 \mid \mu = 10, \sigma = 20) \\
 P_N(Y = 12 \mid \mu = 10, \sigma = 4) & P_N(Y < 12 \mid \mu = 10, \sigma = 4)
 \end{array}$$

e Beta distribution. Find the following probabilities:

$$\begin{array}{ll}
 P_\beta(Q \leq 0.9 \mid n = 10, r = 9) & P_\beta(Q \geq 0.5 \mid n = 5, r = 2) \\
 P_\beta(Q \leq 0.44 \mid n = 30, r = 20) & P_\beta(0.04 < Q \leq 0.38 \mid n = 18, r = 8) \\
 P_\beta(0.12 < Q \leq 0.25 \mid n = 14, r = 4) & P_\beta(Q = 0.76 \mid n = 18, r = 14)
 \end{array}$$

9.6 Find values for z such that

$$\begin{array}{ll}
 P_N(Z \leq z \mid \mu = 0, \sigma = 1) = 0.05 & P_N(Z > z \mid \mu = 0, \sigma = 1) = 0.25 \\
 P_N(Z \leq z \mid \mu = 0, \sigma = 1) = 0.50 & P_N(Z > z \mid \mu = 0, \sigma = 1) = 0.10
 \end{array}$$

9.7 Find the parameters (μ and σ) for a normal distribution whose first and third quartiles are 125 and 275. That is, $P_N(Y \leq 125 \mid \mu, \sigma) = 0.25$ and $P_N(Y \leq 275 \mid \mu, \sigma) = 0.75$. What are μ and σ ?

- 9.8** An exponential distribution has $P_E(T \geq 5 | m) = 0.24$. Find m .
- 9.9** A Poisson distribution has $P_P(X \geq 12 | m) = 0.01$ and $P_P(X \leq 2 | m) = 0.095$. Find m .
- 9.10** A Poisson distribution has $P_P(X = 0 | m) = 0.175$. Calculate m .
- 9.11** Use Appendix E to verify the Empirical Rule (page 327) for normal distributions, which states that the probability is approximately 0.68 that a normal random variable is within one standard deviation of the mean M , and the probability is approximately 0.95 that the random variable is within two standard deviations of the mean.

QUESTIONS AND PROBLEMS

- 9.12** The amount of time that a union stays on strike is judged to follow an exponential distribution with a mean of 10 days.
- a** Find the probability that a strike lasts less than one day. **b** Find the probability that a strike lasts less than six days. **c** Find the probability that a strike lasts between six and seven days.
- d** Find the conditional probability that a strike lasts less than seven days, given that it already has lasted six days. Compare your answer to part a.
- 9.13** A photographer works part-time in a shopping center, two hours per day, four days per week. On the average, six customers arrive each hour, and the arrivals appear to occur independently of one another. Twenty minutes after she arrives one day, the photographer wonders what the chances are that exactly one customer will arrive during the next 15 minutes. Find this probability (i) if two customers just arrived within the first 20 minutes and (ii) if no customers have come into the shop yet on this particular day.
- 9.14** A consumer is contemplating the purchase of a new compact disc player. A consumer magazine reports data on the major brands. Brand A has lifetime (T_A) which is exponentially distributed with $m = 0.2$; and Brand B has lifetime (T_B), which is exponentially distributed with $m = 0.1$. (The unit of time is one year.)
- a** Find the expected lifetimes for A and B. If a consumer must choose between the two on the basis of maximizing expected lifetime, which one should be chosen?
- b** Find the probability that A's lifetime exceeds its expected value. Do the same for B. What do you conclude?
- c** Suppose one consumer purchases a Brand A compact disc, and another consumer purchases a Brand B compact disc. Find the mean and variance of (i) the average lifetime of the two machines and (ii) the difference between the lifetimes of the two machines. (*Hint:* You must use the rules about means and variances of linear transformations that we discussed in Chapter 7.)
- 9.15** On the basis of past data, the owner of an automobile dealership finds that, on average, 8.5 cars are sold per day on Saturdays and Sundays during the months of January and February, with the sales rate relatively stable throughout the day. Moreover, purchases appear to be independent of one another. The dealership is open for 10 hours per day on

each of these days. There is no reason to believe that the sales for the upcoming year will be any different than in the past.

- a On the first Saturday in February, the dealership will open at 9 A.M. Find the probability that the time until the first sale is more than two hours, $P_e(T \geq 2 \text{ hours} \mid m = 8.5 \text{ cars per 10 hours})$.
 - b Find the probability that the number of sales before 11 A.M. is equal to zero, $P_p(X = 0 \text{ in 2 hours} \mid m = 8.5 \text{ cars per 10 hours})$. Compare your answer to that from part a. Can you explain why the answers are the same?
 - c The owner of the dealership gives her salespeople bonuses, depending on the total number of cars sold. She gives them \$20 whenever exactly 13 cars are sold on a given day, \$30 whenever 14 cars are sold, \$50 whenever 15 cars are sold, and \$70 whenever 16 or more cars are sold. On any given Saturday or Sunday in January or February, what is the expected bonus that the dealer will have to pay?
 - d Consider the bonus scheme presented in part c. February contains exactly four Saturdays and four Sundays. What is the probability that the owner will have to pay the \$20 bonuses exactly twice in those days?
- 9.16** Reconsider your assessed 0.05 and 0.95 fractiles in Problem 8.13. If you are perfectly *calibrated* when judging these fractiles, then you would expect that in any given situation the actual value has a 0.90 chance of falling between your assessed fractiles (for that variable).
- a Assuming you are perfectly calibrated, what is the probability that 0, 1, or 2 of the 10 actual values fall between the assessed fractiles?
 - b To justify using the binomial distribution to answer part a, you must assume that the chance of any particular value falling between its fractiles is the same regardless of what happens with the other variables. Do you think this is a reasonable assumption? Why or why not?
- 9.17** In the soft pretzels example concerning the taste test — in which we used binomial probabilities — suppose the results had been that 4 out of the 20 people sampled preferred your pretzel. Use Bayes' theorem to find your posterior probability $P(\text{Hit} \mid r = 4, n = 20)$ for the following pairs of prior probabilities:
- a $P(\text{Hit}) = 0.2, P(\text{Flop}) = 0.8$
 - b $P(\text{Hit}) = 0.4, P(\text{Flop}) = 0.6$
 - c $P(\text{Hit}) = 0.5, P(\text{Flop}) = 0.5$
 - d $P(\text{Hit}) = 0.75, P(\text{Flop}) = 0.25$
 - e $P(\text{Hit}) = 0.90, P(\text{Flop}) = 0.10$
 - f $P(\text{Hit}) = 1.0, P(\text{Flop}) = 0$

Create a graph of the posterior probability as a function of the prior probability. (We did something similar in Chapter 7 in discussing the Hinckley trial. Refer to Figure 7.20.)

- 9.18** In a city, 60% of the voters are in favor of building a new park. An interviewer intends to conduct a survey.
- a If the interviewer selects 20 people randomly, what is the probability that more than 15 of them will favor building the park?
 - b Instead of choosing 20 people as in part a, suppose that the interviewer wants to conduct the survey until he has found exactly 12 who are in favor of the park. What is the

probability that the first 12 people surveyed all favor the park (in which case the interviewer can stop)? What is the probability that the interviewer can stop after interviewing the thirteenth subject? What is the probability that the interviewer can stop after interviewing the eighteenth subject?

- 9.19** In bottle production, bubbles that appear in the glass are considered defects. Any bottle that has more than two bubbles is classified as "nonconforming" and is sent to recycling. Suppose that a particular production line produces bottles with bubbles at a rate of 1.1 bubbles per bottle. Bubbles occur independently of one another.
- What is the probability that a randomly chosen bottle is nonconforming?
 - Bottles are packed in cases of 12. An inspector chooses one bottle from each case and examines it for defects. If it is nonconforming, she inspects the entire case, replacing nonconforming bottles with good ones. This process is called rectification. If the chosen bottle conforms (has two or fewer bubbles), then she passes the case. In total, 20 cases are produced. What is the probability that at least 18 of them pass?
 - What is the expected number of nonconforming bottles in the 20 cases after they have been inspected and rectified using the scheme described in part b?
- 9.20** In our discussion of the Poisson distribution, we used this distribution to represent the process by which customers arrive at the pretzel stand. Is it reasonable to assume that the Poisson distribution is appropriate for finding the probabilities that we need? Why or why not?
- 9.21** You are the mechanical engineer in charge of maintaining the machines in a factory. The plant manager has asked you to evaluate a proposal to replace the current machines with new ones. The old and new machines perform substantially the same jobs, and so the question is whether the new machines are more reliable than the old. You know from past experience that the old machines break down roughly according to a Poisson distribution, with the expected number of breakdowns at 2.5 per month. When one breaks down, \$150 is required to fix it. The new machines, however, have you a bit confused. According to the distributor's brochure, the new machines are supposed to break down at a rate of 1.5 machines per month on average and should cost \$170 to fix. But a friend in another plant that uses the new machines reports that they break down at a rate of approximately 3.0 per month (and do cost \$170 to fix). (In either event, the number of breakdowns in any month appears to follow a Poisson distribution.) On the basis of this information, you judge that it is equally likely that the rate is 3.0 or 1.5 per month.
- Based on minimum expected repair costs, should the new machines be adopted?
 - Now you learn that a third plant in a nearby town has been using these machines. They have experienced 6 breakdowns in 3.0 months. Use this information to find the posterior probability that the breakdown rate is 1.5 per month.
 - Given your posterior probability, should your company adopt the new machines in order to minimize expected repair costs?
 - Consider the information given in part b. If you had read it in the distributor's brochure, what would you think? If you had read it in a trade magazine as the result of an independent test, what would you think? Given your answers, what do you think about using sample information and Bayes' theorem to find posterior probabilities? Should the source of the information be taken into consideration somehow? Could this be done in some way in the application of Bayes' theorem?

- 9.22** In the soft pretzels example for the beta distribution, suppose you assessed the median of Q to be approximately 0.30 and the upper quartile 0.50. In this case you could use a beta distribution with parameters $r = 1$ and $n = 3$ to represent your beliefs.
- Plug the values $r = 1$ and $n = 3$ into the formula for the beta distribution. The expression simplifies considerably. (Recall that $0! = 1$ by definition.)
 - Draw a graph of the distribution of Q .
 - Use what you know about areas of triangles to show that this distribution is a reasonably good representation of your subjective beliefs in the sense that the median and upper quartile are fairly close to your subjectively assessed median and upper quartile.
 - Find the expected profit for the pretzel project with this new distribution.
 - Find the probability that the pretzel project results in a loss under this distribution.
- 9.23** Sometimes we use probability distributions that are not exact representations of the physical processes that they are meant to represent. (For example, we might use a normal distribution for a distribution of individuals' weights, even though no one can weigh less than zero pounds.) Why do we do this?
- 9.24** You are an executive at Procter and Gamble and are about to introduce a new product. Your boss has asked you to predict the market share (Q , a proportion between 0 and 1) that the new product will capture. You are unsure of Q , and you would like to communicate your uncertainty to the boss. You have made the following assessments:
- There is a 1-in-10 chance that Q will be greater than 0.22, and also a 1-in-10 chance that Q will be less than 0.08.
- The value for Q is just as likely to be greater than 0.14 as less than 0.14.
- What should your subjective probabilities $P(0.08 < Q < 0.14)$ and $P(0.14 < Q < 0.22)$ be in order to guarantee coherence?
 - Use Appendix F to find a beta distribution for Q that closely approximates your subjective beliefs.
 - The boss tells you that if you expect that the market share will be less than 0.15, the product should not be introduced. Write the boss a memo that gives an expected value and also explains how risky you think it would be to introduce the product. Use your beta approximation.
- 9.25** Suppose you are considering two investments, and the critical issues are the rates of return (R_1 and R_2). For Investment 1, the expected rate of return (μ_1) is 10%, and the standard deviation (σ_1) is 3%. For the second investment, the expected rate of return (μ_2) is 20%, and the standard deviation (σ_2) is 12%.
- Does it make sense to decide between these two investments on the basis of expected value alone? Why or why not?
 - Does it make sense to represent the uncertainty surrounding the rates of return with normal distributions? What conditions do we need for the normal distribution to provide a good fit?
 - Suppose you have decided to use normal distributions (either because of or in spite of your answer to part b). Use Appendix E to find the following probabilities:

$$P(R_1 < 0\%)$$

$$P(R_2 < 0\%)$$

$$P(R_1 > 20\%)$$

$$P(R_2 < 10\%)$$

- d** How can you find the probability that $R_1 > R_2$? Suppose R_1 and R_2 are correlated (as they would be if, say, both of the investments were stocks). Then the random variable $\Delta R = R_1 - R_2$ is normal with mean $\mu_1 - \mu_2$ and variance $\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2$, where ρ is the correlation between R_1 and R_2 . If $\rho = 0.5$, find $P(R_1 > R_2)$. (*Hint*: Think about it in terms of ΔR , and find $P(\Delta R > 0)$.)
- e** How could you use the information from the various probabilities developed in this problem to choose between the two investments?
- 9.26** Your inheritance, which is in a blind trust, is invested entirely in McDonald's or in U.S. Steel. Because the trustee owns several McDonald's franchises, you believe the probability that the investment is in McDonald's is 0.8. In any one year, the return from an investment in McDonald's is approximately normally distributed with mean 14% and standard deviation 4%, while the investment in U.S. Steel is approximately normally distributed with mean 12% and standard deviation 3%. Assume that the two returns are independent.
- a** What is the probability that the investment earns between 6% and 18% (i) if the trust is invested entirely in McDonald's, and (ii) if the trust is invested entirely in U.S. Steel?
- b** Without knowing how the trust is invested, what is the probability that the investment earns between 6% and 18%?
- c** Suppose you learn that the investment earned more than 12%. Given this new information, find your posterior probability that the investment is in McDonald's.
- d** Suppose that the trustee decided to split the investment and put one-half into each of the two securities. Find the expected value and the variance of this portfolio.
- 9.27** A continuous random variable X has the following density function:

$$f(x) = \begin{cases} 0.5 & \text{for } 3 \leq x \leq 5 \\ 0 & \text{otherwise} \end{cases}$$

$$P_U(X \leq 4.3 \mid a = 3, b = 5) P_U(0.25 \leq X \leq 0.75 \mid a = 0, b = 1)$$

- d** Plot the CDF for the uniform distribution where $a = 0, b = 1$.

e The expected value of a uniform distribution is $E(X) = (b + a)/2$, and the variance is $\text{Var}(X) = (b - a)^2/12$. Calculate the expected value and variance of the uniform density with $a = 3$, $b = 5$.

9.28 The length of time until a strike is settled is distributed uniformly from 0 to 10.5 days. (See the previous problem for an introduction to the uniform density.)

a Find the probability that a strike lasts less than one day. **b** Find the probability that a strike lasts less than six days. **c** Find the probability that a strike lasts between six and seven days.

d Find the conditional probability that a strike lasts less than seven days, given that it already has lasted six days.

e Compare your answers with those of Problem 9.12.

9.29 In a survey in a shopping center, the interviewer asks customer how long their shopping trips have lasted so far. The response (T) given by a randomly chosen customer is uniformly distributed from 0 to 1.5 hours.

a Find the probability that a customer has been shopping for 36 minutes or less.

b The interviewer surveys 18 customers at different times. Find the probability that more than one-half of these customers say that they have been shopping for 36 minutes or less.

9.30 A continuous random variable X has the following density function:

$$f(x) = \begin{cases} \frac{1}{3}x - \frac{2}{3} & \text{for } 2 \leq x \leq 4 \\ -\frac{2}{3}x + \frac{10}{3} & \text{for } 4 \leq x \leq 5 \\ 0 & \text{otherwise} \end{cases}$$

a Draw a graph of this density. Verify that the area under the density function equals 1.

b A density function such as this one is called a *triangular* density. It is almost as easy to work with as the uniform density; probabilities for intervals can be calculated easily by calculating the areas of triangles and quadrilaterals. For example, find $P(3 \leq X \leq 4.5)$ for this distribution.

c Find the value $x_{0.50}$ such that $P(X \leq x_{0.50}) = 0.5$. (That is, find the median of this distribution.)

d Find the upper and lower quartiles of this distribution. That is, find $x_{0.25}$ such that $P(X \leq x_{0.25}) = 0.25$, and find $x_{0.75}$ such that $P(X \leq x_{0.75}) = 0.75$.

9.31 A greeting card shop makes cards that are supposed to fit into 6-inch (in) envelopes. The paper cutter, however, is not perfect. The length of a cut card is normally distributed with mean 5.9 in and standard deviation 0.0365 in. If a card is longer than 5.975 in, it will not fit into a 6-in envelope.

a Find the probability that a card will not fit into a 6-in envelope.

b The cards are sold in boxes of 20. What is the probability that in one box there will be two or more cards that do not fit in 6-in envelopes?

9.32 You are the maintenance engineer for a plant that manufactures consumer electronic goods. You are just about to leave on your vacation for two weeks, and the boss is concerned about certain machines that have been somewhat unreliable, requiring your expertise to keep

them running. The boss has asked you how many of these machines you expect to fail while you are out of town, and you have decided to give him your subjective probability distribution. You have made the following assessments:

- 1 There is a 0.5 chance that none of the machines will fail.
- 2 There is an approximate 0.15 chance that two or more will fail.
- 3 There is virtually no chance that four or more will fail.

Being impatient with this slow assessment procedure, you decide to try to fit a theoretical distribution.

- a Many operations researchers would use a Poisson distribution in this case. Why might the Poisson be appropriate? Why might it not be appropriate?
 - b Find a Poisson distribution that provides a good representation of your assessed beliefs. Give a specific value for the parameter m .
 - c Given your answer to b, what is the expected number of machines that will break down during your absence?
- 9.33** After you have given your boss your information (Problem 9.32), he considers how accurate you have been in the past when you have made such assessments. In fact, he decides you are somewhat optimistic (and he believes in Murphy's Law), so *he assigns a Poisson distribution with $m = 1$* to the occurrence of machine breakdowns during your two-week vacation. Now the boss has a decision to make. He either can close the part of the plant involving the machines in question, at a cost of \$10,000, or he can leave that part up and running. Of course, if there are no machine failures, there is no cost. If there is only one failure, he can work with the remaining equipment until you return, so the cost is effectively zero. If there are two or more failures, however, there will be assembly time lost, and he will have to call in experts to repair the machines immediately. The cost would be \$15,000. What should he do?
- 9.34** Regarding the soft pretzels again, suppose you decide to conduct a taste test of a new recipe at your stand. On average, one person comes to the stand every four minutes, and the arrivals seem to follow a Poisson distribution fairly closely. You decide to check for 30 minutes to see how many customers during that time prefer the new recipe. Suppose the probability is 0.4 that any arriving customer prefers the new recipe over the old one. What is the probability that you will find four or more customers who prefer the new recipe during your 30-minute test period?
- 9.35** A factory manager must decide whether to stock a particular spare part. The part is absolutely essential to the operation of certain machines in the plant. Stocking the part costs \$10 per day in storage and cost of capital. If the part is in stock, a broken machine can be repaired immediately, but if the part is not in stock, it takes one day to get the part from the distributor, during which time the broken machine sits idle. The cost of idling one machine for a day is \$65. There are 50 machines in the plant that require this particular part. The probability that any one of them will break and require the part to be replaced on any one day is only 0.004 (regardless of how long since the part was previously replaced). The machines break down independently of one another.
- a If you wanted to use a probability distribution for the number of machines that break down on a given day, would you use the binomial or Poisson distribution? Why?

b Whichever theoretical distribution you chose in part a, what are appropriate parameters? That is, if you chose the binomial, what are the values for p and n ? If you chose the Poisson, what is the value for m ?

c If the plant manager wants to minimize his expected cost, should he keep zero, one, or two parts in stock? Draw a decision tree and solve the manager's problem. (Do not forget that more than one machine can fail in one day!)

9.36 Another useful distribution that is based on the normal is the *lognormal* distribution. Among other applications, this distribution is used by environmental engineers to represent the distribution of pollutant levels, by economists to represent the distribution of returns on investments, and by actuaries to represent the distribution of insurance claims.

Finding probabilities from a lognormal distribution is "as easy as falling off a log"! If X is lognormally distributed with parameters μ and σ , then $Y = \ln(X)$ is normally distributed and has mean μ and variance σ^2 . Thus, the simplest way to work with a lognormal random variable X is to work in terms of $Y = \ln(X)$. It is easy to obtain probabilities for Y from the normal table. The expected value and variance of X are given by the following formulas:

$$E(X) = e^{\mu + 0.5\sigma^2} \quad \text{Var}(X) = (e^{2\mu})(e^{\sigma^2} - 1)(e^{\sigma^2})$$

For example, if X is lognormally distributed with parameters $\mu = 0.3$ and $\sigma = 0.2$, then Y is normal with mean 0.3 and standard deviation 0.2. Finding probabilities just means taking logs:

$$P_L(X \geq 1.4 \mid \mu = 0.3, \sigma = 0.2) = P_N(Y \geq \ln(1.4) \mid \mu = 0.3, \sigma = 0.2)$$

$$= P_N(Y \geq 0.336 \mid \mu = 0.3, \sigma = 0.2) = P(Z \geq 0.18) = 0.4286$$

The mean and expected value of X are

$$E(X) = e^{0.3 + 0.5(0.2)^2}$$

$$= 1.38$$

$$E(X) = e^{0.3 + 0.5(0.2)^2}$$

$$= 1.38$$

$$\text{Var}(X) = (e^{2(0.3)})(e^{(0.2)^2} - 1)(e^{(0.2)^2})$$

$$= 0.077$$

After all that, here is a problem to work. After a hurricane, claims for property damage pour into the insurance offices. Suppose that an insurance actuary models noncommercial property damage claims (X , in dollars) as being lognormally distributed with parameters $\mu = 10$ and $\sigma = 0.3$. Claims on different properties are assumed to be independent.

a Find the mean and standard deviation of these claims.

b Find the probability that a claim will be greater than \$50,000.

c The company anticipates 200 claims. If the state insurance commission requires the company to have enough cash on hand to be able to satisfy all claims with probability 0.95, how much money should be in the company's reserve? [*Hint:* The total

claims can be represented by the variable $Q = \sum_{i=1}^{200} X_i$, and Q will be approximately normally distributed with mean $200 E(X)$ and variance $200 \text{Var}(X)$.]

C A S E S T U D I E S

OVERBOOKING

Most airlines practice *overbooking*. That is, they are willing to make more reservations than they have seats on an airplane. Why would they do this? The basic reason is simple; on any given flight a few passengers are likely to be "no-shows." If the airline overbooks slightly, then it still may be able to fill the airplane. Of course, this policy has its risks. If more passengers arrive to claim their reservations than there are seats available, the airline must "bump" some of its passengers. Often this is done by asking for volunteers. If a passenger with a reserved seat is willing to give up his or her seat, the airline typically will give a refund as well as provide a free ticket to the same or another destination. The fundamental trade-off is whether the additional expected revenue gained by flying an airplane that is nearer to capacity on average is worth the additional expected cost of refunds and free tickets.

To study the overbooking policy, let us look at a hypothetical situation. Mockingbird Airlines has a small commuter airplane with places for 16 passengers. The airline uses this jet on a route for which it charges \$225 for a one-way fare. Every flight has a fixed cost of \$900 (for pilot's salary, fuel, airport fees, and so on). Each passenger costs Mockingbird an additional \$100. Finally, the no-show rate is 4%. That is, on average approximately 4% of those passengers holding confirmed reservations do not show up. Refunds for unused tickets are made only if the reservation is canceled at least 24 hours before scheduled departure.

How many reservations should Mockingbird be willing to sell on this airplane? The strategy will be to calculate the expected profit for a given number of reservations. For example, suppose that the Mockingbird manager decides to sell 18 reservations. The revenue is \$225 times the number of reservations:

$$\begin{aligned} R &= \$225(18) \\ &= \$4050 \end{aligned}$$

The cost consists of two components. The first is the cost of flying the plane and hauling the passengers who arrive (but not more than the airplane's capacity of 16):

$$C_1 = \$900 + \$100 \times \text{Min}(\text{Arrivals}, 16)$$

The second component is the cost of refunds and free tickets that must be issued if 17 or 18 passengers arrive:

$$C_2 = (\$225 + \$100) \times \text{Max}(0, \text{Arrivals} - 16)$$

In this expression for C_2 , the \$225 represents the refund for the purchased ticket, and the \$100 represents the cost of the free ticket. The $\text{Max}()$ expression calculates the number of excess passengers who show up (zero if the number of arrivals is fewer than 16).

Questions

- 1 Find the probability that more than 16 passengers will arrive if Mockingbird sells 17 reservations ($\text{Res} = 17$). Do the same for 18 and 19.
- 2 Find:

$$E(R \mid \text{Res} = 16)$$

$$E(C_1 \mid \text{Res} = 16)$$

$$E(C_2 \mid \text{Res} = 16)$$

Finally, calculate $E(\text{Profit} \mid \text{Res} = 16) = E(R \mid \text{Res} = 16) - E(C_1 \mid \text{Res} = 16) - E(C_2 \mid \text{Res} = 16)$

Repeat Question 2 for 17, 18, and 19 reservations. What is your conclusion? Should Mockingbird overbook? By how much?

Since the airlines were deregulated in the 1970s, pricing has become more competitive. One of the promotional schemes is the "supersaver" fare that requires early payment and restrictions on refunds. For example, to receive the special fare, a customer may be required to purchase a nonrefundable ticket two weeks in advance, after which changes can be made only by paying a penalty (e.g., \$50). How do you think this policy has affected the airlines' overbooking policy?

EARTHQUAKE PREDICTION

Because of the potential damage and destruction that earthquakes can cause, geologists and geophysicists have put considerable effort into understanding when and where earthquakes occur. The ultimate aim is the accurate prediction of earthquakes on the basis of movements in the earth's crust, although this goal appears to be some way off. In the meantime, it is possible to examine past data and model earthquakes probabilistically.

Fortunately, considerable data exist on the basis of which to model earthquakes as a probabilistic phenomenon. Gere and Shah provide the information shown in Table 9.1. Richter magnitude refers to the severity of the earthquake. For example, if an earthquake is in the 8.0-8.9 category, by definition the ground would shake strongly for 30 to 90 seconds over an area with a diameter of 160 to 320 kilometers. Earthquakes of magnitude less than 4.0 are not dangerous and, for the most part, are not noticed by laypeople.

Table 9.1
Earthquake frequency
data for California.

Richter Magnitude	Average Number of Earthquakes per 100 Years in California
8.0-8.9	1
7.0-7.9	12
6.0-6.9	80
5.0-5.9	400
4.0-4.9	2000

Source: J. M. Gere and H. C. Shah (1984) *Terra Non Firma: Understanding and Preparing for Earthquakes*. Stanford, CA: Stanford Alumni Association.

An earthquake of magnitude 8.0 or greater could cause substantial damage and a large number of deaths if it were to occur in a highly populated part of the world. In fact, the San Francisco earthquake of April 6, 1906, was calculated later as measuring 8.3 on the Richter scale. The resulting fire burned much of the city, and some 700 people died. California is particularly susceptible to earthquakes because the state straddles two portions of the earth's crust that are slipping past each other, primarily along the San Andreas Fault. For this reason, we will consider the probability of a severe earthquake happening again in California in the near future.

Questions

- 1 We can model the occurrence of earthquakes using a Poisson distribution. Strictly speaking, the independence requirement for the Poisson is not met for two reasons. First, the geologic processes at work in California suggest that the probability of a large earthquake increases as time elapses following an earlier large quake. Second, large earthquakes often are followed by aftershocks. Our model will ignore these issues and hence can be viewed only as a first-cut approximation at constructing a probabilistic model for earthquakes.

The data from Gere and Shah indicate that, on average, 2493 earthquakes with magnitude 4.0 or greater will occur in California over a 100-year period. Thus, we might consider using a Poisson distribution with $m = 24.93$ to represent the probability distribution for the number of earthquakes (all magnitudes greater than 4.0) that will hit California during the next year. Use this distribution to find the following probabilities:

$$P_P(X \leq 10 \text{ in Next Year} \mid m = 24.93 \text{ Earthquakes per Year})$$

$$P_P(X \leq 7 \text{ in Six Months} \mid m = 24.93 \text{ Earthquakes per Year})$$

$$P_P(X > 3 \text{ in Next Month} \mid m = 24.93 \text{ Earthquakes per Year})$$

- 2 We also can model the probability distribution for the magnitude of an earthquake. For example, the data suggest that the probability of an earthquake in California of magnitude 8.0 or greater is $1/2493$, or approximately 0.0004. If we

use an exponential distribution to model the distribution of magnitudes, assuming that 4.0 is the least possible, then we might use the following model. Let M denote the magnitude, and let $M' = M - 4$. Then, using the exponential formula, we have $P(M \geq 8) = P(M' \geq 4) = e^{-4m} = 0.0004$. Now we can solve for m :

$$\begin{aligned} e^{-4m} &= 0.0004 \\ \ln(e^{-4m}) &= \ln(0.0004) \\ -4m &= -7.824 \\ m &= 1.96 \end{aligned}$$

Thus, the density function for M is given by

$$f(M) = 1.96e^{-1.96(M-4)}$$

Plot this density function.

We now can find the probability that any given earthquake will have a magnitude within a specified range on the Richter scale. For example, use this model to find

$$\begin{aligned} P_E(M \leq 6.0 \mid m = 1.96) & P_E(5.0 \leq \\ M \leq 7.5 \mid m = 1.96) & P_E(M \geq 6.4 \\ \mid m = 1.96) \end{aligned}$$

You may find it instructive to use this distribution to calculate the probability that an earthquake's magnitude falls within the five ranges of magnitude shown in Table 9.1. Here is a sensitivity-analysis issue: How might you find other reasonable values for m ? What about a range of possible values for m ?

3 We now have all of the pieces in the puzzle to find the probability of at least one severe (8.0 magnitude or more) earthquake occurring in California in the near future, say, within the next six months. Our approach will be to find the probability of the complement

$$P(X_{8+} \geq 1) = 1 - P(X_{8+} = 0)$$

where X_{8+} is used to denote the number of earthquakes having magnitude 8.0 or greater. Now expand $P(X_{8+} = 0)$ using total probability:

$$\begin{aligned} P(X_{8+} = 0) &= P(X_{8+} = 0 \mid X = 0) P(X = 0) + P(X_{8+} = 0 \mid X = 1) P(X = 1) \\ &+ \dots + P(X_{8+} = 0 \mid X = k) P(X = k) + \dots \\ &= \sum_{k=0}^{\infty} P(X_{8+} = 0 \mid X = k) P(X = k) \end{aligned}$$

The probabilities $P(X = k)$ are just the Poisson probabilities from Question 1: $P(X = k) = P_p(X = k \mid m = 12.5)$

where $m = 12.5$ because we are interested in a 6-month period. The probability of no earthquakes of magnitude 8.0 out of the k that occur is easy to find. If $k = 0$, then $P(X_{8+} = 0 \mid X = 0) = 1$. If $k = 1$, then

$$\begin{aligned} P(X_{8+} = 0 \mid X = 1) &= P_E(M < 8.0 \mid m = 1.96) \\ &= 1 - e^{-1.96(8-4)} \\ &= 0.9996 \end{aligned}$$

Likewise, if $k = 2$, then

$$P(X_{8+} = 0 \mid X = 2) = (0.9996)^2 = 0.9992$$

because this is just the probability of two independent earthquakes each having magnitude less than 8.0. Generalizing,

$$P(X_{8+} = 0 \mid X = k) = (0.9996)^k$$

Now we can substitute these probabilities into the formula: $P(X_{8+} \geq$

$$0) = 1 - P(X_{8+} = 0)$$

$$\begin{aligned} &= 1 - \sum_{k=0}^{\infty} P(X_{8+} = 0 \mid X = k) P(X = k) \\ &= 1 - \sum_{k=0}^{\infty} (0.9996)^k P_P(X = k \mid m = 12.5) \end{aligned}$$

To calculate this, you must calculate with k until the Poisson probability is so small that the remaining probabilities do not matter. It turns out that the probability of at least one earthquake of magnitude 8.0 or more within six months is approximately 0.005.

Now that you have seen how to do this, try calculating the probability of at least one earthquake of magnitude 8.0 or more (i) within the next year and (ii) within the next five years. For these, you may want to use a computer program to calculate the Poisson probabilities. How does the probability of at least one severe earthquake vary as you use the different reasonable values for m from your exponential model in Question 2? [Hint: Calculating the Poisson probabilities may be difficult, even on an electronic spreadsheet, because of the large exponential and factorial terms. An easy way to calculate these probabilities is to use the recursive equation

$$P_P(X = k + 1 \mid m) = \frac{m}{k + 1} P_P(X = k \mid m)$$

[This equation can be used easily in an electronic spreadsheet or calculator without having to calculate large factorial and exponential terms.]

Using the probability model described above, it turns out that the probability of at least one earthquake of magnitude 8.0 or more within the next 20 years in California is approximately 0.2 (or higher, depending on the value used for m in the exponential distribution for the magnitude, M). That is a 1-in-5 chance. Now imagine that you are a policy maker in California's state government charged with making recommendations regarding earthquake preparedness. How would this analysis affect your recommendations? What kinds of issues do you think should be considered? What about the need for more research regarding precise earthquake prediction at a specific location? What about regulations regarding building design and construction? What other issues are important?

Table 9.2
Probabilities for major earthquakes in California from two different probability models.

Time	USCS Probability	Poisson Model Probability
Next 5 years	0.27	0.29
Next 10 years	0.49	0.50
Next 20 years	0.71	0.75
Next 30 years	0.90	0.87

Source for USGS probabilities: U.S. Geological Survey (1988). "Probabilities of Large Earthquakes Occurring in California on the San Andreas Fault," by the Working Group on California Earthquake Probabilities. USGS Open-File Report No. 88-398, Menlo Park, CA.

The probabilistic model that we have developed using the information from Gere and Shah is based on a very simplistic model and does not account for geologic processes. Geologists do, however, use probability models in some cases as a basis for earthquake predictions. For example, as mentioned at the beginning of Chapter 8, a recent U.S. Geological Survey report concluded that the probability of an earthquake of 7.5-8.0 magnitude along the southern portion of the San Andreas Fault within the next 30 years is approximately 60%. The authors of the report actually constructed separate probability models for the occurrence of large quakes in different segments of major faults using data from the individual segments. Rather than a Poisson model, they used a lognormal distribution to model the uncertainty about the time between large earthquakes. Although their approach permits them to make probability statements regarding specific areas, their results can be aggregated to give probabilities for at least one major earthquake in the San Francisco Bay Area, along the southern San Andreas Fault in Southern California, or along the San Jacinto Fault. Table 9.2 compares their probabilities, which were developed for "large" earthquakes with expected magnitudes of 6.5-8.0, with our Poisson model probabilities of at least one earthquake having magnitude of 7.0 or greater. It is comforting to know that our model, even with its imperfections, provides probabilities that are not radically different from the geologists' estimates.

MUNICIPAL SOLID WASTE

Linda Butner considered her task. As the risk analysis expert on the city's Incineration Task Force (ITF), she was charged with reporting back to the ITF and to the city regarding the risks posed by constructing an incinerator for disposal of the city's solid waste. It was not a question of whether such an incinerator would be constructed. The city landfill site would be full within three years, and no alternative sites were available at a reasonable cost.

In particular, the state Department of Environmental Quality (DEQ) required information regarding levels of pollutants the incinerator was expected to produce.

DEQ was concerned about organic compounds, metals, and acid gases. It was assumed that the plant would incorporate appropriate technology and that good combustion practices would be followed. Residual emissions were expected, however, and the officials were interested in obtaining close estimates of these. Linda's task was to provide any analysis of anticipated emissions of dioxins and furans (organic compounds), particulate matter (PM, representing metals), and sulfur dioxide (SO₂, representing the acid gases). She figured that a thorough analysis of these substances would enable her to answer questions about others.

The current specifications called for a plant capable of burning approximately 250 tons of waste per day. This placed it at the borderline between small- and medium-sized plants according to the Environmental Protection Agency's (EPA) guidelines. In part, this size was chosen because the EPA had proposed slightly different permission levels for these two plant sizes, and the city would be able to choose the plant size that was most advantageous. A smaller (less than 250 tons/day) plant would be expected to have an electrostatic precipitator for reducing particulate matter but would not have a specified SO₂ emission level. A larger plant would have a fabric filter instead of an electrostatic precipitator and would also use dry sorbent injection — the injection of chemicals into the flue — to control the SO₂ level. A summary of the EPA's proposed emission levels is shown in Table 9.3.

Standard practice in environmental risk analysis called for assessment and analysis of "worst case" scenarios. But to Linda's way of thinking, this kind of approach did not adequately portray the uncertainty that might exist. Incineration of municipal solid waste (MSW) was particularly delicate in this regard, because the levels of various pollutants could vary dramatically with the content of the waste being burned. Moreover, different burning conditions within the incineration chamber (more or less oxygen, presence of other gasses, different temperatures, and so on) could radically affect the emissions. To capture the variety of possible emission levels for the pollutants, Linda decided to represent the uncertainty about a pollutant-emission level with a probability distribution.

The lognormal distribution makes sense as a distribution for pollutant-emission levels. (See Problem 9.36 for an introduction to the lognormal.) After consulting the available data for the content of pollutants in MSW and the pollutant-emission levels for other incinerators, Linda constructed a table (Table 9.4) to show the parameters

Table 9.3
Proposed pollutant
emission levels.


Pollutant	Plant Capacity (Tons of Waste per Day)	
	Small (Less than 250)	Medium (250 or More)
Dioxms/furam (ng/Nm ³)	500	125
PM (mg/dscm)	69	69
SO ₂ (ppmdv)	—	30

Notes: ng/Nm³ = nanograms per normal cubic meter; mg/dscm = milligrams per dry standard cubic meter; ppmdv = parts per million, by dry volume.

Table 9.4
Lognormal distribution
parameters μ and σ
for pollutants.

Pollutant	μ	σ
Dioxins/furans	3.13	1.20
PM	3.43	0.44
SO ₂	3.20	0.39

Figure 9.12
*Lognormal density
function for SO₂ emis-
sions from incineration
plant.*



for the lognormal distributions for the three pollutants in question. Figure 9.12 illustrates the lognormal distribution for the SO₂ emissions.

As Linda looked at this information, she realized that she could make certain basic calculations. For example, it would be relatively straightforward to calculate the probability that the plant's emissions would exceed the proposed levels in Table 9.3. Having these figures in hand, she felt she would be able to make a useful presentation to the task force.

Questions

- 1 The plant will be required to meet established emission levels for dioxins/furans and PM on an annual basis. Find the probabilities for exceeding the small-plant levels specified in Table 9.3 for these two pollutants. Repeat the calculations for the medium-plant emission levels.
- 2 If the plant is subject to SO₂ certification, its emissions of this pollutant will be monitored on a continual basis and the average daily emission level must remain below the specified level in Table 9.3. The numbers in Table 9.4, however, refer to the probability distribution of a single observation. That is, we could use the specified lognormal distribution to find the probability that a single reading of the SO₂

level exceeds the specified level. Finding the probability that an average daily emission exceeds the specified level, though, requires more analysis.

Let us assume that the average emission level will be calculated by taking n observations and then calculating the *geometric mean*. To do this, we multiply the n observations together and then take the n th root of the product. In symbols, let G denote the geometric mean:

$$G = \left(\prod_{i=1}^n X_i \right)^{1/n}$$

It turns out that if each X_i is drawn independently from a distribution which is lognormal with parameters μ and σ , then G has a lognormal distribution with parameters μ and σ/\sqrt{n} . In our case we will take the 24th root of a product of 24 hourly observations of emission levels.

Find the probability that the geometric mean of the 24 hourly emission observations exceeds the SO₂ limit specified in Table 9.3. Compare this to the probability that a single observation will exceed the same limit.

- 3 Discuss the issues that the city should consider in deciding whether to build a small- or medium-sized plant.

Sources: J. Marcus, and R. Mills (1988) "Emissions from Mass Burn Resource Recovery Facilities," *Risk Analysis*, No. 8, 315-327; and (1989) "Emission Guidelines: Municipal Waste Combustors," *Federal Register*, 54 (243), 52209.

REFERENCES

In this chapter we have only scratched the surface of theoretical probability distributions, although we have discussed most of the truly useful probability models. Theoretical distributions are widely used in operations research and in the construction of formal models of dynamic and uncertain systems. For additional study as well as many illustrative examples, consult the texts by DeGroot (1970), Feller (1968), Olkin, Gleser, and Derman (1980), or Winkler (1972). Johnson and Kotz (1969, 1970a, 1970b, and 1972) have compiled encyclopedic information on a great variety of probability distributions.

DeGroot, M. (1970) *Optimal Statistical Decisions*. New York: McGraw-Hill.

Feller, W. (1968) *An Introduction to Probability Theory and Its Applications, Vol. i*, 3rd ed. New York: Wiley.

Johnson, N. L., and S. Kotz (1969) *Distributions in Statistics: Discrete Distributions*. New York: Houghton Mifflin.

Johnson, N. L., and S. Kotz (1970a) *Distributions in Statistics: Continuous Univariate Distributions I*. Boston: Houghton Mifflin.

Johnson, N. L., and S. Kotz (1970b) *Distributions in Statistics: Continuous Univariate Distributions II*. Boston: Houghton Mifflin.

Johnson, N. L., and S. Kotz (1972) *Distributions in Statistics: Continuous Multivariate Distributions*. New York: Wiley.

Olkin, L., L. J. Gleser, and C. Derman (1980) *Probability Models and Applications*. New York: Macmillan.

Winkler, R. L. (1972) *Introduction to Bayesian Inference and Decision*. New York: Holt.

EPILOGUE

The case study on earthquake prediction was written in early October 1989, just two weeks before an earthquake of magnitude 7.1 occurred near Santa Cruz, California, on the San Andreas Fault. Strong ground shaking lasted for approximately 15 seconds. The results were 67 deaths, collapsed buildings around the Bay Area, damage to the Bay Bridge between San Francisco and Oakland, and the destruction of a one-mile stretch of freeway in Oakland. This was a small earthquake, however, relative to the 8.3 magnitude quake in 1906. The "big one" is still to come, and may cause even more damage.

CHAPTER 10. Using Data

We have discussed subjective judgments and theoretical distributions as sources for probabilities when modeling uncertainty in a decision problem. In this chapter we consider an obvious source for information about probabilities; historical data. It is possible to use data alone to develop probability distributions; we cover the development of discrete and continuous distributions in the first part of the chapter. We also consider the use of data to understand and model relationships among variables. It is possible to use data in conjunction with theoretical probability models; that is, the data can provide information to help refine assessments regarding the parameters of a theoretical distribution. In the last part of the chapter, we discuss the use of data in conjunction with some of the theoretical models discussed in Chapter 9.

Using Data to Construct Probability Distributions

Using past data when it is available is a straightforward idea, and most likely you have done something like this in the past, at least informally. Suppose, for example, that you are interested in planning a picnic at the Portland Zoo on an as-yet-undetermined day during February. Obviously, the weather is a concern in this case, and you want to assess the probability of rain. If you were to ask the National Weather Service for advice in this regard, forecasters would report that the probability of rain on any given day in February is approximately 0.47. They base this estimate on analysis of weather during past years; on 47% of the days in February over the past several years rain has fallen in Portland.

We can think about developing both discrete and continuous probability distributions on the basis of empirical data. For the discrete situation, the problem really becomes one of creating a relative frequency histogram from the data. In the case of a continuous situation, we can use the data to draw an empirically based CDF. We will look briefly at each of these.

Histograms

Imagine that you are in charge of a manufacturing plant, and you are trying to develop a maintenance policy for your machines. An integral part of the analysis leading to policy selection most likely would involve an examination of the frequency of machine failures. For example, you might collect the following data over 260 days:

- No Failures 217 days
- One Failure 32 days
- Two Failures 11 days

These data lead to the following relative frequencies, which could be used as estimates of probabilities of machine failures in your analysis:

No Failures $0.835 = 217/260$

One Failure $0.123 = 32/260$

Two Failures $0.042 = 11/260$

Thus, we would have a histogram that looks like Figure 10.1, and in a decision tree, we would have a chance node with three branches like that in Figure 10.2.

Our treatment of histograms and estimation of discrete probabilities is brief precisely because the task is simply a matter of common sense. The only serious consideration to keep in mind is that you should have enough data to make a reliable estimate of the probabilities. (One nuclear power plant accident is not enough to

Figure 10.1
Relative frequency histogram for machine failure data.





Figure 10.2
The decision-tree representation of uncertainty regarding machine failures.



develop a probability distribution!) The data requirements depend on the particular problem, but the minimum should be approximately five observations in the least likely category. The other categories, of course, will have more observations.

Finally, keep in mind that your probability estimates are just that — estimates. The goal is to model uncertainty. Always ask yourself whether the probabilities estimated on the basis of the data truly reflect the uncertainty that you face. If you are not satisfied with the representation based on the data, you may need to model your uncertainty using subjective assessment methods. In particular, this might be the case if you think that the past may not indicate what the future holds.

Empirical CDFs

Continuous probability distributions are common in decision-making situations. We have seen how to assess a continuous probability distribution subjectively through assessment of a CDF. Our use of data to estimate a continuous distribution will follow the same basic strategy. To illustrate the principles, we will use an example and data concerning the costs related to correctional halfway houses.

HALFWAY HOUSES

A halfway house is a facility that provides a residence and some supervision and support for individuals who have been recently released from prison. The purpose of these halfway houses is to ease the transition from prison life to normal civilian life, with the ultimate goal being to improve an ex-convict's chance of successful integration into society. The National Advisory Committee on Criminal Justice Standards and Goals is responsible for providing information pertaining to the costs and resource implications of correctional standards that are related to halfway houses.

The main purpose of the reports made by the advisory committee is to provide cost information to state and local decision makers regarding the many services that halfway houses perform. One important variable is yearly per-bed rental costs. Denote this

Table 10.1
Yearly bed-rental costs
for 35 halfway houses.

Rental Costs (\$)				
52	205	303	400	64
16	250	313	402	69
100	257	317	408	73
136	264	325	417	74
137	280	345	422	75
186	282	373	472	79
196	283	384	480	89


Source: T. Sincich, (1989) *Business Statistics by Example*, 3rd ed. San Francisco: Dellen.

variable by C Table 10.1 shows yearly per-bed rental costs (in dollars) — values of C — for a random sample of 35 halfway houses. The data are arranged in ascending order.

To create a smooth CDF from these data, recall the idea of a cumulative probability. Suppose we look at the middle value, 325. Eighteen of the 35 values are less than or equal to 325. It also is true that 18 values are less than or equal to 326, 327, ..., up to 344.99. So $P(C \leq 325) = P(C \leq 326) = \dots = P(C \leq 344.99) = 18/35$, or 0.514. So how can we estimate the 0.514 fractile? Take 335 as the best estimate of the 0.514 fractile; 335 is the value that is halfway between 325 and 345. Figure 10.3 shows what we are doing in terms of a CDF graph. The data points define the ends of the flat steps on the graph. Our estimate, the smooth curve, will go through the centers of the flat steps.

To get a CDF, we do the same thing for all data points. The procedure first rank orders the data and then calculates the centers of the flats; this is just a matter of calculating the halfway points between adjacent data points. For example, consider the

Figure 10.3
Building the CDF by drawing a smooth curve through the centers of the flat steps.



halfway point between the first and second data points. The halfway point is $(52 + 76)/2 = 64$. Call this x_1 . Now calculate x_2, \dots, x_{n-1} , where n is the number of data points. Now associate with each x_m its approximate cumulative probability. The value x_m is in the m th position, and so it has cumulative probability estimated as m/n . In symbols, $P(X \leq x_m)$ is approximately m/n . In our example, $P(C \leq 64)$ is approximately $1/35$, or 0.029 . Likewise, $x_{15} = 308$, so $P(C \leq 308)$ is approximately $15/35$, or 0.429 . Table 10.2 shows the calculations for all 35 data points. The final step is to plot these points as in Figure 10.4. The tails of the CDF are sketched as smooth extrapolations of the curve.


Although all of the calculations that must be done to obtain a CDF look complicated, these are the kinds of calculations that can be done easily on an electronic spreadsheet and then plotted on a graph.

Once we have the CDF, we can use it in the same way as before. For example, we could use it to make probability statements about the uncertain quantity. With the halfway house, for example, we could say that there is a 50% chance that the yearly bed-rental cost will fall between \$240 and \$450. Likewise, we could say that there is

Table 10.2
Estimated cumulative probabilities for the halfway-house data.

Obs. No.	Cost	x_m	Cumulative Probability	Obs. No.	Cost	x_m	Cumulative Probability
1	52	64.0	0.029	19	345	359.0	0.543
2	76	88.0	0.057	20	373	378.5	0.571
3	100	118.0	0.086	21	384	392.0	0.600
4	136	136.5	0.114	22	400	401.0	0.629
5	137	161.5	0.143	23	402	405.0	0.657
6	186	191.0	0.171	24	408	412.5	0.686
7	196	200.5	0.200	25	417	419.5	0.714
8	205	227.5	0.229	26	422	447.0	0.743
9	250	253.5	0.257	27	472	476.0	0.771
10	257	260.5	0.286	28	480	561.5	0.800
11	264	272.0	0.314	29	643	668.0	0.829
12	280	281.0	0.343	30	693	712.5	0.857
13	282	282.5	0.371	31	732	740.5	0.886
14	283	293.0	0.400	32	749	749.5	0.914
15	303	308.0	0.429	33	750	770.5	0.943
16	313	315.0	0.457	34	791	841.0	0.971
17	317	321.0	0.486	35	891		
18	325	335.0	0.514				

Figure 10.4
*Estimated CDF for the
 halfway-house data.*




a 25% chance that the cost would fall below \$240 and a 25% chance that it would fall above \$450.

Another way we could use the empirical CDF would be to derive a discrete approximation (Chapter 8) for inclusion in a decision tree. In the case of the halfway-house data, Figure 10.5 shows a three-point discrete approximation that has been constructed using the extended Pearson-Tukey method described in Chapter 8. Recall that the 0.05 and 0.95 fractiles each are given probability 0.185 and the median is given probability 0.63. These fractiles are indicated in Figure 10.4.

There are other alternatives for using data to approximate a continuous distribution. One way that may be familiar to you is to split the data into reasonable groups and create a relative frequency histogram of the grouped data. This procedure is straightforward and similar to the discussion above regarding discrete probabilities. Again, no category should contain fewer than five observations, and typically the widths of the intervals that define the categories should be the same. If you need more direction on the use of histograms for continuous distributions, consult any introductory statistics textbook. The text by Sincich, from which the halfway-house data were taken, is an excellent source.

Figure 10.5
*Three-point approxi-
 mation for the
 halfway-house data.*



Using Data to Fit Theoretical Probability Models

Another way to deal with data is simply to fit a theoretical distribution to it. Typically, this involves two steps. First, we must decide what kind of distribution is appropriate (binomial, Poisson, normal, and so on). The choice should be based both on an understanding of the situation (for example, whether the uncertain quantity must be between 0 and 1, which suggests a beta distribution, or whether a discrete distribution is appropriate) as well as an inspection of the distribution of the data. For example, it makes little sense to fit a normal distribution, which is symmetric, to data whose distribution is highly skewed.

Having chosen the kind of distribution, the next step is to choose the values of the distribution parameters. In the case of the binomial, for instance, we need values for n and p . For the normal, μ and σ are required. Parameter values can be chosen by calculating some summary statistics (mean, standard deviation, and so on) for a sample and then simply using those as the parameter values. For example, it is possible to calculate the sample mean (\bar{x}) and sample variance (s^2) for the 35 halfway-house observations:

$$n = 35$$

$$\bar{x} = \frac{\sum_{i=1}^n x_i}{n} = 380.4$$

$$s^2 = \frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n - 1} = 47,344.3$$

$$s = \sqrt{47,344.3} = 217.6$$

Thus, we might choose a normal distribution with mean $\mu = 380.4$ and standard deviation $\sigma = 217.6$ to represent the distribution of the yearly bed-rental costs.

Another possibility is to fit a theoretical distribution using fractiles. That is, find a theoretical distribution whose fractiles match as well as possible with the fractiles of the empirical data. This is the data-based counterpart of the procedure discussed in Chapter 9 for fitting a theoretical distribution to a subjectively assessed distribution. In this case we would be fitting a theoretical distribution to a data-based distribution.

For most initial attempts to model uncertainty in a decision analysis, it may be adequate to use the sample mean and variance as estimates of the mean and variance of the theoretical distribution and to establish parameter values in this way. Refinement of the probability model may require more careful judgment about the kind of distribution as well as more care in fitting the parameters. Statisticians have devised many clever parameter-estimation methods. A discussion of these techniques is beyond the scope of this treatment but may be found in advanced textbooks in statistics.

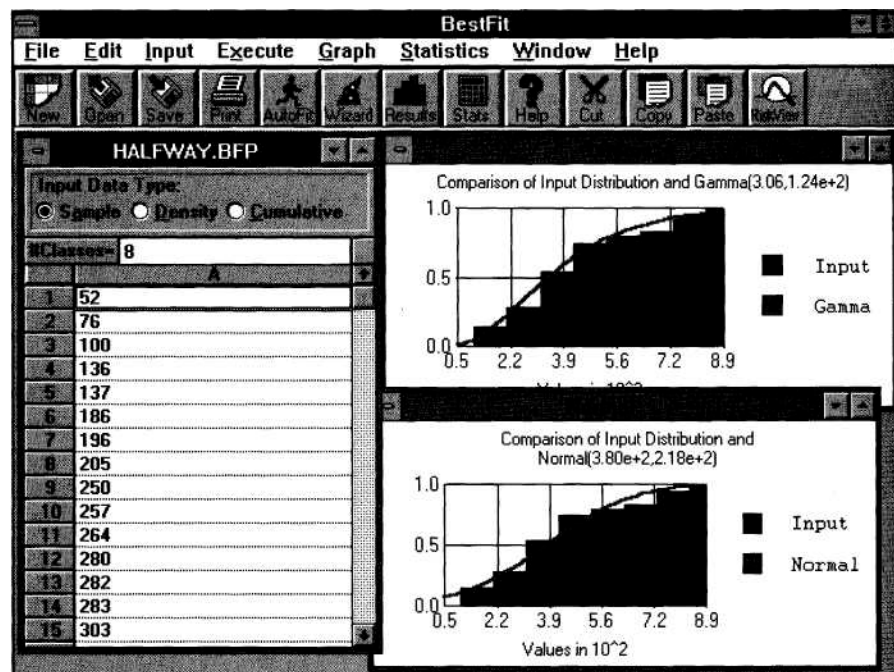
Software for Fitting Distributions: BestFit

Excellent software exists to help a decision analyst fit theoretical distributions to data. BestFit, a program published by Palisade Software, provides an easy-to-use personal-computer environment for fitting distributions. This program allows the user to estimate parameters for many different families of distributions, including all of the ones discussed in Chapter 9. (In addition, BestFit can handle cumulative probability information, thereby providing a way for an analyst to fit theoretical distributions to subjective assessments elicited from experts as in Chapter 8.) For the power user, the program displays elaborate statistics, and the analyst can control many aspects of the analytical process. For the less sophisticated, BestFit can guide the process or even go through a fully automatic distribution-fitting procedure.

Figure 10.6 shows a computer screen from BestFit, which runs in a Windows environment. The halfway-house data are listed in the spreadsheet on the left side of the screen, and the two graphs to the right show the results of fitting two kinds of distributions to these data. For example, the normal distribution has a mean of \$380 and standard deviation \$218; this represents the "best fitting" normal distribution for these data. The other distribution is from the gamma family, which is closely related to the exponential and Poisson distributions. Although neither of these models is a perfect fit, it turns out that the gamma distribution fits somewhat better than the normal.

A companion program called RiskView allows the user to display and analyze the chosen distribution. For example, the analyst could use the program to extract the appropriate fractiles from the gamma distribution in Figure 10.6 in order to create a

Figure 10.6
Fitting distributions to the halfway-house data with BestFit.



discrete approximation or to derive fractiles for making probability statements about the uncertain quantity.

Using Data to Model Relationships

One of the most important things that we do with data is to try to understand the relationships that exist among different phenomena in our world. The following list includes just a few examples:

Causes of Cancer. Our smoking habits, diet, exercise, stress level, and many other aspects of our lives have an impact on our overall health as well as the risk of having cancer. Scientists collect data from elaborate experiments to understand the relationships among these variables.

Sales Revenue. For a business, perhaps the most crucial issue is understanding how various economic conditions, including its own decisions, can impact demand for products and hence revenue. Firms try to understand the relationships among sales revenue, prices, advertising, and other variables.

Economic Conditions. Economists develop statistical models to study the complex relationships among macroeconomic variables like disposable income, gross domestic product, unemployment, and inflation.

Natural Processes. Much scientific work is aimed at understanding the relationships among variables in the real world. A few examples include understanding the relationships among weather phenomena, movements in the earth's crust, changes in animal and plant populations due to ecological changes, and causes of violence in society.

In most cases, the motivation for studying relationships among phenomena that we observe is to gain some degree of control over our world. In many cases we hope to make changes in those areas where we have direct control in order to accomplish a change in another area. For example, on occasion the U.S. Federal Reserve Board buys or sells securities in order to have an impact on interest rates in the U.S. money market. Doing so requires an understanding of how their operations in the market can affect interest rates. Firms want to set prices and advertising budgets so that they result in an optimal level of profits, but doing so requires an understanding of what affects demand for their products and services. Scientific studies of natural processes often are used as inputs in government policy and regulation decisions.

In this section, we will focus on the problem of using data on a number of auxiliary variables (which we will denote as X_1, \dots, X_k) to determine the distribution of some other variable of interest (Y) that is related to the X 's. Y is sometimes called a *response variable*, because its probability distribution changes in response to changes in the X 's. Likewise, the X 's sometimes are called *explanatory variables*, because they can be used to help explain the changes in Y . (Y and the X 's are also sometimes called *dependent* and *independent* variables, although this terminology can be misleading. Y does not necessarily *depend* on the X 's in a causal sense, and the X 's can be highly dependent — in a probabilistic sense — among themselves.)

In a business context, a common example is forecasting sales, and to come up with a conditional distribution for sales (Y), we might want to use explanatory variables such as price (X_1), advertising (X_2), a competitor's price (X_3), or other important micro-economic variables. These variables could be either quantities that are uncertain (like the competitor's price) or variables that are under the control of the decision maker (e.g., price and advertising). Often it is possible to obtain data indicating the amount of sales that was associated with a particular price, amount of advertising, and so on. How to use such data to create a model of the relationships among the variables is our topic.

The use of data to understand relationships is not trivial. Consider the influence diagram in Figure 10.7. The brute force approach would require obtaining enough data to estimate the conditional distribution for the particular variable of interest (Y) for every possible combination of values for its conditioning or predecessor variables (X_1 and X_2)! In addition, of course, we would need to know what are feasible values for the decision variable (X_1), and we would have to assess a distribution for the possible values for the uncertain variable (X_2). That would require a lot of data, and even in simple problems this could be a tedious or infeasible task.

One possible simplification, of course, is to use approximations. For example, in Figure 10.7 we could reduce the continuous distribution for X_2 to an extended Pearson-Tukey three-point approximation, and we could consider three possible values for X_1 (low, medium, and high, say). Even in this case, we would need to assess nine different conditional probability distributions for Y based on the possible scenarios. An approach like this may be possible in a simple situation, but the data requirements again become unwieldy as soon as there are more than just a few conditioning arcs. For example, suppose we were to add two more variables, X_3 and X_4 , as in Figure 10.8. Using three-point approximations for the uncertain variables (X_2 and X_3) and considering only three possible values for each of the two decision variables (X_1 and X_4), we would still need to come up with 81 conditional distributions.

Figure 10.7
An influence diagram
for modeling
relationships among
uncertain quantities
 X_1 , X_2 , and Y .





Figure 10.8
An influence
diagram relating two
uncertain quantities
and two decision
variables to Y .



(There are 81 conditional distributions because there are $3^4 = 81$ different combinations of the values for the X 's.) Clearly, we need a way to streamline this process!

The Regression Approach

One way to approach this problem is to split it up into two pieces. First is determining the conditional expected value of Y given the X 's, $E(Y | X_1, \dots, X_k)$. The second step is to consider the conditional probability distribution around that expected value. Statisticians refer to this general modeling approach as *regression*.

We focus here on the simplest regression model, whose assumptions naturally lead it to be called *linear regression*. Beginning with the expected value:

1 *The conditional expected value of Y is linear in the X 's.* In symbols,

$$E(Y | X_1, \dots, X_k) = \beta_0 + \beta_1 X_1 + \dots + \beta_k X_k$$

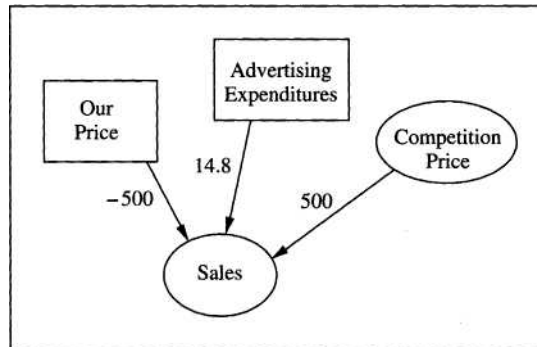
The β 's are coefficients, and they serve the purpose of combining the X values to obtain a conditional expected value for Y . Moreover, every unit increase in X_i leads to a β_i change (positive or negative, as indicated by the sign of β_i) in the conditional expected value of Y . For example, suppose we have an equation relating regional expected sales (in \$ 1000s) of a high-volume consumer electronics item (say, a portable stereo) to the amount of money spent on advertising in that region (in \$ 1000s), the price set forth for that item in the ads (\$), and the current price of a competitive product (\$):

$$\begin{aligned} E(\text{Sales} | \text{Advertising, Price, Competition Price}) \\ = 2000 + 14.8(\text{Advertising}) - 500(\text{Price}) + 500(\text{Competition Price}) \end{aligned}$$

An influence diagram portraying this relationship is shown in Figure 10.9. Note that we have placed the coefficients on the arrows, indicating the effect that each variable has on the expected value of Sales. For example, every additional \$1000 spent on Advertising leads to an increase in expected Sales of \$14,800 (14.8 times \$1000, the units for Sales), given that Price and Competition Price stay the same. The opposite signs on the coefficients for Price and Competition Price indicate that these variables are expected to have opposite effects; an increase in our price should decrease our expected sales (everything else being equal), as indicated by the negative coefficient (-500). But an increase in Competition Price (again, with everything else held constant) is good news for us and should increase our expected sales, as indicated by the positive coefficient ($+500$) on Competition Price. Moreover, the price effects are quite strong, reflecting the competitive nature of the market; an increase in Price of \$1 leads to a \$500,000 decrease in expected Sales. Likewise, a \$1 increase in the Competition Price leads to an increase in expected Sales of \$500,000.

It is important to remember that the equation defines a relationship between the explanatory variables and the *expected* Y (Sales). The *actual* Y value will be above or below this expected value to some extent; this is where the uncertainty and the con-

Figure 10.9
Relating sales to three
explanatory variables.



ditional probability distribution of Y come into play. The regression approach makes the following assumption about this distribution:

- 2 *The distribution around the conditional expected value has the same shape regardless of the particular X values.*

A convenient way to think about this is in terms of random errors (or "noise") added to the conditional expected value. Denote a random error as ε , and let us say that these errors are, on average, equal to zero: $E(\varepsilon) = 0$. Now we can write an expression for an actual Y value as follows:

$$Y = E(Y | X_1, \dots, X_k) + \varepsilon$$

This implies that the conditional distribution (and the corresponding density) of Y , given the X 's, has the same shape as the distribution (or density) of the errors, but it is just shifted so that the distribution is centered on the expected value $E(Y | X_1, \dots, X_k)$.

Take the Sales example. Suppose that the errors — the unexplainable factors that make Sales fall above or below its conditional expected value — have a CDF as shown in Figure 10.10. Now, suppose we decide to spend \$40,000 on Advertising (X_1) and to set Price (X_2) at \$97.95, and Competition Price (X_3) turns out to be \$94.99. Then expected Sales (Y), conditional on these values, would be

$$\begin{aligned} E(Y | X_1, X_2, X_3) &= 2000 + 14.8(40) - 500(97.95) + 500(94.99) = \\ &= 1112(\$1000\text{s}) \end{aligned}$$

Suppose in an entirely different scenario we had Advertising equal to \$70,000, Price equal to \$93.95, and Competition Price equal to \$98.99. Then conditional expected Sales would be

$$\begin{aligned} E(Y | X_1, X_2, X_3) &= 2000 + 14.8(70) - 500(93.95) + 500(98.99) \\ &= 5556(\$1000\text{s}) \end{aligned}$$

The conditional CDFs for Sales, given the three conditioning variables in each scenario and the distribution for errors, are displayed in Figure 10.11. As you can see, the distributions have the same shape as the distribution for the errors shown in

Figure 10.10
A CDF for error in the sales example.




Figure 10.11
Conditional CDFs for sales in two scenarios.




Figure 10.10. In Figure 10.11, though, the location of the distributions are different, and this is due entirely to differences in the explanatory variables.

Our two assumptions — a linear expression for conditional expected value and the constant shape for the conditional distribution — take us quite a long way toward being able to use data to study relationships among variables. You should keep in mind, however, that this is only one way to model the relationships and that it is quite possible for the expected-value relationships to be nonlinear or for the distribution to change shape. An example of a nonlinear relationship would be the degree of acidity in a lake and the number of microorganisms. If the lake is highly acidic, then reducing the acidity can lead to an increase in microorganisms up to a point, beyond which further changes are detrimental to the microorganisms' population. Changes in the distribution can be somewhat more subtle. Consider a retailer who is interested in entering a new market segment. Although the retailer may be fairly sure of the amount of sales resulting from ads aimed at her traditional market, the same

advertising dollars spent in an effort to attract the new customers might lead to a great deal of uncertainty.

Although our approach has some limitations, it will provide an excellent base that can be extended via additional modeling techniques. We will mention some of these techniques later. For now, though, let us turn to the problem of using data within the regression framework we have established.

Estimation: The Basics

The discussion above has assumed that we know the β coefficients and the distribution of the errors. Of course, we may be able to make subjective judgments of these quantities, but if data are available, we may want to use that information in a systematic way to estimate the β 's and to construct a data-based distribution of errors. To understand the process, let us simplify the sales example that we have been using. Suppose that all we have are data for Sales and Advertising, and we want to study the relationship between these two variables alone. The data, observations of Advertising and Sales during 36 different promotions in the eastern United States, are shown in Table 10.3. Figure 10.12 shows a *scatterplot* of these data; each point on the graph represents one of the observations.

Using the linear-regression assumptions from the previous section, we establish the equation relating expected Sales (Y) and Advertising (X_1):

$$E(Y | X_1) = \beta_0 + \beta_1 X_1$$

Figure 10.12
A scatterplot of
advertising versus
sales.




Table 10.3
Data for sales
example.

Promotion	Advertising (\$1000s)	Sales (\$1000s)
1	366	10541
2	377	8891
3	387	5905
4	418	8251
5	434	11461
6	450	6924
7	457	7347
8	466	10972
9	467	7811
10	468	10559
11	468	9825
12	475	9130
13	479	5116
14	479	7830
15	481	8388
16	490	8588
17	494	6945
18	502	7697
19	505	9655
20	529	11516
21	532	11952
22	533	13547
23	542	9168
24	544	11942
25	547	9917
26	554	10666
27	556	9717
28	560	13457
29	561	10319
30	566	9731
31	566	10279
32	582	7202
33	609	12103
34	612	11482
35	617	11944
36	623	9188

We use our data to estimate β_0 and β_1 . Estimating β_0 and β_1 amounts to finding a line that passes through the cloud of points in the scatterplot; Figure 10.13 shows two different candidates with their corresponding expressions. The expressions show what those particular lines use as estimates of β_0 and β_1 . The dashed line, for example, estimates β_0 to be 1900 and β_1 to be 15. Clearly, no single line can pass precisely through all of the points at once, but we would like to find one that in some sense is the "best fitting" line. And, as you can see from the graph, there are many reasonable estimates for β_0 and β_1 .

How will we choose the best-fitting line through the data points? Although there are many possible answers to this question, we will choose a way that many statisticians have found convenient and useful. In simple terms, we will choose the line that minimizes the sum of the squared vertical distances between the line and each point. Figure 10.14 shows how this will work for two sample points, A and B. Between the line and any given point (x_i) there is a certain vertical distance, which we will call the *residual* (e_i). For any given line, we can calculate the residuals for all of the points. Then we square each residual and add up the squares. And the problem is to find the line that minimizes this total of squared residuals. (The decision to minimize squared residuals is somewhat arbitrary. We could just as well choose estimates that minimize the sum of the absolute values of residuals, for example. Any error measure that penalizes large errors above or below the line will work. Mathematicians have focused on the sum of squared residuals because this sum is easy to work with mathematically.)

Figure 10.13
Two possible lines
relating expected sales
and advertising.





Figure 10.14
Calculating residuals.



In symbols, let b_0 and b_1 represent the estimates of β_0 and β_1 . With this we can calculate the i th residual as

$$e_i = y_i - (b_0 + b_1 x_i)$$

Squaring and summing over the n residuals gives the total (which we denote by SSE for the sum of squared errors):

$$SSE = \sum_{i=1}^n [y_i - (b_0 + b_1 x_i)]^2$$

We want to choose b_0 and b_1 so that SSE is minimized. While this sounds like a complicated problem, it turns out to be a rather straightforward calculus problem. And even more good news is that you will never need to do the calculations yourself! Many computer programs are available that do the calculations for you automatically.

Figure 10.15 shows the spreadsheet in Microsoft Excel that results from doing the calculations using Excel's built-in "Regression" procedure in the Sales-Advertising example. The figure highlights Excel's estimates $b_0 = 3028.18$ and $b_1 = 12.95$. These are the coefficients that define the minimum-SSE line. (In fact, $SSE = 117,853,513$ and is shown in the column labeled "Sum of Squares.") Excel does many other calculations that are useful in drawing inferences from the data. For our purposes now, that of modeling relationships among variables, we will stick to

Figure 10.15
Using Excel to estimate regression coefficients.

Sales & Advertising						
	A	B	C	D	E	F
61						
62	Regression Statistics					
63						
64	Multiple R	0.43				
65	R Square	0.18				
66	Adjusted R Square	0.16				
67	Standard Error	1,861.79				
68	Observations	36				
69						
70	Analysis of Variance					
71		df	Sum of Squares	Mean Square	F	Significance F
72	Regression	1	26,180,794	26,180,794	7.55	0.01
73	Residual	34	117,853,513	3,466,280		
74	Total	35	144,034,307			
75						
76		Coefficients	Standard Error	t Statistic	P-value	Lower 95%
77						Upper 95%
78	Intercept	3,028.18	2,414.97	1.25	0.22	(1,879.64)
79	Advertising	12.95	4.71	2.75	0.01	3.37
80						22.53

the basics. All we require at this point are b_0 and b_1 , the estimates of β_0 and β_1 . At the end of the chapter we will briefly discuss the idea of statistical inference and how it relates to decision-analysis modeling.

We can write out the expression for the line defined by the estimates b_0 and b_1 :

$$E(Y | X_1) = 3028.18 + 12.95X_1$$

This equation makes sense. In general, we expect Sales to increase when we spend more on Advertising, as indicated by the positive coefficient 12.95. In particular, this coefficient can be interpreted as follows: For every additional \$1000 spent on advertising, Sales are expected to increase by \$12,950. (Recall that Advertising and Sales are measured in \$1000s.)

For our simple model, we now have an indication of how expected Sales change as Advertising changes. How can we use the data to construct a model of the distribution of errors? This step is straightforward; for every data point we have a residual, which can be thought of as an estimate of the error associated with that particular point. Thus, the distribution of the residuals should be a good approximation to the distribution of the errors. Table 10.4 shows the calculation of the residuals, and Figure 10.16a displays a cumulative distribution constructed from those residuals in the same way that we constructed the distribution of halfway-house bed costs earlier (Figure 10.4). Figure 10.16b is a histogram of the residuals.

The CDF can be used to make probability statements about the errors. For example, we can see that the quartiles are approximately -1300 and $+1300$, so there is approximately a 50% chance that the error falls between these values. Or we can use the CDF to construct a discrete approximation. The dashed lines in Figure 10.16a for

Table 10.4
Residuals for the sales
example.

Observation	Sales (Y)	Advertising (X)	$E(Y X) =$ 3028.18 +12.95X	Residual = Y - E(Y X)
1	10541	366	7768	2773
2	8891	377	7911	980
3	5905	387	8040	-2135
4	8251	418	8442	-191
5	11461	434	8649	2812
6	6924	450	8856	-1932
7	7347	457	8947	-1600
8	10972	466	9063	1909
9	7811	467	9076	-1265
10	10559	468	9089	1470
11	9825	468	9089	736
12	9130	475	9180	-50
13	5116	479	9232	-4116
14	7830	479	9232	-1402
15	8388	481	9258	-870
16	8588	490	9374	-786
17	6945	494	9426	-2481
18	7697	502	9530	-1833
19	9655	505	9568	87
20	11516	529	9879	1637
21	11952	532	9918	2034
22	13547	533	9931	3616
23	9168	542	10048	-880
24	11942	544	10074	1868
25	9917	547	10112	-195
26	10666	554	10203	463
27	9717	556	10229	-512
28	13457	560	10281	3176
29	10319	561	10294	25
30	9731	566	10358	-627
31	10279	566	10358	-79
32	7202	582	10566	-3364
33	12103	609	10915	1188
34	11482	612	10954	528
35	11944	617	11019	925
36	9188	623	11097	-1909

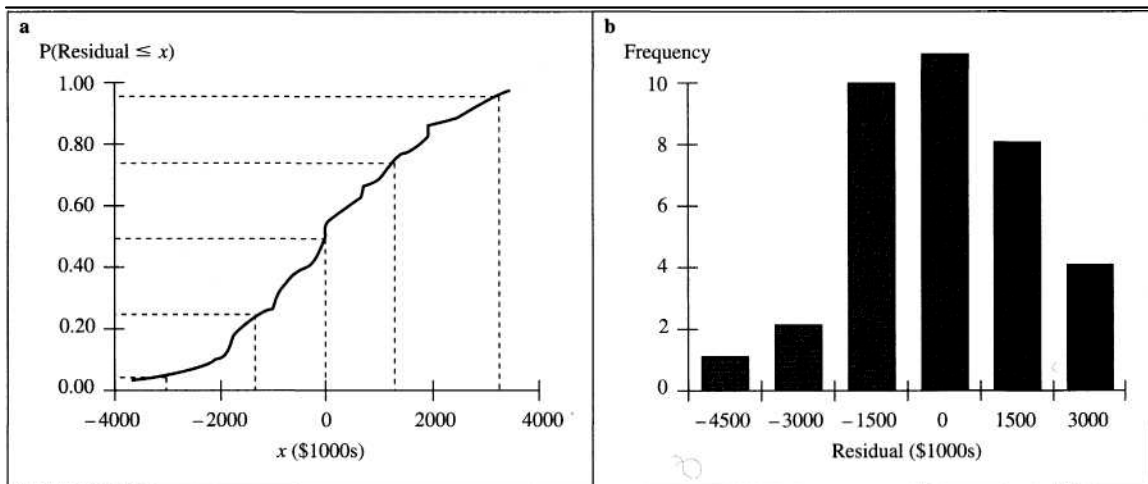



Figure 10.16 CDF and histogram of residuals.

Figure 10.17
Extended Pearson-Tukey approximation
for error distribution.



the 0.05, 0.50, and 0.95 fractiles represent the required numbers for a three-point extended Pearson-Tukey approximation as shown in Figure 10.17.

The CDF can also be used as a basis for representing the uncertainty in Sales given Advertising. To create a conditional distribution for Sales given Advertising, simply add the conditional expected Sales to the horizontal scale. For example, if Advertising = \$110,000, then $E(\text{Sales} \mid \text{Advertising}) = \$4,452,680$; add this amount to the horizontal scale to get the distribution for Sales given Advertising = \$110,000 as shown in Figure 10.18. This distribution can in turn be used to generate conditional probability statements about Sales. Or we could create a discrete approximation of the distribution for use in an influence diagram or decision tree. Finally, we could also fit a theoretical distribution. Typically in regression analysis, a normal distribution is fit to the residuals, but virtually any continuous distribution that makes sense for the situation could be used. (For example, the residuals could be used as data in BestFit to find the best-fitting distribution.)

All of our effort has paid off in modeling terms. With the coefficient estimates and the distribution of errors, we have a complete, if simplified, model of the uncertainty in Sales given its conditioning variable, Advertising. Figure 10.19 shows an

Figure 10.18
CDF for sales given advertising = \$110,000.


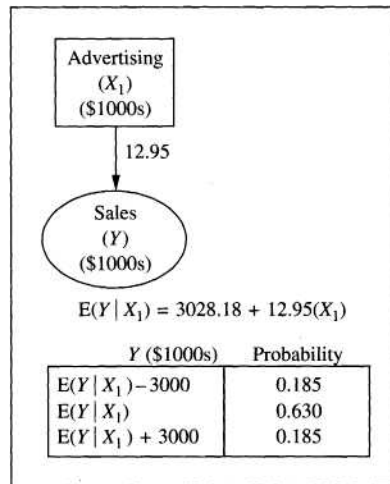


Figure 10.19
Influence diagram model showing relationship between sales and advertising.



influence diagram that represents the simple model we have created, including a three-point approximation for Sales given Advertising.

Estimation: More than One Conditioning Variable

Our advertising-and-sales example above demonstrated the basics. The model that we came up with, though, really is quite simplistic. In fact, with Advertising as the

only explanatory variable, we implicitly assume that nothing else matters. But that clearly is too simple, and we have already argued that both the price of our product and our competitor's price also matter.

An augmented data set is shown in Table 10.5. For each of the 36 occasions, we now also have information on Price (X_2) and Competition Price (X_3). Can we incorporate these data into our analysis? The answer is yes, and the procedure is essentially the same as it was before.

We want to come up with estimates for β_0 , β_1 , β_2 , and β_3 in the following expression:

$$E(Y | X_1, X_2, X_3) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3$$

As before, we will let b_0 , b_1 , b_2 , and b_3 denote the estimates of the corresponding β coefficients. Analogous to the one-variable version of the problem, we can calculate expected Sales, given the values of the conditioning variables. In turn, we can calculate the residual by subtracting the expected sales calculation from the actual sales for that observation. And, as before, we want to specify the b 's so that the sum of the squared residuals is minimized. Accomplishing this is just a slightly more general version of the one-variable problem described above.

The same computer program used for the one-variable problem can be used here. Figure 10.20 shows the Excel screen with the results of the regression analysis with all three conditioning variables. As before, b_0 , b_1 , b_2 , and b_3 are in the column labeled "Coefficients." With these values, we can now write out the formula for expected Sales, given the three conditioning variables:

$$E(Y | X_1, X_2, X_3) = 2199.34 + 15.05X_1 - 503.76X_2 + 499.67X_3$$

This expression makes sense. As before, every additional \$1000 spent on Advertising increases expected Sales, but now in a more complete model this effect is estimated to be \$15,050. The coefficients on Price and Competition Price are interpreted similarly; a \$1 Price increase would lead to a decrease in expected Sales of \$503,760, whereas a \$1 increase in the Competition Price would bring expected Sales up by \$499,670.

Also as before, we can calculate the residuals and use them to construct a CDF for the errors. Table 10.6 gives the residuals for our full model, and Figure 10.21 displays the CDF of the residuals. From the graph (especially the horizontal scale), you can see that the distribution of errors is much tighter in the full model than it was in the model above in which we used only Advertising. In fact, an estimate of the standard deviation (often called the *standard error* and usually abbreviated as s_e) of the current distribution is \$459,100, and is given in Cell B134 in Figure 10.20. For the single-variable model, the standard deviation was over \$1.8 million, which is shown in the corresponding cell in Figure 10.15. The interpretation is straightforward; Price and Competition Price are variables that are relevant (and useful!) for understanding the uncertainty about Sales. Without them in the previous model, we were ignoring important information. With them in the current model, the result ultimately is less uncertainty about Sales. For example, with Advertising set at \$110,000, Price at \$94.50, and Competition Price at \$96.69,

Table 10.5
Augmented data set
for sales example.

Promotion	Advertising (\$1000\$)	Price (\$)	Competition Price (\$)	Sales (\$ 1000s)
1	366	90,99	96.95	10541
2	377	90.99	93.99	8891
3	387	94.99	90.99	5905
4	418	96.99	97.95	8251
5	434	92.99	97.95	11461
6	450	95.95	93.95	6924
7	457	93.95	90,99	7347
8	466	91.95	96.95	10972
9	467	96.95	94.99	7811
10	468	92.95	96.95	10559
11	468	97.99	98.95	9825
12	475	91,95	90.99	9130
13	479	99.95	91.95	5116
14	479	96,99	95.95	7830
15	481	91.95	90.95	8388
16	490	96.99 ..	96-99	8588
17	494	96.95	91.95	6945
18	502	98.95	95.95	7697
19	505	94.99	96.99	9655
20	529	93.99	97.95	* 11516
21	532	91.99	95,99	11952
22	533	92,99	97,99	13547
23	542	93.99	92,95	9168
24	544	90.95	95.95	11942
25	547	94.99	93.95	9917
26	554	89.95	90.95	10666
27	556	96,95	95.95	9717
28	560	91,99	97.95	13457
29	561	98,99	97.95	10319
30	566	93.95	91.99	9731
31	566	94.99	94.99.	10279
32	582	98,99	91.99	7202
33	609	89.95	92.99	12103
34	612	92.95	92.99	11482
35	617	92.95	94.95	11944
36	623	94.99	91,99	9188

Figure 10.20
Regression analysis
for full model.

Regression Analysis						
	A	B	C	D	E	F
128						
129	Regression Statistics					
130						
131	Multiple R	0.98				
132	R Square	0.95				
133	Adjusted R Square	0.95				
134	Standard Error	459.10				
135	Observations	36				
136						
137	Analysis of Variance					
138		df	Sum of Squares	Mean Square	F	Significance F
139	Regression	3	137,289,638	45,763,213	217.12	0.00
140	Residual	32	6,744,669	210,771		
141	Total	35	144,034,307			
142						
143		Coefficients	Standard Error	t Statistic	P-value	Lower 95%
144						Upper 95%
145	Intercept	2199.34	3839.74	0.57	0.57	-5621.94
146	Advertising	15.05	1.17	12.83	0.00	12.66
147	Price	-503.76	28.34	-17.77	0.00	-561.50
148	Comp Price	499.67	30.56	16.35	0.00	437.42
149						

conditional expected Sales would be \$4,562,600. This is almost the same value that we had before when we used only Advertising. Now, though, you can see in Figure 10.22 that the uncertainty about Sales is much less, as reflected in the tighter distribution.

The tightness of the distribution, as indicated by the CDF of the errors and the standard deviation of that distribution, is the best available indication of the value of the model. For example, a product manager might look at the \$1.8 million standard deviation for the previous single-variable model and scoff. Such a broad distribution makes planning difficult at best. On the other hand, the \$459,100 standard deviation of the residuals in the full model may represent enough accuracy that this model could be used for planning production and distribution for a major promotion. And further model enhancements (adding other variables or improved modeling through advanced techniques) might reduce the uncertainty even further, thereby making the model that much more useful as a forecasting and planning tool.

The influence diagram in Figure 10.23 shows the entire model, including a three-point approximation of the error distribution to represent the conditional uncertainty about Sales.

Regression Analysis and Modeling: Some Do's and Don't's

The examples and discussion above give you a good idea of how regression analysis can be used to create models in decision analysis. However, you should be aware that we have barely scratched the surface of this powerful statistical technique. Complete coverage of regression is beyond the scope of this text, but many

Table 10.6
Residuals for the full
model.

Observation	Sales(Y)	$E(Y X_1, X_2, X_3)$	Residual
1	10541	10312	229
2	8891	8999	-108
3	5905	5635	270
4	8251	8572	-321
5	11461	10827	634
6	6924	7578	-654
7	7347	7212	135
8	10972	11333	-361
9	7811	7850	-39
10	10559	10859	-300
11	9825	9320	505
12	9130	8490	640
13	5116	5000	116
14	7830	8490	-660
15	8388	8561	-173
16	8588	9175	-587
17	6945	6737	208
18	7697	7849	-152
19	9655	10408	-753
20	11516	11753	-237
21	11952	11826	126
22	13547	12337	1210
23	9168	9450	-282
24	11942	12511	-569
25	9917	9521	396
26	10666	10667	-1
27	9717	9669	48
28	13457	13227	230
29	10319	9716	603
30	9731	9352	379
31	10279	10327	-48
32	7202	7054	148
33	12103	12514	-411
34	11482	11047	435
35	11944	12102	-158
36	9188	9686	-498

Figure 10.21
*CDF of residuals for
 the full model*




Figure 10.22
*CDF of sales given
 advertising =
 \$110,000, price =
 \$94.50, and competi-
 tion price = \$96.69.*


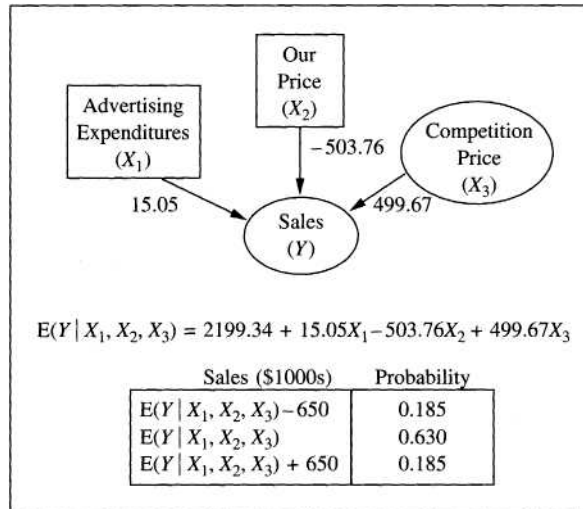


Figure 10.23
Influence diagram model of sales given advertising, price, and competition price.



excellent statistical textbooks are available. This section provides a brief survey of some of the issues in regression modeling, especially as they relate to decision analysis.

We have seen a number of approaches to creating uncertainty models, including subjective assessment, the use of theoretical distributions, and data. The use of data can be very powerful; a statistical model and analysis, when based on an appropriate data set of adequate size, is very persuasive. The drawback is that data collection can take time and resources and may not always be feasible. In such cases, an analyst might be better off relying on theoretical distributions and expert subjective judgments. Still, when data are available, they can be valuable.

Although you should by now have an understanding of regression basics (which will improve with practice on some problems), there is much more to learn. Many courses in statistics are available, as are innumerable textbooks, some of which are listed in this chapter's reference section. Also, become familiar with at least one computer program for doing regression. The examples have used the simple (in some ways primitive) regression facility that is available with Excel. Packages developed specifically for statistical analysis may provide a more complete set of analysis options, but the flexible spreadsheet environment can be very useful to a decision analyst.

Creating an uncertainty model with regression can be quite powerful. It does have some important limitations, however. Two in particular warrant consideration in the context of decision analysis. First, the data set must have an adequate number of observations. Just what "adequate" means is debatable; many rules of thumb have been proposed, and some work better than others. The approach we have described above for decision-analysis modeling will perform poorly without enough data to estimate the coefficients and the error distribution. Although "enough data" may

mean different things in different contexts, a conservative rule of thumb would be to have at least 10 observations for each explanatory variable and never less than 30 observations total.

The main reason for having adequate data is that many data points may be necessary to fully represent all of the different ways in which the variables can occur together. You would like to have as good a representation as possible of all the possible combinations. In the sales example, we would like to have occasions when prices are high, low, and in between, and for each different price level, advertising expenditures should range from high to low. Note that there is no guarantee that a large data set will automatically provide adequate coverage of all possibilities. (Inadequate coverage can result in nonsensical coefficient estimates. Advanced techniques can diagnose this situation and provide ways to get around it to some extent.) On the other hand, very small data sets almost by definition will not provide adequate coverage.

Another way to view this problem is to realize that a small number of observations means that we cannot be very sure that the b 's are accurate or reliable estimates of the β 's. When we use the b 's to calculate the conditional expected value, however, we are in essence assuming that the b 's are good estimates and that virtually all of the uncertainty in the response variable comes from the error term. In fact, with a small data set, we may also be very unsure about what the conditional expected value itself should be. In some cases you might want to model the uncertainty about the estimates explicitly and incorporate this additional uncertainty by an appropriate broadening of Y 's conditional probability distribution. Statistical procedures for generating appropriate probability intervals under these circumstances are included in most computer regression programs (although not in Excel).

Even with an adequate data set and a satisfactory model and analysis, there remains an important limitation. Recall that the regression model has a particular form; it is a linear combination of the explanatory variables. And the coefficient estimates are based on the particular set of observations in the data set. The upshot is that your model may be a terrific approximation of the relationship for the variables in the neighborhood of the data that you have. But if you try to predict the response variable outside of the range of your data, you may find that your model performs poorly. For example, our sales example above was based on prices that ranged from about \$89 to \$100. If the product manager was contemplating reducing the price to \$74.95, it is unlikely that an extrapolation of the model would give an accurate answer; there is no guarantee that the linear approximation used in the first place would extend so far outside of the range of the data. The result can be an inaccurate conditional expected value and a poor representation of the uncertainty, usually in the form of a too-narrow distribution, resulting in turn in poor planning. In practical terms, this means that the company could be taken by surprise by how many (or how few) units are sold during the promotion.

The limitation with the range of the data goes somewhat further yet when we use multiple explanatory variables, and this is related to the discussion of the data requirement above. You may try to predict the response variable for a combination of

the explanatory variables that is poorly represented in the data. Even though the value of each explanatory variable falls within its own range, the combination for which you want to create a forecast could be very unusual. In the sales example, it would not be unreasonable to find in such a competitive situation that both Price and Competition Price move together. Thus, we might be unlikely to observe these prices more than a few dollars apart. Using Price = \$89.99 and Competition Price = \$99.95 (or visa versa) could turn out to be a situation poorly represented by the data, and in this case the analysis may provide a very misleading representation of the expected Sales and the uncertainty. Although there is no simple way to avoid this problem, one helpful step is to prepare scatterplots of all pairs of explanatory variables and to ensure that the point for which you wish to predict the response variable lies within the data cloud in each of the scatterplots. Thus, in the sales example, we would be sure that the combination of prices and advertising that we wanted to analyze fell within the data clouds for each of three different scatterplots: Price versus Competition Price, Price versus Advertising, and Competition Price versus Advertising.


Regression Analysis: Some Bells and Whistles

As mentioned, our coverage of regression is necessarily limited. In this section, we look briefly at some of the ways in which regression analysis can be extended.

Nonlinear Models. First, we can easily get past the notion that the conditional expected value is a linear function of the explanatory variables. For example, suppose an analyst encounters the scatterplot in Figure 10.24, which relates average temperature (X) and energy consumption (Y), measured in kilowatt-hours, for small residences. A linear relationship between these two quantities clearly would not work well, and the reason is obvious: When the temperature deviates from a comfortable average temperature (around 20° C), we consume additional energy for either heating or air conditioning. To model this, a quadratic expression of the form $E(Y | X) = \beta_0 + \beta_1 X + \beta_2 X^2$ might work well. Implementation is easy: For each observation calculate X^2 , thereby creating a second explanatory variable. Now run the regression with two explanatory variables, X and X^2 . The coefficient estimates obtained define a parabolic curve that fits the data best in a least-squares sense. A similar approach works for fitting exponential or logarithmic functional forms, some of which have very compelling and useful economic interpretations. A textbook on regression modeling such as Neter, Wasserman, and Kutner (1989) discusses such modeling techniques in depth.

Categorical Explanatory Variables. Up to this point, all of the explanatory variables have been continuous. But not all useful measurements occur along a continuum. We might want to treat sales regions or sexes differently, for example. Or it might be useful to classify people into different age categories (child, adolescent, young adult, middle age, senior citizen) for marketing purposes. Is it possible to deal with such situations in a regression model?

Figure 10.24
A nonlinear relationship between temperature and energy consumption.



The answer is yes, and the technique is straightforward. Suppose we have marketing data (X_1, \dots, X_k) for both men and women, but we would like to treat them separately. We can create a categorical variable (X_{k+1}) that equals 1 for men and 0

$$X_{k+1} = \begin{cases} 1 & \text{if male} \\ 0 & \text{if female} \end{cases}$$

This adds an extra column to the data matrix, and in the regression equation we include the new categorical variable:

$$E(Y | X_1, \dots, X_k, X_{k+1}) = \beta_0 + \beta_1 X_1 + \dots + \beta_k X_k + \beta_{k+1} X_{k+1}$$

Now consider how this equation works for men and women separately. Women have X_{k+1} coded as a 0, so the last term on the right-hand side drops out. Thus, for women the expression becomes

$$\text{Women: } E(Y | X_1, \dots, X_k, X_{k+1} = 0) = \beta_0 + \beta_1 X_1 + \dots + \beta_k X_k$$

For men, on the other hand, $X_{k+1} = 1$, and so the last term becomes β_{k+1} :

$$\begin{aligned} \text{Men: } E(Y | X_1, \dots, X_k, X_{k+1} = 1) &= \beta_0 + \beta_1 X_1 + \dots + \beta_k X_k + \beta_{k+1} \\ &= (\beta_0 + \beta_{k+1}) + \beta_1 X_1 + \dots + \beta_k X_k \end{aligned}$$

In other words, the coefficients for X_1 through X_k are exactly the same, but the expressions have different constant terms to represent the different genders. The interpretation of β_{k+1} is that this is the incremental difference in expected Y for a man as compared to a woman. For example, if Y represents mail-order clothing purchases per year, then $\beta_{k+1} = -\$50$ means that, everything else (age, income, and so on) being equal, a man is expected to spend \$50 less per year on mail-order clothing than a woman.

Actually doing the analysis is easy. Create the categorical variable, include it as one of the explanatory variables in the regression analysis, and obtain the coefficient for the categorical variable from the computer printout. The rest is the same as before: To generate a prediction for a particular person, plug in the appropriate demographic data for all of the variables (including gender), calculate $E(Y | X_1, \dots, X_k, X_{k+1})$, and use the error distribution to construct the conditional probability distribution for Y .

Note that we had two categories (male and female), and we were able to distinguish between them with just one categorical variable. What if you have more than one category? The rule is that you should use *one less categorical variable than there are categories*. For example, take the age example that we mentioned. There are five categories: child, adolescent, young adult, middle age, and senior citizen. (These categories must, or course, be defined clearly enough so that any individual can be placed unambiguously into one of the categories.) Now define four categorical variables:

$$\begin{aligned} X_{k+1} &= \begin{cases} 1 & \text{if child} \\ 0 & \text{otherwise} \end{cases} \\ X_{k+2} &= \begin{cases} 1 & \text{if adolescent} \\ 0 & \text{otherwise} \end{cases} \\ X_{k+3} &= \begin{cases} 1 & \text{if young adult} \\ 0 & \text{otherwise} \end{cases} \\ X_{k+4} &= \begin{cases} 1 & \text{if adult} \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Each individual will have a 1 or a 0 for each of these variables. Imagine seeing the values of these variables for some particular person. If you see that $X_{k+3} = 1$ and the others are zero, you know that this person is a young adult, and likewise for the other variables. In fact, there will never be more than one nonzero value, and that nonzero value will be a 1, indicating the person's category. But what if you see zeros for all four of these variables? Simple: This person falls into none of the four categories associated with the variables and therefore must be a senior citizen.

The rule of having one less categorical variable than you have categories is easy to forget. It may seem natural to set up one variable for each category. There is a fail-safe mechanism, however. Your computer program will not run with the extra categorical variable! It will most likely produce some cryptic error message like ANALYSIS FAILED — MATRIX SINGULAR, which is its way of indicating that it is unable to solve the required equations. Although this type of message can indicate other problems, if you have categorical variables, it almost surely means that you have one too many.

Categorical Response Variable (Y) The regression procedure we have presented is designed especially for situations in which F is a continuous variable. Often, though, you want to use data to predict whether an observation falls into a particular category (e.g., whether a particular customer will place another order or whether a particular individual will develop cancer). In decision-analysis terms, you would want to specify the conditional probability of falling into that category, given the values of the conditioning variables.

Statistical procedures are available to use data for this very purpose. However, these procedures have special requirements, and it is beyond the scope of this short introduction to describe how to use them. Suffice it to say that the standard approach laid out above for continuous variables is *not* appropriate for categorical response variables. If you would like to use data to model conditional probabilities for such a variable, consult one of the advanced regression texts mentioned at the end of this chapter.

Regression Modeling: Decision Analysis Versus Statistical Inference

Our approach in this section has very much reflected the decision-analysis paradigm. We have used data as a basis for constructing a model of the uncertainty surrounding a response variable. Our focus has been particularly on coming up with a conditional probability distribution for F , given values of conditioning variables. The motivation has been that understanding and modeling the relationships among the explanatory and response variables can be of use in making decisions.

If you have been exposed to regression analysis before, perhaps in a statistics course, you no doubt realize that the perspective taken here is somewhat different from the conventional approach. Although our approach is not entirely inconsistent with conventional thinking, there are some key differences:

- Statistical inference often focuses on various types of *hypothesis tests* to answer questions like, "Is a particular coefficient equal to zero?" Or, "Is this set of explanatory variables able to explain any of the variation in Y ?" As decision analysts, such questions would more likely be framed as, "Is it reasonable to include X_i as an explanatory variable, or should it be left out?" A decision analyst might answer such a question with standard sensitivity-analysis techniques, which would permit consideration of the context of the current decision, including costs, benefits, and other risks. The conventional hypothesis-testing approach ignores these issues.
- Conventional regression modeling and the inference that is typically done (including hypothesis tests and the calculation of confidence and prediction intervals for Y) rely on the assumption that the errors follow a normal distribution. While the analysis can be quite powerful when this condition is met, it can be limiting or even misleading for situations in which normality is inappropriate. As decision analysts, we have focused on using the data to create a useful model of the uncertainty in the system for forecasting and decision-making purposes, which includes coming up with a suitable error distribution.

In the standard regression approach, a question left unaddressed is how to model uncertainty in the explanatory variables. For those variables that are themselves uncertain (like Competition Price in our example), it may be important from a decisions-analysis point of view to model that uncertainty. The influence-diagram approach that we have developed in this chapter provides the basis for such modeling. Either historical data or subjective judgments can be used to assign a probability distribution to an uncertain explanatory variable. Thus included in the model, that uncertainty becomes an integral part of any further analysis that is done.

An Admonition: Use with Care

With a little knowledge of regression, you have a very powerful data-analysis tool. A temptation is to apply it indiscriminately to a wide variety of situations. Part of its appeal is that it is using "objective" data, and hence the analysis and results are free of subjective judgments and bias. This is far from true! In fact, you should appreciate the broad range of judgments and assumptions that we made: linear expected values, the error probability distribution that does not change shape, the inclusion of particular explanatory variables in the model. More fundamentally, we implicitly make the judgment that the past data we use are appropriate for understanding current and future relationships among the variables. When we use a conditional probability distribution based on regression analysis, all of these judgments are implicit. Thus, the admonition is to use regression with care (especially with regard to its data requirements) and to be cognizant of its limitations and implicit assumptions. Regression is a rich and complex statistical tool. This section has provided only a brief introduction: further study of regression analysis will be especially helpful for developing your skill in using data for modeling in decision analysis.

Natural Conjugate Distributions (Optional)

So far we have considered modeling uncertainty using subjective judgments, theoretical models, and data. It is clear that we can mix these techniques to varying degrees. We have seen examples in which we have used a theoretical model to represent a subjective probability distribution (the use of the beta distribution in the soft pretzel problem), others in which data were used to estimate parameters for a theoretical distribution (the halfway-house data), and still others in which data were used to modify subjective probabilities regarding the parameter of a theoretical distribution (the soft pretzel examples using the binomial and Poisson distributions). Our topic now continues in this vein. We will see how to use Bayes' theorem in a systematic way to use data to update beliefs in the context of theoretical models.

Recall the Poisson example in Chapter 9 (page 322) in which we wondered about the quality of a new location for the soft pretzel kiosk. The new location could have been good, bad, or dismal, with $P(\text{Good}) = 0.7$, $P(\text{Bad}) = 0.2$, and $P(\text{Dismal}) = 0.2$.

Each characterization is associated with a specific value of m , the expected number of pretzels sold per hour; a good location implies $m = 20$, bad implies $m = 10$, and dismal implies $m = 6$. The uncertainty about the nature of the location also can be thought of as uncertainty about the parameter m : $P(m = 20) = 0.7$, $P(m = 10) = 0.2$, and $P(m = 6) = 0.1$. Furthermore, we might have had more than three possible values for m ; in fact, the location might have been anything between dismal and wonderful, or might have had m values anywhere from 0 to, say, 50 (or even higher).

Uncertainty about a parameter (such as m) often can be modeled with a continuous probability distribution. As before, this would be the prior probability distribution. Then, as in the Poisson example, we may obtain data that provide information about the parameter of the process. Bayes' theorem provides the mechanism for using the data to update the prior probability distribution in order to arrive at a posterior probability distribution.

In this section, we will explore this process in some detail. It turns out that the mathematical details are such that performing the calculations can be quite difficult unless the prior distribution and the distribution for the data match in a particular way. We will look at two such situations. The first is when the decision maker is uncertain about the parameter p in a binomial distribution and represents that uncertainty with a beta prior distribution. The decision maker then observes data, the outcome of a binomial random variable, and uses these data to update the prior distribution. In the second situation, the decision maker is uncertain about the mean μ of a normal distribution and represents that uncertainty with another normal distribution for μ . Now the decision maker takes a sample, calculates a sample mean, and uses this information to update the prior distribution for μ .

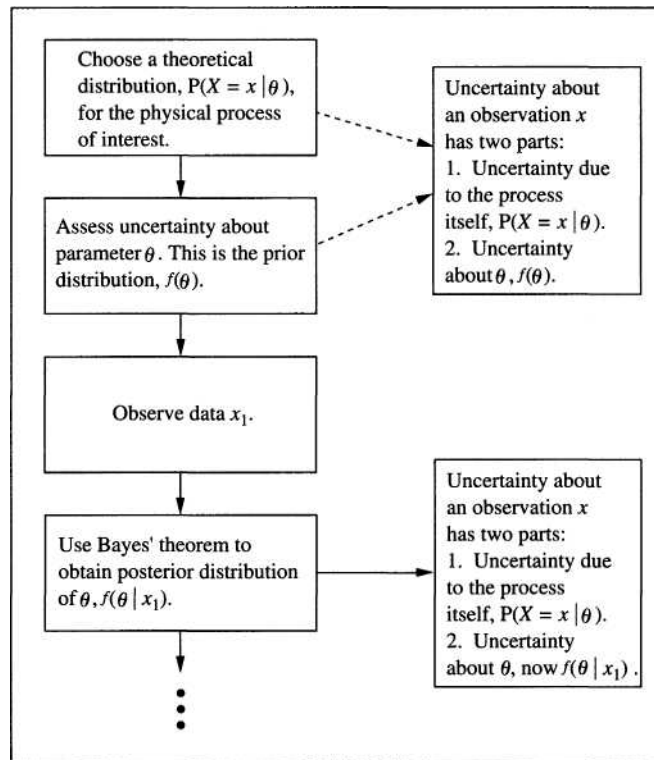
Our plan of attack is as follows. First, we will discuss the overall process. The process is not particularly complicated, but the concepts involved are somewhat demanding. A thorough understanding of the process requires attention to several conceptual issues. Next, we will show how the process works in the two specific situations that were mentioned above, and we will see examples. We will end with a discussion of how the uncertainty about a distribution parameter influences one's prediction or forecast.

Uncertainty About Parameters and Bayesian Updating

The flowchart (*not* an influence diagram) in Figure 10.25 shows the process that we will go through. The symbol θ is used to represent a parameter of a theoretical distribution in which we may be interested. The process goes as follows.

First, the decision maker chooses some theoretical distribution to model his or her beliefs regarding an uncertain quantity X (for example, a normal distribution to model the scores of students on a standardized exam or a Poisson distribution to model the number of customers who arrive at a soft pretzel kiosk). Let $P(X = x | \theta)$ denote this theoretical distribution, which has some parameter θ about which the decision maker is uncertain. Being uncertain about θ , the decision maker assesses a *prior probability distribution* $f(\theta)$ for the parameter. At this point, $P(X = x | \theta)$ and $f(\theta)$ together embody all of the decision maker's beliefs about the possible outcomes

Figure 10.25
Using data to update uncertainty about parameter θ via Bayes' theorem.



of the uncertain quantity X . If necessary, the decision maker can use this model of the uncertainty in solving a decision problem. Our concern here, however, is how to use data. Thus, the next step in the process is the acquisition of data and the incorporation of the new information into the model.

Suppose that the decision maker, having assessed $P(X = x | \theta)$ and $f(\theta)$, now observes an x , an outcome from the physical system (for example, a sample of exam scores or actual arrivals at the soft pretzel kiosk over a period of time). Label the data as x_1 . The data contain information that the decision maker can use to refine his or her probabilities. In particular, Bayes' theorem allows for the updating of the prior distribution $f(\theta)$ so that it becomes a posterior distribution $f(\theta | x_1)$:

$$f(\theta | x_1) = \frac{P(x_1 | \theta) f(\theta)}{\int_{-\infty}^{+\infty} P(x_1 | \theta) f(\theta) d\theta} \tag{10.1}$$

Equation (10.1) is simply Bayes' theorem when $f(\theta)$ is a continuous distribution. The output of Bayes' theorem in this case is a posterior distribution of θ after having

observed x_1 . That is, $f(\theta | x_1)$ represents the decision maker's uncertainty about θ after having seen data x_1 .

The term $P(x_1 | \theta)$ in Equation (10.1) is called the *likelihood function*. We calculated likelihoods for the binomial and Poisson soft pretzel examples in Chapter 9: We did it when we calculated the probability of observing the data for each possible parameter value. In the binomial case, we calculated binomial probabilities for the observed data depending on whether the pretzels were a hit ($p = 0.3$) or a flop ($p = 0.1$). In the Poisson case, we calculated Poisson probabilities for the observed number of arrivals depending on whether the new location was good ($m = 20$ per Hr), bad ($m = 10$ per Hr), or dismal ($m = 6$ per Hr).

The flowchart in Figure 10.25 is shown with three dots at the bottom, indicating that the process can continue. That is, the decision maker now has uncertainty about θ , as represented by $f(\theta | x_1)$. New data (x_2) might be observed, and the decision maker again can incorporate this new information via Bayes' theorem. The result would be $f(\theta | x_1, x_2)$. This sequential updating process can continue indefinitely.

In principle, Equation (10.1) can be used regardless of the form of the likelihood function or the proper distribution $f(\theta)$. With specific mathematical expressions for these functions, it always is possible to have a computer crunch the numbers to generate a posterior distribution $f(\theta | x_1)$. It may not be possible to write down a mathematical expression for $f(\theta | x_1)$, but the computer could grind out the calculations, evaluate $f(\theta | x_1)$ for many different values of θ , and finally plot the distribution. But an easier way to do this would be nice.

Fortunately, there is an easier way in certain situations. Recall that we are moving from a prior distribution to a posterior distribution. In certain cases, the prior and posterior distributions have the same form. This form is called the *natural conjugate distribution* for the theoretical distribution that represents the physical process. Moreover, in these cases relatively simple rules guide the use of observed data in adjusting the prior distribution to obtain the posterior. In other words, this elegant result reduces the seemingly complicated Equation (10.1) to a few simple formulas in these cases. Two such cases involve the binomial and normal distributions. We will discuss each one in turn, using examples.

Binomial Distributions: Natural Conjugate Priors for p

In this case, the chosen theoretical distribution to represent the physical process is binomial with parameters n and p . That is, the decision maker knows that the process generates many observations that are either successes or failures, and that in a group of n observations there will be some number X of successes. Thus, $P(X = x | \theta)$ in this case is equal to the binomial probability $P_B(X = r | n, p)$, which is given in Equation (9.1).

Of course, $P_B(X = r | n, p)$ is easy to calculate for any value of r as long as p is known. In this case, however, we want to incorporate uncertainty about p . The natural conjugate prior distribution for p is a beta distribution. That is, if the decision maker uses the beta distribution to model the prior uncertainty about p and then updates that beta prior on the basis of an observation from the process, then the posterior distribution of p also will be a beta distribution. Note that the beta distribution is a reasonable model, because p must be between 0 and 1. Here is how the updating process works:

The decision maker considers a physical process and concludes that its uncertainty can be represented with a binomial distribution. That is, out of n trials, the probability of r successes is $P_B(X = r | n, p)$.

The decision maker assesses a beta prior distribution for p , $f_\beta(p | r_0, n_0)$. That is, the decision maker carefully considers the uncertainty about p and concludes that it can be adequately represented with a beta distribution having parameters r_0 and n_0 . Based on this prior distribution for p , a good estimate for p would be the mean of this distribution, r_0/n_0 .

The decision maker observes n_1 independent trials, and r_1 of these are successes. Based on these data alone, of course, a good estimate for p would be r_1/n_1 . The issue is how to combine this information with the prior distribution. Combining the data with the prior distribution via Bayes' theorem gives a posterior distribution for p that is also a beta distribution, $f_\beta(p | r^*, n^*)$, where

$r^* = r_0 + r_1$	(10.2a)
and	
$n^* = n_0 + n_1$	(10.2b)

That is, the decision maker only has to add the r 's and n 's to find the parameters of the posterior distribution for p .


In addition, we can see how this process could continue as new data are observed. The posterior $f_\beta(p | r^*, n^*)$ becomes the prior distribution in the next round. Data r_2 and n_2 are observed. Thus, the parameters for the posterior distribution after observing r_2 and n_2 are $r^{**} = r^* + r_2 = r_0 + r_1 + r_2$ and $n^{**} = n^* + n_2 = n_0 + n_1 + n_2$.

In Chapter 9 we discussed briefly that in subjectively assessing a beta distribution, the decision maker may think of r and n as representing "equivalent information." If the decision maker believes that the available information is equivalent to having observed 5 successes in 10 trials, then appropriate beta parameters would be $r = 5$ and $n = 10$. In performing the updating of a prior distribution via Bayes' theorem here, we can see just how appropriate this analogy is. Essentially, the r 's and n 's are added together as if they had all resulted from sample observations.

As an example, let us return to the soft pretzel example in Chapter 9, specifically pages 334-335. There we considered whether to proceed with the marketing of your pretzels. Of course, if you are convinced that your new pretzel will be a great seller, you certainly would undertake the project. The problem is that the proportion (Q) of the market that your pretzel might capture is uncertain. In the example, this uncertainty about Q is represented with a beta distribution having parameters $r = 1$ and $n = 4$. Let us take this as a prior distribution, $f_\beta(Q | r_0 = 1, n_0 = 4)$. The expected value of this distribution is $E(Q) = 1/4 = 0.25$, which would be a good estimate for Q based on the prior distribution.

Now suppose you run a taste test. You bake 20 pretzels, and 7 of 20 tasters prefer your pretzel over the competition's. Thus, $r_1 = 7$ and $n_1 = 20$. Combining these

Figure 10.26
 Prior and posterior distributions for the market proportion captured by the soft pretzel



data with the prior distribution, the posterior probability distribution for Q is $f_{\beta}(Q | r^* = 1 + 7 = 8, n^* = 4 + 20 = 24)$. The posterior expected value of Q is $8/24 = 0.33$. The data have considerably affected your beliefs about Q . Figure 10.26 shows both the prior and posterior distributions. The additional data have greatly reduced the uncertainty about Q .

Normal Distributions: Natural Conjugate Priors for μ

Suppose the decision maker chooses a normal distribution to represent the physical process. That is, the decision maker knows that the outcomes (the x 's) that arise from the process are numbers that can take on any value in a continuum, and believes that a normal distribution provides a good representation for the x 's distribution. In this case, $P(x | \theta)$ is the expression $f_N(x | \mu, \sigma)$ from Equation (9.4).

We will assume that the standard deviation of the process, σ , is known. The decision maker, however, does not know μ but is willing to assess a probability distribution for it. The natural conjugate prior distribution for μ also turns out to be a normal distribution. That is, if the decision maker chooses a normal prior distribution to represent the uncertainty about μ and then updates that prior based on observations from the process, the resulting posterior distribution for x also will be normal. Let us suppose that the decision maker assesses a normal distribution for μ , $f_N(\mu | m_0, \sigma_0)$.

This situation can be a bit confusing because so many different normal distributions are floating around. First, the original distribution is normal, $f_N(x | \mu, \sigma)$. But we also have normal distributions for the parameter μ itself. The prior is $f_N(\mu | m_0, \sigma_0)$, and the posterior also will be normal, $f_N(\mu | m^*, \sigma^*)$. Here are the steps in the updating process:

- 1 The decision maker considers a physical process and concludes that the uncertainty of the process can be represented with a normal distribution. That is, the distribution of observations from the process is $f_N(x | \mu, \sigma)$.
- 2 The decision maker assesses a normal prior distribution for μ , $f_N(\mu | m_0, \sigma_0)$. That is, the decision maker carefully considers the uncertainty about μ and concludes

that it can be modeled with a normal distribution having parameters m_0 and σ_0 . Based on this prior distribution for μ , a good estimate for μ would be the mean m_0 .

The decision maker actually observes n_1 independent observations from the process, and the average of these sample observations is \bar{x}_1 . Based on these data alone, of course, a good estimate for μ would be \bar{x}_1 .

Combining the data with the prior distribution via Bayes' theorem gives a posterior distribution for μ that also is a normal distribution, $f_N(\mu | m^*, \sigma^*)$, where

$m^* = \frac{m_0\sigma^2/n_1 + \bar{x}_1\sigma_0^2}{\sigma^2/n_1 + \sigma_0^2} \tag{10.3a}$
<p>and</p>
$\sigma^* = \sqrt{\frac{\sigma_0^2\sigma^2/n_1}{\sigma^2/n_1 + \sigma_0^2}} \tag{10.3b}$


That is, the decision maker combines the prior information and the data via Equations (10.3a) and (10.3b) to obtain the posterior distribution of μ .

As with the binomial model, one can see how the sequential incorporation of data would be done in this case. Observing the first sample would lead to the posterior as given above. In the next round, m^* and σ^* become the parameters for a new prior distribution that is updated on the basis of a second set of observations, using Equations (10.3a) and (10.3b) with m^* and σ^* in place of m_0 and σ_0 . Further sequential updating would follow the same pattern.

In the binomial situation above, it was easy to see how the prior information could be interpreted as "equivalent" sample information. The same is true here, although it is not evident from (10.3a) and (10.3b). Suppose, however, that we write $\sigma_0^2 = \sigma^2/n_0$. Then n_0 is interpreted as our "equivalent sample size." Having made this transformation, Equations (10.3a) and (10.3b) become

$m^* = \frac{n_0m_0 + n_1\bar{x}_1}{n_0 + n_1} \tag{10.4a}$
<p>and</p>
$\sigma^* = \sqrt{\frac{\sigma^2}{n_0 + n_1}} \tag{10.4b}$

Figure 10.27
 Prior and posterior
 distributions for x in
 the halfway-house
 example.



From these expressions, it is clear how n_0 and n_1 are interacting. The posterior mean m^* is simply a weighted average of m_0 and \bar{x}_1 , where the weights are based on the "sample sizes" n_0 and n_1 . Posterior variance σ^{*2} has the same form as σ_0^2 , with the cumulative sample size $n_0 + n_1$ used as the divisor of the process variance σ^2 .

As an example of the updating process for the normal distribution, let us again look at the halfway-house data. Suppose that a decision maker believes that yearly bed-rental costs are normally distributed with mean μ and standard deviation $\sigma = \$220$. The expected value μ is unknown, but the decision maker assesses a normal distribution with mean $m_0 = \$345$ and standard deviation $\sigma_0 = \$50$, $f_N(\mu | m_0 = \$345, \sigma_0 = \$50)$. For example, this means that the decision maker believes there is roughly a 68% chance that μ is between \$295 and \$395, and roughly a 95% chance that μ is between \$245 and \$445.

Now suppose that the decision maker obtains the 35 observations listed in Table 10.1. The sample mean of these observations is \$380.40. We have $n_1 = 35$ and $\bar{x}_1 = \$380.40$. Applying Equations (10.3a) and (10.3b), we obtain the posterior distribution of μ , a normal distribution with parameters $m^* = \$367.80$ and $\sigma^* = \$29.80$. The prior and posterior distributions are illustrated in Figure 10.27. Notice that, with the added information, the distribution for μ is tighter, reflecting less uncertainty about μ .

Predictive Distributions

We have represented the uncertainty about the physical process under investigation with the distribution $P(X = x | \theta)$. That is, if we only knew θ , we would have everything we need of the distribution of the x 's. We do not know θ , however, but we do have a probability distribution for it, $f(\theta)$. To give a distribution for the x 's, we must combine $P(X = x | \theta)$ and $f(\theta)$. As previously mentioned, these two kinds of uncertainty, taken together, provide a complete model of the uncertainty that the decision maker faces regarding the x 's. We simply must figure out how to put these pieces together.

To see this, imagine the binomial process. The probability of r successes in n trials is $P_b(X = r | n, p)$. But we need to know p in order to calculate probabilities using Equation (9.1). What we can do, however, is take something like an average over all possible values for p . This is exactly the approach we will take.

If we want a distribution for the x 's that does not depend on θ , we want $P(X = x)$, not $P(X = x | \theta)$. We will call $P(X = x)$ the *predictive distribution* because it describes our uncertainty about X — which x values are most likely, the probability that X will fall into different intervals, and so on. To get $P(X = x)$, we must "integrate out" θ :

$$P(X = x) = \int_{-\infty}^{+\infty} P(X = x | \theta) f(\theta) d\theta.$$

You might recognize this expression as the denominator of Bayes' theorem as it was written in Equation (10.1). Figuring out this integral will not be any easier than it was in Equation (10.1). It will always be possible to use a computer to calculate this integral numerically, but numerical integration techniques are beyond the scope of this text. In cases in which we have a natural conjugate prior distribution for q , it may be possible to find a simple expression for this distribution. We will take a look at the normal and binomial processes.

Predictive Distributions: The Normal Case

If the physical process follows a normal distribution, the integral becomes

$$\begin{aligned} P(X = x) &= \int_{-\infty}^{+\infty} f_N(x | \mu, \sigma) f(\mu | m_0, \sigma_0) d\mu \\ &= f_N\left(x | m_0, \sigma_p = \sqrt{\sigma^2 + \sigma_0^2}\right) \end{aligned}$$

We will express this in words. Taking into account the prior uncertainty about μ , the distribution of the next sampled observation from the process is normal with mean m_0 and variance $\sigma_p^2 = \sigma^2 + \sigma_0^2$. The fact that the mean for x is m_0 is reasonable; our best guess for x would be \bar{x} , and our best guess for μ is m_0 . The fact that the variance is $\sigma^2 + \sigma_0^2$ shows that indeed the overall uncertainty about x is composed of uncertainty about the process that generates the x 's (σ^2) and uncertainty about M (σ_0^2). If we knew exactly what \bar{x} was, then the variance of the distribution simply would be σ^2 .

What we have just done can be generalized to any normal distribution for μ . If some data already have been collected and the decision maker has a posterior distribution for μ , $f_N(\mu | m^*, \sigma^*)$, then the predictive distribution for x is normal with mean m^* and variance $\sigma_p^2 = \sigma^2 + \sigma^{*2}$.

Continuing our halfway-house example, suppose that, having collected the data and updated the prior distribution for μ , the decision maker would like to express the uncertainty about the next yearly bed rental (x). The distribution desired is the predictive distribution. We know that $\sigma = \$220$ and that the posterior distribution for μ was normal with mean $m^* = \$367.80$ and standard deviation $\sigma^* = \$29.80$. Thus, the predictive distribution will be normal with mean $\$367.80$ and variance $\sigma_p^2 = (220)^2 + (29.8)^2$. Performing the calculations, we find that $\sigma_p = \$222$, which is just barely more than the process standard deviation of $\$220$. Thus, as we might have suspected, most of the uncertainty about the bed-rental cost is simply the result of fluctuation in these costs. A little bit of uncertainty results from the uncertainty about μ .

Predictive Distributions: The Binomial Case

If the physical process is one that can be represented with a binomial distribution, then the matter is slightly more complicated. We really want to know how likely we would be to get different numbers of successes in n trials. The integral to find the predictive distribution $P(X = r)$ becomes

$$\begin{aligned} P(X = r) &= \int_0^1 P_B(r | n, p) f_\beta(p | r_0, n_0) dp \\ &= \frac{(r + r_0 - 1)! (n + n_0 - r - r_0 - 1)! n! (n_0 - 1)!}{r! (r_0 - 1)! (n - r)! (n_0 - r_0 - 1)! (n + n_0 - 1)!} \end{aligned}$$

This rather unusual looking expression is called a *beta-binomial* probability distribution. The number of successes r can take on only discrete values of $0, 1, \dots, n$. The probability for any of these outcomes can be calculated by plugging values for r, n, r_0 , and n_0 into the beta-binomial expression.

As with the normal distribution, this result generalizes as long as the distribution of p is a beta distribution. For example, in the case of our soft pretzels, the posterior distribution for p was a beta distribution with parameters $r^* = 8$ and $n^* = 24$. The predictive distribution of r (the number of tasters out of n who would prefer the new pretzel) would be the same as given above, but with r^* and n^* replacing r_0 and n_0 . For example, the probability that 6 out of the next 12 tasters will prefer the new pretzel would be

$$\begin{aligned} P(X = 6) &= \frac{(r + r^* - 1)! (n + n^* - r - r^* - 1)! n! (n^* - 1)!}{r! (r^* - 1)! (n - r)! (n^* - r^* - 1)! (n + n^* - 1)!} \\ &= \frac{(6 + 8 - 1)! (12 + 24 - 6 - 8 - 1)! 12! (24 - 1)!}{6! (8 - 1)! (12 - 6)! (24 - 8 - 1)! (12 + 24 - 1)!} \\ &= \frac{13! 21! 12! 23!}{6! 7! 6! 15! 35!} \\ &= 0.1116 \end{aligned}$$

Probabilities for other possible values of r can be calculated in the same way. The use of data to update information in a systematic way via Bayes' theorem can be very useful in decision analysis. We have spent much time in discussing nat-

ural conjugate distributions here, and they can be quite helpful. What is perhaps more critical, however, is that you understand the concept that one can be uncertain about a parameter and encode that uncertainty as a probability distribution. This uncertainty then has an impact on the decision maker's probability of future outcomes. Thus, the underlying ideas of distributions on parameters and predictive distributions for future outcomes are quite basic.

In the next chapter, we will take this process a step further. In simulation we will generate random variables from theoretical probability distributions with specific parameters. This is a standard technique in decision analysis and risk analysis. But we also will be able to specify our parameters as uncertain. By specifying parameters themselves as probabilistic, we will be doing the simulation counterpart of the natural conjugate analysis that we have discussed here. Furthermore, we will not be limited to certain families of distributions. We will be able to specify any kind of probability distribution that seems reasonable; the computer will do the calculations for us.

A Bayesian Approach to Regression Analysis (Optional)

Can the Bayesian approach described above be applied to regression analysis as well as the normal and binomial models described in the last section? The answer is yes, although certain assumptions must be made. The most important is that some sort of theoretical distribution for the errors must be assumed; this provides the basis for defining the likelihood function for the data. A typical assumption is that the errors follow a normal distribution with variance σ_e^2 .

The general Bayesian approach is to model uncertainty about the model parameters. In a regression model, the parameters would be the error variance σ_e^2 and the coefficients β_0, \dots, β_k . It turns out that there are natural conjugate prior distributions for these parameters and that data can be used to update these priors. Moreover, once the updating has been done, one can generate predictive distributions for future observations of the Y variable, given the X 's. Thus, the Bayesian approach provides a systematic way to incorporate uncertainty about the coefficients as well as the error term, thereby giving a more complete accounting of uncertainty in the system. For more discussion of Bayesian regression, see Zellner, (1971) or Berry (1995).

SUMMARY

In this chapter we have seen some ways in which data can be used in the development of probabilities and probability distributions for decision analysis. We began with the basics of constructing histograms and empirically based CDFs. A short discussion concerned the use of data to estimate parameters for theoretical distributions. The last part of the chapter discussed the use of data to update prior beliefs about parameters of a distribution. This process is based on Bayes' theorem. We saw how natural conjugate distributions worked for both normal and binomial cases. Finally, we discussed predictive distributions, or the distribution for future outcomes of a process when the parameters are uncertain.

EXERCISES

- 10.1** Explain in your own words the role that data can play in the development of models of uncertainty in decision analysis.
- 10.2** Why might a decision maker be reluctant to make subjective probability judgments when historical data are available? In what sense does the use of data still involve subjective judgments?
- 10.3** Suppose that an analyst for an insurance company is interested in using regression analysis to model the damage caused by hurricanes when they come ashore. The response variable is Property Damage, measured in millions of dollars, and the explanatory variables are Diameter of Storm, Barometric Pressure, Wind Speed, and Time of Year.
- What subjective judgments, both explicit and implicit, must the analyst make in creating and using this model?
 - If X_1 , Diameter of Storm, is measured in miles, what is the interpretation of the coefficient β_1 ?
 - Suppose the analyst decides to introduce a categorical variable X_5 , which equals 1 if the eye of the hurricane comes ashore near a large city (defined as a city with population $\geq 500,000$) and 0 otherwise. How would you interpret β_5 ?
- 10.4** Estimate the 0.65 and 0.35 fractiles of the distribution of yearly bed-rental costs, based on the CDF in Figure 10.4.
- 10.5** Choose appropriate intervals and create a relative frequency histogram based on the half way-house data in Table 10.1.
- 10.6** Estimate the 0.25 and 0.75 fractiles of the conditional distribution for Sales in Figure 10.22.

QUESTIONS AND PROBLEMS

- 10.7** It was suggested that five is the minimum number of data points for the least likely category in constructing histograms. In many cases, however, we must estimate probabilities that are extremely small or for which relatively few data are available. In the example about machine failures at the beginning of the chapter (pages 356-357), suppose that we had one day out of 260 in which three failures occurred. What do you think we should do in such a situation?
- 10.8** As discussed in the text, it often is possible to use a theoretical distribution as an approximation of the distribution of some sample data. It is always important, however, to check to make sure that your data really do fit the distribution you propose to use.

Consider the halfway-house data again. We calculated the sample mean $\bar{x} = 380.4$ and the sample standard deviation $s = 217.6$. Taking these as approximately equal to μ and σ , use the Z table in Appendix E to find normal cumulative probabilities for dollar costs of 200, 300, 400, 500, 600, and 700. How do these theoretical normal probabilities com-

pare with the data-based probabilities from the CDF in Figure 10.4? Do you think that the normal distribution is an appropriate distribution to use in this case? You might sketch the normal distribution CDF and superimpose it on Figure 10.4 to obtain a visual feel for how the normal distribution fits the data.

10.9 A scientist collected the following weights (in grams) of laboratory animals:

9.79	9.23	9.11
9.62	8.73	11.93
10.39	8.68	9.76
9.59	11.49	9.86
11.41	9.60	7.24

- a Use these data to create a data-based CDF for the distribution of weights for lab animals. Estimate the probability that an animal's weight will be less than 9.5 grams.
 - b Fit a normal distribution to these data on the basis of the sample mean and sample variance. Use this normal distribution to estimate the probability that an animal's weight will be less than 9.5 grams.
 - c Do you think the normal distribution used in part b is a good choice for a theoretical distribution to fit these data? Why or why not?
- 10.10** A plant manager is interested in developing a quality-control program for an assembly line that produces light bulbs. To do so, the manager considers the quality of the products that come from the line. The light bulbs are packed in boxes of 12, and the line produces several thousand boxes of bulbs per day. To develop baseline data, some of the workers test all the bulbs in 100 boxes. They obtain the following results:

No. of Defective Bulbs/Box	No. of Boxes
0	68
1	27
2	3
3	2

- a Plot a histogram of these results.
 - b What kind of theoretical distribution might fit this situation well? Explain.
 - c Estimate parameters for the distribution you chose in part b by calculating the sample mean and variance. Use the theoretical distribution with its estimated parameters to estimate P(0 Defective), P(1 Defective), P(2 Defective), and P(3 Defective). How do these estimates compare to the relative frequency of these events in the data?
- 10.11** A retail manager in a discount store wants to establish a policy of the number of cashiers to have on hand and also when to open a new cash register. The first step in this process is to determine the rate at which customers arrive at the cash register. One day, the manager observes the following times (in minutes) between arrivals of customers at the cash registers:

0.1	2.6	2.9	0.5
1.2	1.8	4.8	3.3
1.7	0.2	1.5	2.0
4.2	0.6	1.0	2.6
0.9	3.4	1.7	0.4

- Plot a CDF based on these data.
- What kind of theoretical distribution do you think would be appropriate for these data? Why?
- Calculate the sample mean and sample standard deviation of the data, and use these to estimate parameters for the theoretical distribution chosen in part b.
- Plot a few points from the CDF of your theoretical distribution, and draw a smooth curve through them. How does this theoretical curve compare to the data-based CDF?

10.12 An ecologist studying the breeding habits of birds sent volunteers from the local chapter of the Audubon Society into the field to count nesting sites for a particular species. Each team was to survey five acres of land carefully. Because she was interested in studying the distribution of nesting sites within a particular kind of ecological system, the ecologist was careful to choose survey sites that were as similar as possible. In total, 24 teams surveyed five acres each, and reported the following numbers of nesting sites in each of the five-acre parcels:

7	12	6	9
5	2	9	9
7	3	9	9
5	1	7	10
1	8	6	3
4	5	3	13

- Plot a histogram for these data.
- What kind of theoretical distribution might you select to represent the distribution of the number of nesting sites? Why?
- Estimate the parameters of the theoretical distribution by calculating the sample mean and sample variance.
- Plot the probability distribution based on your theoretical model from parts b and c. How does it compare to your histogram from part a?

10.13 Decision analyst Sandy Baron has taken a job with an up-and-coming consulting firm in San Francisco. As part of the move, Sandy will purchase a house in the area. There are two houses that are especially attractive, but their prices seem high. To study the situation a bit more, Sandy obtained data on 30 recent real-estate transactions involving properties roughly similar to the two candidates. The data are shown in Table 10.7. The Attractiveness Index measure is a score based on several aspects of a prop-

Table 10.7
Data for 30 real-estate transactions in Problem 10.13.

Property	House Size <Sq. Ft)	Lot Size (Acres)	Attractiveness Index	Sale Price (\$1000s)
1	3000	3.6	64	550
2	2300	1.2	69	461
3	3300	1.3	72	501
4	2100	3.2	71	455
5	3900	1.1	40	503
6	3100	2.0	74	529
7	3600	1.6	69	478
8	2900	2.5	85	562
9	2000	2.6	70	417
10	3500	1.3	74	566
11	3100	2.3	79	494
12	3200	1.5	75	515
13	2800	1.3	62	490
14	3300	3.3	62	537
15	3000	3.9	70	527
16	3400	2.4	81	577
17	2800	1.7	77	490
18	2000	3.4	67	486
19	2400	2.9	68	450
20	3600	2.9	84	674
21	2400	1.9	75	454
22	3000	2.8	63	523
23	2200	3.6	78	469
24	3600	2.4	73	628
25	2900	1.1	85	570
26	3000	4.4	69	564
27	3100	1.8	54	444
28	2200	2.1	75	494
29	2500	3.9	61	479
30	2900	1.1	74	477

erty's character, including overall condition and features (e.g., swimming pool, view, seclusion).

- a Run a regression analysis on the data in Table 10.7 with Sale Price as the response variable and House Size, Lot Size, and Attractiveness as explanatory variables. What are the coefficients for the three explanatory variables? Write out the

expression for the conditional expected Sale Price given House Size, Lot Size, and Attractiveness.

b Create a CDF based on the residuals from the regression in part a. **c**

The two properties that Sandy is considering are as follows:

House	List Price	House Size	Lot Size	Attractiveness
1	575	2700	1.6	75
2	480	2000	2.0	80

What is the expected Sale Price for each of these houses according to the regression analysis from part a?

d Create a probability distribution of the Sale Price for each house in part c. Where does the List Price for each house fall in its respective distribution? What advice can you give Sandy for negotiating on these houses?

- 10.14** Reilly Global Communications (RGC), Inc., sells communications systems to companies that have worldwide operations. Over the next year, RGC will be instituting important improvements in its manufacturing operations. The CEO and majority stockholder, Terry Reilly, has been extremely optimistic about the outlook for the firm and made the statement that sales will certainly exceed \$6 million in the first half of 1996.

Semiannual data for RGC over the past 19 years are presented in Table 10.8. The Spending Index is related to the amount of disposable income that consumers have, adjusted for inflation and averaged over the industrial nations. System Price is the average sales price of systems that RGC sells, Capital Investment refers to capital investment in the business, and Advertising and Marketing and Sales are self-explanatory.

- a** Use the data in Table 10.8 to run a regression analysis with Sales as the response variable and the other four variables as explanatory variables. Write out the expression for Expected Sales, conditional on the other variables. Interpret the coefficients.
- b** The forecast for the Spending Index in the first half of 1996 is 45.2. The company has committed \$145,000 for capital improvements during this time. In order to make his goal of \$6 million in sales, Reilly has indicated that the company will offer discounts on systems (to the extent that the average system price could be as low as \$70,000) and will spend up to \$90,000 on advertising and various marketing programs. If the firm drops the price as much as possible and spends the full \$90,000 on Advertising and Marketing, what is the expected value for Sales over the next six months? Estimate the probability that Sales will exceed \$6 million.
- c** Given that the company spends the full \$90,000 on Advertising and Marketing, how low would Price have to be in order to ensure a 90% chance that Sales will exceed \$6 million? What advice would you give Terry?

- 10.15** Before performing the experiment, the scientist in Problem 10.9 thought a bit about the animals that he typically used in his lab. He knew (based on information from the animals' supplier) that the standard deviation of their weights was 1.5 grams. But the supplier was unable to specify the average weight precisely. The company did say that there

Table 10.8
 Twenty years of data
 for Reilly Global
 Communications, Inc.

Year	Spending Index	System Price (\$1000s)	Capital Investment (\$1000s)	Advertising and Marketing (\$1000s)	Sales (\$1000s)
1977	39.8	56.2	49.9	76.9	5540
	36.9	59.0	16.6	88.8	5439
1978	26.8	56.7	89.2	51.3	4290
	48.4	57.8	106.7	39.6	5502
1979	39.4	59.1	142.6	51.7	4872
	33.2	60.1	61.3	20.5	4708
1980	33.6	59.8	-30.4	40.2	4628
	38.3	60.1	-44.6	31.6	4110
1981	28.5	63.1	-28.4	12.5	4123
	27.7	62.3	75.7	68.3	4842
1982	45.6	64.9	144.0	52.5	5740
	35.5	64.9	112.9	76.7	5094
1983	36.4	63.4	128.3	96.1	5383
	32.0	65.6	10.1	48.0	4888
1984	31.1	67.0	-24.8	27.2	4033
	36.2	66.9	116.7	72.7	4942
1985	40.8	66.2	120.4	62.3	5313
	43.3	67.9	121.8	24.7	5140
1986		68.9	71.1	73.9	4397
	47.6	71.4	-4.2	63.3	5149
1987	41.5	69.3	-46.9	28.7	5151
	42.0	69.7	7.6	91.4	4989
1988	53.6	73.2	127.5	74.0	5927
	43.2	73.4	-49.6	16.2	4704
1989	43.6	73.1	100.1	43.0	5366
	41.5	74.9	-40.2	41.1	4630
1990	46.2	73.2	68.2	92.5	5712
	42.9	74.2	88.0	83.3	5095
1991	51.7	74.3	27.1	74.9	6124
	32.8	77.1	59.3	87.5	4787
1992	41.8	78.6	142.0	74.5	5036
	51.5	77.1	126.4	21.3	5288
1993	41.2	78.2	29.6	26.5	4647
	45.5	77.9	18.0	94.6	5316
1994	55.4	81.0	42.4	92.5	6180
	44.1	79.9	-21.6	50.0	4801
1995	41.7	80.6	148.4	83.2	5512
	46.1	82.3	-17.6	91.2	5272

was a 68% chance that the average weight was between 9.0 and 9.8 grams and a 95% chance that the average weight was between 8.6 and 10.2 grams.

a Use the stated probabilities to find a natural conjugate prior distribution for the average weight of the animals. What is $P(\mu \geq 10 \text{ grams} \mid m_0, \sigma_0)$?

b Find the posterior distribution for the average weight of the animals after having seen the data in Problem 10.9. What is $P(\mu \geq 10 \text{ grams} \mid m^*, \sigma^*)$?

10.16 Continuing Problem 10.15:

a Based on the prior probability distribution for μ , find the predictive probability that a single lab animal will weigh more than 11 grams, $P(x \geq 11 \text{ grams} \mid m_0, \sigma_0)$.

b After having seen the data and updating the beliefs about μ , find the predictive probability that a lab animal weighs more than 11 grams, $P(x \geq 11 \text{ grams} \mid m^*, \sigma^*)$?

10.17 The plant manager in Problem 10.10 actually began her investigation by assessing her subjective probability distribution for Q , the proportion of defective bulbs coming off the assembly line. Her assessed distribution was a beta distribution with parameters $r_0 = 1$, $n_0 = 20$; that is, her distribution was $f_\beta(q \mid r_0 = 1, n_0 = 20)$.

a Plot this prior distribution for Q .

b What would her posterior distribution for Q be after having observed the data in Problem 10.10?

10.18 A political analyst was interested in the proportion C of individuals who would vote for a controversial ballot measure. While he thought that it would be a close call, he was unsure of the precise value for C . He assessed a beta distribution for C with parameters $r_0 = 3$, $n_0 = 6$.

a Plot this prior distribution for C .

b The analyst is about to ask four individuals about their preferences. What is the probability that more than two of these individuals will express their support for the ballot measure?

c Having questioned the four individuals, the analyst found that three would indeed vote for the ballot measure. Find the analyst's posterior distribution for C . Plot this posterior distribution and compare it with the prior distribution plotted in part a.

d The analyst is now about to survey another 10 people. What is the probability that more than five of these people will support the ballot measure in this new poll?

e Suppose that 6 of the 10 people surveyed said they would vote for the ballot measure. The analyst now must write up his results. What is his probability that the ballot measure will pass?

10.19 A comptroller was preparing to analyze the distribution of balances in the various accounts receivable for her firm. She knew from studies in previous years that the distribution would be normal with a standard deviation of \$1500, but she was unsure of the mean μ . She thought carefully about her uncertainty about this parameter and assessed a normal distribution for μ with mean $m_0 = \$10,000$ and $\sigma_0 = \$800$.

Over lunch, she discussed this problem with her friend, who also worked in the accounting division. Her friend commented that she also was unsure of μ but would have placed it somewhat higher. The friend said that "better" estimates for m_0 and σ_0 would have been \$12,000 and \$750, respectively.

- a Find $P(\mu > \$11,000)$ for both prior distributions.
- b That afternoon, the comptroller randomly chose nine accounts and calculated $\bar{x} = \$11,003$. Find her posterior distribution for μ . Find the posterior distribution of M for her friend. Calculate $P(\mu > \$11,000)$ for each case.
- c A week later the analysis had been completed. Of a total of 144 accounts (including the nine reported in part b), the average was $\bar{x} = \$11,254$. Find the posterior distribution for μ for each of the two prior distributions. Calculate $P(M > \$11,000)$ for each case.
- d Discuss your answers to parts a, b, and c. What can you conclude?

CASE STUDIES

TACO SHELLS

Martin Ortiz, purchasing manager for the True Taco fast food chain, was contacted by a salesperson for a food service company. The salesperson pointed out the high breakage rate that was common in the shipment of most taco shells. Martin was aware of this fact, and noted that the chain usually experienced a 10% to 15% breakage rate. The salesperson then explained that his company recently had designed a new shipping container that reduced the breakage rate to less than 5%, and he produced the results of an independent test to support his claim.

When Martin asked about price, the salesperson said that his company charged \$25 for a case of 500 taco shells, \$1.25 more than True Taco currently was paying. But the salesperson claimed that the lower breakage rate more than compensated for the higher cost, offering a lower cost per usable taco shell than the current supplier. Martin, however, felt that he should try the new product on a limited basis and develop his own evidence. He decided to order a dozen cases and compare the breakage rate in these 12 cases with the next shipment of 18 cases from the current supplier. For each case received, Martin carefully counted the number of usable shells. The results are shown below:

New Supplier		Usable Shells		
		Current Supplier		
468	467	444	441	450
474	469	449	434	444
474	484	443	427	433
479	470	440	446	441
482	463	439	452	436
478	468	448	442	429

Questions

- 1 Martin Ortiz's problem appears to be which supplier to choose to achieve the lowest expected cost per usable taco shell. Draw a decision tree of the problem, assuming he orders one case of taco shells. Should you use continuous or discrete chance nodes to represent the number of usable taco shells in one case?
- 2 Develop CDFs for the number of usable shells in one case for each supplier. Compare these two CDFs. Which appears to have the highest expected number of usable shells? Which one is riskier?
- 3 Create discrete approximations of the CDFs found in Question 2. Use these approximations in your decision tree to determine which supplier should receive the contract.
- 4 Based on the sample data given, calculate the average number of usable tacos per case for each supplier. Use these sample means to calculate the cost per usable taco for each supplier. Are your results consistent with your answer to Question 3? Discuss the advantages of finding the CDFs as part of the solution to the decision problem.
- 5 Should Martin Ortiz account for anything else in deciding which supplier should receive the contract?

Source: This case was adapted from W. Mendenhall, J. Reinmuth, and R. Beaver (1989) *Statistics for Management and Economics*, 6th ed. Boston: PWS-KENT.

FORECASTING SALES

Sales documents were scattered all over Tim Hedge's desk. He had been asked to look at all of the available information and to try to forecast the number of microwave ovens that NewWave, Inc., would sell over the upcoming year. He had been with the company for only a month and had never worked for a microwave company before. In fact, he had never worked for a company that was involved in consumer electronics. He had been hired because the boss had been impressed with his ability to grasp and analyze a wide variety of different kinds of decision problems.

Tim had dug up as much information as he could, and one of the things he had found was that one of NewWave's salespeople, Al Morley, had kept detailed records over the past 14 years of his own annual forecasts of the number of microwaves that the company would sell. Over the years, Morley had been fairly accurate. (His sales performance had been pretty good, too.) Of course, Morley had been only too happy to provide his own forecast of sales for the upcoming year. Tim thought this was fine, and reported back to his boss, Bill Maught.

"Yes," said Bill after he had listened to Tim, "I am aware that Al has been making these forecasts over the years. We have never really kept track of his forecasts, though. Even though he claims to have been fairly accurate over the years, I can remember one or two when he was not very close. For example, I think last year he was off by about

6000 units. Of course, if you ask him about those cases, he can explain them in hindsight. But there is always the possibility of unusual unforeseen circumstances in any given year. Besides, I think he forecasts low. When he does that, his sales quota is set lower, and he's more likely to get a bonus. I'd say there's about a 95 percent chance that on average he underforecasts by 1000 units or more. In fact, I'd bet even money that his average forecast error [sales — forecast] is above 1700 units."

Tim asked the obvious question. "Would you like me to look into this more? Maybe I could get a handle on just how good he is."

"Fine with me," answered Bill. "Besides, if we really could get a handle on how good he is, then his forecast might be a good basis for us to work with each year."

After that discussion, Tim had visited the accounting department to find out the number of units actually sold in each year for which Morley had made a forecast. He now had assembled the data and considered the numbers:

Forecast (Units)	Actual Sales	Error (Sales — Forecast)
39,000	41,553	2553
44,000	46,223	2223
46,000	49,351	3351
54,000	55,393	1393
60,000	61,607	1607
59,000	68,835	9835
99,000	101,647	2647
124,000	123,573	-427
149,000	156,473	7473
145,000	146,333	1333
159,000	155,668	-3332
169,000	167,168	-1832
171,000	171,477	477
179,000	185,529	6529

- 1 Calculate the average and standard deviation of Morley's forecasting errors. What do you think about Bill Maught's opinion of Morley's forecasting?
- 2 Plot a CDF based on these data. Based on these data alone, what is the probability that Morley's forecast this year will be too low by 1700 units or more?
- 3 Assume that Morley's forecast errors follow a normal distribution with standard deviation $\sigma = 3000$. Translate Bill Maught's probability assessments into a prior distribution for Morley's average error, μ . On the basis of these data, what would be Maught's posterior distribution? What is his posterior probability that Morley's average error is greater than 1700 units?

- 4 Morley has forecast sales of 187,000 units for the coming year. Based on Maught's posterior distribution for λ (Morley's average error), what is the probability distribution for sales for this coming year? What is the probability that sales will be greater than 190,000 units? Then sketch a CDF for this distribution.

OVERBOOKING, PART II

Consider again the overbooking issue as discussed in the case study at the end of Chapter 9 (pages 345-346). Suppose the Mockingbird Airlines-operations manager believes that the no-show rate (N) is approximately 0.04 but is not exactly sure of this value. She assesses a probability distribution for N and finds that a beta distribution with $r_0 = 1$ and $n_0 = 15$ provides a good fit to her subjective beliefs.

Questions

- 1 Suppose that Mockingbird sells 17 reservations on the next flight. On the basis of this prior distribution, find the predictive probability that all 17 passengers will show up to claim their reservations. Find the probability that 16 will show up, 15, 14, and so on.
- 2 Should Mockingbird have overbooked for the next flight? If so, by how much? If not, why not? Support your answer with the necessary calculations.
- 3 Now suppose 17 reservations are sold, and 17 passengers show up. Find the operations manager's posterior distribution for N .
- 4 On the basis of this new information and the manager's posterior distribution, should Mockingbird overbook on the next flight? If so, by how much? If not, why not? Support your answers with the necessary calculations. [*Hint*: You may have to be careful in calculating the factorial terms. Many calculators and electronic spreadsheets will not have sufficient precision to do these calculations very well. One trick is to cancel out as many terms in the numerator and denominator as possible. Another is to do the calculation by alternately multiplying and dividing. For example, $5!/6! = (5/6)(4/5)(3/4)(2/3)(1/2)$.]

REFERENCES

Using data as the basis for probability modeling is a central issue in statistics. Some techniques that we have discussed, such as creating histograms, are basic tools; any basic statistics textbook will cover these topics. The text by Vatter et al. (1978) contains an excellent discussion of the construction of a data-based CDF. Fitting a theoretical distribution using sample statistics such as the mean and standard deviation is also a basic and commonly used technique (for example, Olkin, Gleser, and Derman 1980), although it is

worth noting that statisticians have many different mathematical techniques for fitting distributions to empirical data.

Virtually every introductory statistics text covers regression analysis at some level. More advanced texts include Chatterjee and Price (1977), Draper and Smith (1981), and Neter, Wasserman, and Kutner (1989). The treatment in Chapter 9 of Vatter et al. (1978) is similar to the one presented here.

Bayesian updating and the use of natural conjugate priors is a central element of Bayesian statistics. Winkler (1972) has the most readable fundamental discussion of these issues. For a more complete treatment, including discussion of many more probabilistic processes and analysis using natural conjugate priors, see Raiffa and Schlaifer (1961) or DeGroot (1970). Zellner (1971) discusses Bayesian regression analysis.

Berry, D. (1995). *Statistics: A Bayesian Perspective*. Belmont, CA: Duxbury. Chatterjee, S., and B. Price (1977) *Regression Analysis by Example*. New York: Wiley. DeGroot, M. (1970) *Optimal Statistical Decisions*. New York: McGraw-Hill. Draper, N., and H. Smith (1981) *Applied Regression Analysis*, 2nd ed. New York: Wiley.

Neter, T., W. Wasserman, and M. Kutner (1989). *Applied Linear Regression Models*, 2nd ed. Homewood, IL: Irwin.

Olkin, I., L. J. Gleser, and C. Derman (1980) *Probability Models and Applications*. New York: Macmillan.

Raiffa, H., and R. Schlaifer (1961) *Applied Statistical Decision Theory*. Cambridge, MA: Harvard University Press.

Vatter, P., S. Bradley, S. Frey, and B. Jackson (1978) *Quantitative Methods in Management: Text and Cases*. Homewood, IL: Irwin.

Winkler, Robert L. (1972) *Introduction to Bayesian Inference and Decision*. New York: Holt.

Zellner, Arnold (1971) *An Introduction to Bayesian Inference in Econometrics*. New York: Wiley.

Monte Carlo Simulation

The problems we have dealt with so far have allowed us to calculate expected values or to find probability distributions fairly easily. In real-world situations, however, many factors may be subject to some uncertainty. You can imagine what becomes of a decision tree that involves many uncertain events. The only way to prevent it from being a bushy mess is to present a skeleton version as in Figure 11.1, in which A and B represent alternative courses of action, each affected by many different uncertain quantities. We know that consideration of the uncertainties involved is important, but how will we deal with this much uncertainty? It is not at all clear that we will be able to use the techniques we have learned so far. We could, of course, painstakingly develop discrete approximations for all of the continuous distributions and construct the decision tree or influence diagram. In many cases this will work out fine. The decision tree, however, may become extremely complex.

Figure 11.1
Decision tree representing a complex decision situation with many sources of uncertainty.

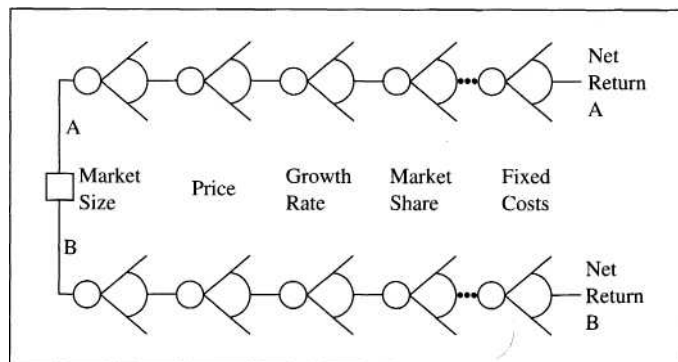
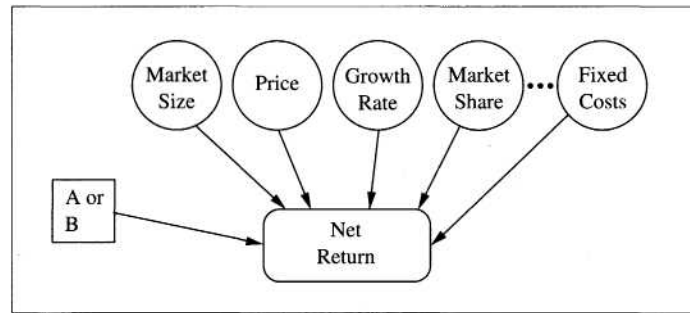


Figure 11.2
Influence diagram corresponding to decision tree in Figure 11.1
 Interdependences also may exist among the uncertainty nodes, thus complicating the influence diagram further.



An influence diagram can help somewhat in this kind of a situation. As a more compact representation of a decision problem, an influence diagram can provide a clear picture of the way that multiple sources of uncertainty affect the decision. Figure 11.2 is an influence diagram that corresponds to the skeleton decision tree in Figure 11.1. Careful structuring can lead to a complete influence diagram, especially if the probability distributions can be represented adequately by discrete distributions in the chance nodes. But if there are multiple interrelated uncertain quantities represented by complicated continuous distributions, then even an influence-diagram approach may not be adequate.

In many decision-analysis problems, the ultimate goal is the calculation of an expected value. If a complicated uncertainty model includes several continuous distributions with expected values that are easily found, then analyzing the problem may be relatively straightforward. But this is not always true. In the following realistic example, it may not be obvious how to calculate an expected value.

FASHIONS

Janet Dawes is in a quandary. She is the purchaser for a factory that produces fashion clothes. Her current task is to choose a supplier who can furnish fabric for a new line of garments. To ensure its supply for the upcoming year, her company must sign a contract with one of several textile suppliers. After a few inquiries, she has narrowed the choice to two specific suppliers. The first will supply as much fabric as she needs during the upcoming year for \$2.00 per yard and will guarantee supply through the year at this price.

The second supplier has a price schedule that depends on how much Janet orders over the next year. The first 20,000 yards will cost \$2.10 per yard. The next 10,000 yards will cost \$1.90 per yard. After this, the price drops to \$1.70 per yard for the next 10,000 yards, and then \$1.50 per yard for anything more than 40,000 yards.

After carefully considering the uncertainty about sales for the new garment line, Janet decides to model her uncertainty about the amount of fabric required over the next year as a normal distribution with mean 25,000 yards and standard deviation 5000 yards.

If the new line of garments is successful, Janet Dawes can save considerable money by going with the second supplier. On the other hand, if she signs the contract with Supplier 2 and the new line of garments does not prove successful, the cost of materials will have been higher than it might have been with Supplier 1, thus adding to the cost of an already expensive experiment. What should she do? Her decision tree is shown in Figure 11.3.

In Janet Dawes's problem, it is relatively easy to figure out the expected cost for Supplier 1: It is \$2.00 per yard times 25,000 yards, or \$50,000. But it is not so easy to figure out the expected cost for Supplier 2. To do this, we must know (a) the probability that X is within each interval and (b) the expected value of X for each interval. Part a is no problem; we could figure out the probabilities using the normal distribution as described in Chapter 9. It is not so clear, however, how to find the conditional expected values for X within each interval.

One approach to dealing with complicated uncertainty models such as that discussed above is through computer simulation. Think of the entire decision situation as an uncertain event, and "play the game" represented by the decision tree many times. In Janet Dawes's problem, we could imagine the computer making a random drawing from a normal distribution to find the amount of fabric required. Based on this specific amount, the computer then would calculate the total cost using the expressions given in the decision tree. After performing this procedure many times, it would be possible to draw a histogram or risk profile of the cost figures and thus to calculate the average cost. On the basis of these results, and comparing them with the distribution of costs for Supplier 1 (normal with mean \$50,000 and standard deviation \$10,000), a choice can be made.

For another example, imagine using a computer simulation to evaluate Alternative A in Figure 11.1. We would let the computer "flip a coin" to find the market size, then

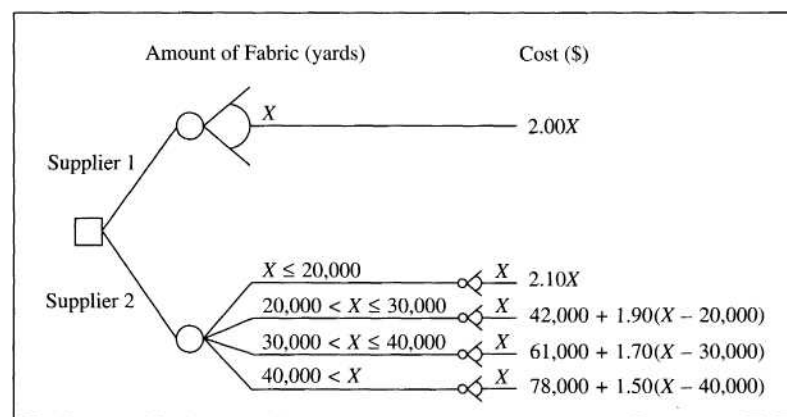


Figure 11.3 Decision tree for the fabric buyer. The expressions at the ends of the branches for Supplier 2 include the discounts. For example, if X is between 30,000 and 40,000 yards, the total cost is \$42,000 for the first 20,000 yards, plus \$19,000 for the next 10,000 yards, plus \$1.70 per yard over 30,000.

flip another to find the price, another to find the growth rate, and so on, until all uncertain quantities are chosen, each according to its own probability distribution. Of course, having the computer flip a coin means having it choose randomly the market size, price, growth rate, and so on. Once all of the necessary values had been determined in this random fashion, the computer then could calculate the net return. We then would repeat the entire procedure many times and track the results. At the end, it would be possible to graph the distribution of net returns and to examine descriptive statistics such as the mean, standard deviation, and probability of a negative return.

The term *Monte Carlo simulation* often is used to refer to this process because, as in gambling, the eventual results depend on random selections of values. The basic idea is straightforward: If the computer "plays the game" long enough, we will have a very good idea of the distribution of the possible results. Although the concept is easy, its implementation requires attention to many details.

We can use Monte Carlo simulation to cope with situations in which uncertainty abounds; the objective is to represent the uncertainty surrounding the possible payoffs for the different alternatives. How does simulation fit into our perspective of building models to represent decision situations? When we put together all probability distributions for all uncertain quantities, we are building a simulation model that we believe captures the relevant aspects of the uncertainty in the problem. After running the simulation many times, we have an approximation of the probability distribution for the payoffs from the different alternatives. The more simulations we can do, the more accurate that approximation. Finally, the results, both risk profiles and average outcomes, can be used in the decision analysis to make an appropriate decision.

For the main example in this chapter, we will continue the soft pretzel problem from Chapter 9. You are thinking about manufacturing and marketing soft pretzels using a new recipe, but you face many sources of uncertainty. In the process of analyzing the decision, you might ask whether choosing to go into the soft pretzel business would have a positive expected value. By using the Monte Carlo simulation, we can examine the probability distribution of returns that would be associated with your pretzel business.

Let us suppose that the market size is unknown, but your subjective beliefs about the market's size can be represented as a normal random variable with mean 100,000 pretzels and standard deviation 10,000. The proportion (P) of the market that you will capture is unknown and, strictly speaking, is a continuous random variable that could range anywhere between 0 and 100%. You have decided, however, that your beliefs can be modeled adequately with the following discrete distribution:

Proportion (%)	Probability
16	0.15
19	0.35
25	0.35
28	0.15

Your pretzel's selling price is known to be \$0.50. Variable costs are a uniform random variable between \$0.08 and \$0.12 per pretzel. Fixed costs also are random, with a distribution we will describe later.

Putting all of the pieces together to calculate net contribution yields

$$\text{Net Contribution} = (\text{Size} \times P/100) \times (\text{Price} - \text{Variable Cost}) - \text{Fixed Cost}$$


From this equation you can see the overall strategy; for each iteration, values for Size, P , Variable Cost, and Fixed Cost will be chosen from their respective distributions and then combined to get a figure for Net Contribution. Doing so for many iterations will yield a distribution for Net Contribution.

Using Uniform Random Numbers as Building Blocks

If we are going to construct random numbers from various different probability distributions, we will need a place to start. The starting point is a random variable that is uniformly distributed between 0 and 1. That is, all values between 0 and 1 are possible and equally likely. (See Problems 9.27-9.29.) Let x denote the number generated from such a distribution, and suppose we require a random number y^* from the distribution (CDF) $F(y) = P(F \leq y)$. An example CDF is plotted in Figure 11.4. Now generate a uniform x^* between 0 and 1. Locate that x^* on the vertical axis. Read across from x^* to the CDF and down to the horizontal axis. The number on the horizontal is the required y^* .

This approach is straightforward and intuitive; we first generate a uniform random number and then work backward through the CDF to get y . In principle, we can do this for any distribution. Unfortunately, it is not possible in some cases because of the mathematical form of $F(y)$. In these cases, other methods can be used.

Figure 11.4
Generating a random number y^* on the basis of a uniform random number x^* .



General Uniform Distributions

Suppose that we want to generate a uniform random number between 1 and 2 instead of between 0 and 1. The solution would be simply to calculate

$$y = x + 1$$

Then y would be uniformly distributed between 1 and 2. If we wanted a random number between 0 and 2, we simply multiply x by 2. This would "stretch out" the uniform distribution so that it would cover the interval from 0 to 2.

In general, suppose you want a uniform random number between a and b ($a < b$). The procedure is to generate x , then move it to the right place and stretch out the distribution to cover the desired interval. The formula is

$$y = a + x(b - a)$$


Multiplying x by $(b - a)$ stretches out the distribution, and adding a moves the distribution to the right place. If x turns out to be 0, then $y = a$, and if x is 1, then $y = b$.

In our soft pretzel example, variable cost is uniformly distributed between \$0.08 and \$0.12. Thus, the calculation is

$$\text{Variable Cost} = 0.08 + x(0.12 - 0.08)$$

This procedure for generating general uniform random numbers, which we developed intuitively, is fully compatible with the general simulation approach described above. The CDF for a uniform distribution between a and b is a straight line as shown in Figure 11.5. The line has the equation $F(y) = (y - a)/(b - a)$. (Can you verify this?) To work backward through this CDF, we set the uniform x^* equal to $F(y^*) = (y^* - a)/(b - a)$ and solve for y^* . That is, we need to know what y^*

Figure 11.5
Generating a uniform
 y^* between a and b .



corresponds to our uniform x^* between 0 and 1. Solving for y^* , we obtain $y^* = a + x^*(b - a)$.

Exponential Distributions

Another continuous random variable that is easy to simulate using our general approach is an exponential random variable T . Suppose you want to generate a random number that is exponentially distributed with rate m . The CDF for this distribution is $F(t) = 1 - e^{-tm}$. Thus, we generate a uniform x^* and set $x^* = F(t^*) = 1 - e^{-tm}$. Solving for t^* gives

$$t^* = \frac{-\ln(1 - x^*)}{m}$$


Figure 11.6 shows the CDF and the graphical equivalent of solving for t^* .

The formula derived above can be programmed easily into an electronic spreadsheet. But we can do a little bit better. If x^* is from a uniform distribution between 0 and 1, then so is $1 - x^*$. Thus, we can simplify the procedure by calculating

$$t' = \frac{-\ln(x^*)}{m}$$

The random number t' also has an exponential distribution with rate m . The reason for the simplification is to avoid unnecessary calculations that would slow down the computer simulation.

Figure 11.6
Generating an exponential random number.



Discrete Distributions

Discrete distributions also are easy to handle in a way that is similar to using the CDF. Suppose, for example, that we want to simulate flips of a fair coin. We want the probability of a head to be 0.50. First we generate x . Then if $x \leq 0.50$, we will say that we have a head, and if $x > 0.50$, we will have a tail. Because x is from a uniform distribution, it has a 50% chance of being between 0 and 0.50 and a 50% chance of being between 0.50 and 1.

The basic strategy is to split the interval from 0 to 1 into smaller intervals corresponding to the outcomes of the discrete event. The outcome of the event then corresponds to the interval into which x falls when it is generated. The probabilities are controlled by the widths of the intervals for x ; smaller intervals have smaller probabilities, and larger intervals have larger probabilities.

In our soft pretzel example, we need to generate the market proportion captured (P) using discrete probabilities. The four intervals we need are as follows:

- If $x \leq 0.15$, then $P = 16\%$.
- If $0.15 < x \leq 0.50$, then $P = 19\%$.
- If $0.50 < x \leq 0.85$, then $P = 25\%$.
- If $0.85 < x$, then $P = 28\%$.

Generating x , determining the interval in which it falls, and choosing P accordingly will generate discrete values of P with the appropriate probabilities.

Other Distributions

As you know, there are many other distributions that one might wish to use in a simulation. The procedure we have described above can be used for many of these distributions, but there are other approaches as well. For example, a normal random variable can be created by adding up a number of independent uniform random variables. Another approach to generating normal random variables uses a pair of independent uniforms and, with complex transformations, produces a pair of independent normals. For more information on simulating random variables, see Law and Kelton (1991).

Computer Software for Simulation

Monte Carlo simulation has long been a useful tool for management scientists, and specialized programs for performing complex simulations have been available for some time. The recent development of simulation software that runs in conjunction with electronic spreadsheets, though, has put powerful simulation tools in the hands of anyone

with a personal computer. Here we describe two such programs, Crystal Ball and @RISK, and show how to analyze the soft pretzel problem using each of them.

Both Crystal Ball and @RISK run as "add-ins" in Lotus 1-2-3 and Microsoft Excel spreadsheets, which make them exceptionally easy to use. The illustrations in this chapter use Excel as the underlying spreadsheet program, and Figure 11.7 shows the Excel spreadsheet for the soft pretzel problem. To start, typical values have been placed in the cells for Market Size, Market Proportion, Variable Cost, and Fixed Cost. Recall that Price is fixed at \$0.50 per pretzel.

With the spreadsheet set up, we are ready to put the simulation software to work. In both Crystal Ball and @RISK, performing a simulation requires four steps: defining the input distributions, identifying the outputs for which we want distributions, running the simulation, and analyzing the results.

CRYSTAL BALL

To assign probability distributions to inputs using Crystal Ball, the analyst selects a cell and chooses a particular distribution to assign. Crystal Ball provides 16 different families of distributions, including a do-it-yourself custom distribution that can be used to create a distribution of virtually any shape. In Figure 11.7, a normal distribution with mean 100,000 and standard deviation 10,000 has been assigned to the Market Size Cell (B5), and a uniform distribution from \$0.08 to \$0.12 has been assigned to Variable Cost in cell B8. A discrete distribution, using the custom distribution option with the values and probabilities on the right side of the spreadsheet, has been assigned to cell B6 for Market Proportion.

For Fixed Cost, we will use a distribution called a *triangular distribution* (see Problem 9.30). A triangular distribution is defined by three points: the minimum, the maximum, and the most likely. Figure 11.8 shows the window from Crystal Ball in

Figure 11.7
The soft pretzel
spreadsheet.

Soft Pretzel Simulation						
	A	B	C	D	E	F
1						
2						
3					Proportion	Probability
4	ASSUMPTIONS:			Market	16	0.15
5	Market Size	100,000		Proportion	19	0.35
6	Market Proportion	22		Probabilities	25	0.35
7	Price	\$0.50			28	0.15
8	Variable Cost	\$0.10				
9	Fixed Cost	\$8,000				
10						
11						
12	FORECAST:					
13	Net Contribution	\$800				
14						

which a triangular distribution is being defined for Fixed Cost. We have indicated that Fixed Cost could be as low as \$6500, as high as \$9000, but the most likely value is \$8000.

With all four probability distributions assigned, the next step is to indicate that the output (called a *forecast* by Crystal Ball) desired is Net Contribution. At this point, the simulation is ready to run. Figure 11.9 shows the histogram of Net

Figure 11.8
Defining a triangular distribution in Crystal Ball





Figure 11.9
Crystal Ball's results from 1000 iterations of the soft pretzel model.



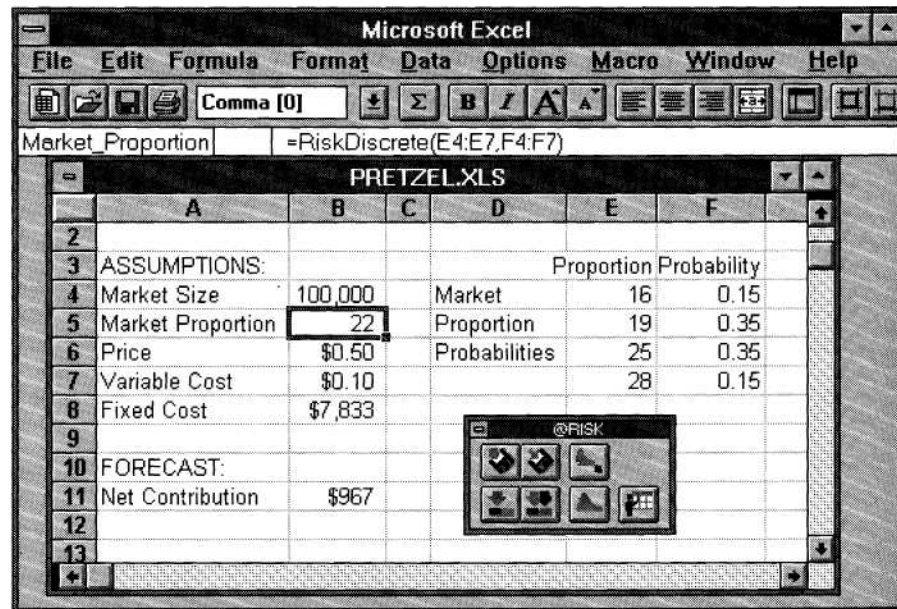
Contribution that resulted from 1000 iterations of the simulation, indicating a 36% chance that the soft pretzel venture will end up losing money. The results can also be viewed as a CDF, and a full complement of statistics about the distribution are available, including the ability to find the probability that Net Contribution falls into any specified interval.

@RISK

Assigning probability distributions to input cells in @RISK is no different than building a standard spreadsheet formula in Excel; @RISK's special random-number generators are entered into cell formulas just like any other Excel function. @RISK has many built-in functions that generate random variables according to specified distributions. For example, Figure 11.10 shows the assignment of the discrete distribution for Market Proportion.

After entering the distribution functions and indicating that Net Contribution is the output that we wish to track, we perform the simulation. A useful feature of @RISK is the ability to monitor the rate at which the output distribution is changing

Figure 11.10
Assigning a discrete distribution with an @RISK function. The formula indicates that the values and probabilities are to be found in cells E4 through F7.




and to stop the simulation when the changes are small enough (according to user-specified parameters). After the simulation is finished, @RISK generates graphs similar to those of Crystal Ball. Figure 11.11, for example, shows the CDF for Net Contribution after 1000 iterations. As with Crystal Ball, the program produces a complete array of statistics on the output distribution.

The figures and discussion above show the basic process of running simulations in Crystal Ball and @RISK. Of course, if you try this problem on your own computer with one of these programs (and you should!), you will probably obtain slightly different results; your computer will probably choose different random values for the variables.

How do we put the results from a Monte Carlo Simulation to work? If we were simply interested in comparing the expected values of two projects, we might compare the averages of the simulated results. The point of the simulation, however, has been to obtain an idea of the uncertainty that surrounds the ultimate payoff. Thus, the resulting histograms and CDFs, which are no more than risk profiles, can be used as discussed in Chapter 4. For example, visually comparing the graphs associated with two different risky alternatives can reveal the extent to which one project is riskier than another, even though they may have similar expected values. You may be able to decide between the two projects simply on the basis of comparing the graphs or by considering $P(\text{Profit} < 0)$, say. Finally, a discrete approximation of the distribution could be placed into a decision tree for a rough approximation of the distribution. This strategy might be particularly appropriate if a simulation has been done as an auxiliary analysis to obtain a probability model for part of a complicated decision model.

Figure 11.11
@RISK's results from
1000 iterations of the
soft pretzel model.



Simulation and Sensitivity Analysis

Imagine this scenario. You have constructed your simulation model, and as you show it to a colleague, the assumptions that you have made are called into question. Suppose your colleague wonders what the effect would be if your normal distribution for Market Size had a mean of 120,000 instead of 100,000. Or what about charging a price other than \$0.50? Of course, it is a relatively simple matter to make a change in the spreadsheet or the assumptions and rerun the simulation. But both @RISK and Crystal Ball have some built-in sensitivity-analysis features.

Let us take up the issue of sensitivity to price. It would be possible, of course, to charge a different price, and reasonable prices might range from \$0.30 up to \$0.70. Of course, charging a higher price is better, everything else being equal. Unfortunately, everything else will not be equal; you can anticipate less sales as your price increases. In our model, this means that the Market Proportion will decrease. The question is whether there is a Price-Market Proportion combination that is best.

We need a simple way to model the dependence between Price and Market Proportion. In Figure 11.12, the original spreadsheet has been modified to show how the Market Proportion changes (up or down) with changes in Price. You can see that each \$0.10 change in price implies a five-point change in the Market Proportion. Net Contribution is calculated for each of the Price scenarios in cells D14-D18. You can see that \$0.50 appears to be the best price. But this is just for the values in the spreadsheet. What happens when we run the simulation?

Figure 11.12
Spreadsheet for price
sensitivity analysis.

Soft Pretzel Trend						
	A	B	C	D	E	F
1						
2						
3					Proportion	Probability
4	ASSUMPTIONS:			Market	16	0.15
5	Market Size	100,000		Proportion	19	0.35
6	Market Proportion	22		Probabilities	25	0.35
7	Price	\$0.50			28	0.15
8	Variable Cost	\$0.10				
9	Fixed Cost	\$8,000				
10						
11		Change in				
12		Market				
13	Price	Proportion		Net contribution		
14	\$0.30	10		-\$1,600		
15	\$0.40	5		\$100		
16	\$0.50	0		\$800		
17	\$0.60	-5		\$500		
18	\$0.70	-10		-\$800		
19						

Figure 11.13
Trend chart showing
changes in risk profile
as price changes.


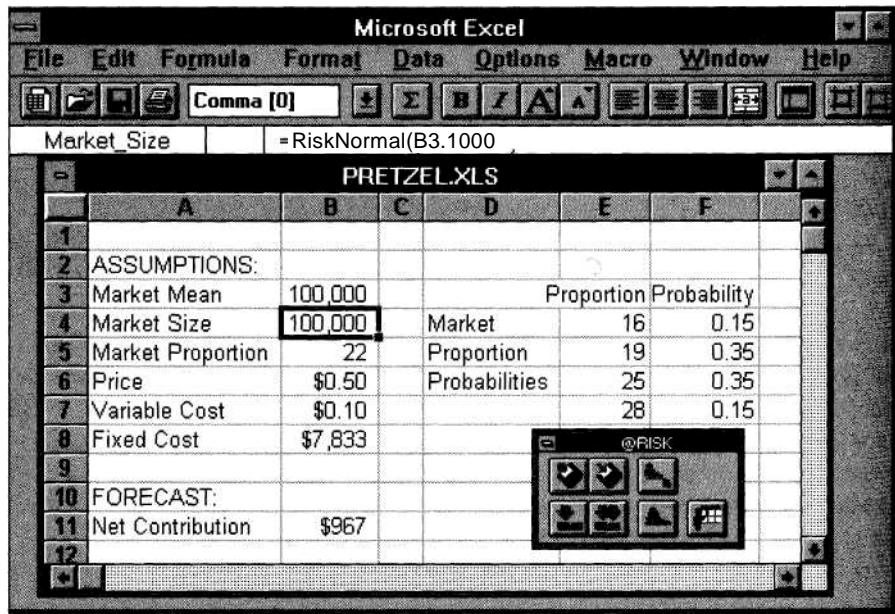


Figure 11.13 is a Trend Chart from Crystal Ball. (@RISK can produce a similar graph.) It shows how the risk profiles change as price changes. You can see clearly the change in the location of the distribution as the price increases. We begin with a distribution showing mostly losses when Price is \$0.30; this price is simply too low. The median profit increases but then decreases with the higher prices as Market Proportion decreases. In addition, the distribution becomes broader as price increases, the reason being that the higher price amplifies the uncertainties in Market Size and Market Proportion. It appears that a price around \$0.50 is the best choice, based on this model.

Distributions on Parameters (Optional)

Suppose in the soft pretzel problem that you cannot settle on a single value for the mean of the distribution for the market size. Your uncertainty about this parameter can also be modeled with a probability distribution. This is similar to what we did in Chapter 10 in our discussion of natural conjugate prior distributions for parameters, but here there is no need to restrict ourselves to specific kinds of distributions. In a sense we will be placing a distribution on another distribution; such models are sometimes called *hierarchical models*, and the uncertainty about the distribution parameter is sometimes called *second-order uncertainty*.

Figure 11.14
Incorporating uncertainty about the mean of the distribution for market size.



After carefully considering your beliefs about the mean of the market-size distribution, let us suppose that you decide to model your uncertainty as a uniform random variable between 90,000 and 110,000 pretzels. How would this change the model? Making the modification is straightforward; this uncertain mean simply becomes a new random variable in the simulation. Figure 11.14 illustrates this with @RISK, showing that the formula for Market Size in cell B4 now refers to cell B3, which in turn contains an @RISK function to generate a uniform random number between 90,000 and 110,000. On each iteration @RISK generates a new mean for the market-size distribution in cell B3 and then a specific market size from that distribution in cell B4.

Simulation, Decision Trees, and Influence Diagrams

Clearly, Monte Carlo simulation provides another modeling tool for decision analysis. Indeed, simulation is an important tool, and probably will become more widely used because of the ease with which small simulations can be performed in a spreadsheet environment. Because of this, it is worthwhile to spend a bit of time thinking about how simulation relates to the other modeling tools in decision analysis.

Simulation is an excellent tool for developing a model of uncertainty. We have used it here to develop risk profiles for decision alternatives; with risk profiles for different alternatives in hand, a decision maker would have some basis in choosing

one alternative over the other. As mentioned above, however, simulation also could be used as a subsidiary modeling tool to construct a probability model for a particularly messy part of a problem. For example, if a policy maker is attempting to evaluate alternatives regarding chemical spills, he or she may ask an analyst to develop a probability model for accidents that lead to spills. Such a model can be developed in the context of a simulation, and once an appropriate probability distribution is constructed, it can be used within a larger analysis.

The ease with which simulation can be performed, along with the flexibility of the simulation environment, makes it an attractive analytical tool. But this ease of use and flexibility does not mean that the decision maker can get away with less effort. In fact, subtle issues in simulation require careful thought. For example, we typically build the simulation in such a way that many of the random numbers are independent draws from their respective distributions. And the analyst may even model uncertainty about parameters through the specification of distributions on those parameters, but still have the parameters independent from one another. But would they be? If the analyst has been optimistic in assessing one parameter, perhaps other estimates have been subject to the same optimism.

The problem of modeling dependence among random variables in a simulation is not trivial and can have a substantial impact on the output distributions. Both @RISK and Crystal Ball have ways to incorporate dependence into their models.

A somewhat less subtle issue in simulation modeling is that an analyst may be tempted to include all possible sources of uncertainty in the model. This is relatively easy to do, after all, and somewhat more effort is required to do the sensitivity analysis to determine whether an uncertain quantity really matters in the outcome of the model. Hence, there may be a tendency with simulation not to take certain analysis steps that can lead to real insights as to what matters in a model, what issues should be considered more fully, or what uncertainties may demand more attention, in either more careful assessment or information acquisition.

For example, recall the Eagle Airlines case that we studied in detail in Chapter 5. Recall that Dick Carothers was trying to decide whether to purchase a new airplane for his business, and the decision was complicated by uncertainty about many different variables. In Chapter 5 we showed via sensitivity analysis that the variables which had the greatest impact on Profit were Operating Cost, Hours Flown, Capacity, and Charter Price. A simple simulation approach might place triangular distributions on all of the variables in the problem (see Table 5.1), with the base value being the most likely and the low and high values determining the upper and lower bounds of the distributions. Running this simulation with all variables uncertain yields results that are not very different from a simulation in which only the four most sensitive variables are specified as uncertain with the same triangular distribution. Thus, it is clear that these four variables drive the analysis; doing the sensitivity analysis first generates important insights about which variables are the most important in this problem.

In short, constructing a Monte Carlo simulation model requires the same careful thought that is required in any decision modeling. Simulation does have its own advantages (flexibility and ease of use) as well as disadvantages (rampant independence

assumptions and a tendency to solve problems with brute force) and hence, to some extent, leads to some special problems. But the same is true of simulation that is true of decision modeling in general. The decision maker and the decision analyst still are required to think clearly about the problem at hand and to be sure that the decision model addresses the important issues appropriately. Clear thinking is the key, not fancy quantitative modeling. The objective with any decision-analysis tool is to arrive at a requisite model of the decision, one that appropriately addresses all essential elements of the decision problem. It is through the process of constructing a requisite model, which includes careful thought about the issues, that the decision maker will gain insight and understanding about the problem.

SUMMARY

As we have seen in this chapter, Monte Carlo simulation is another approach to dealing with uncertainty in a decision situation. The basic approach is to construct a model that captures all of the relevant aspects of the uncertainty, and then to translate this model into a form that a computer can use; we focused on the development of such models within the environment of electronic spreadsheets and the use of @RISK and Crystal Ball to perform the simulation and analyze the results. We discussed how sensitivity analysis and simulation can be used together, the possibility of including second-order uncertainty for distribution parameters, and the role of simulation in creating a requisite decision-analysis model.


EXERCISES

- 11.1 Explain in your own words how Monte Carlo simulation may be useful to a decision maker.
- 11.2 Explain how the simulation process works to produce results that are helpful to a decision maker.
- 11.3 A simulation model has produced the three cumulative risk profiles displayed in Figure 11.15. What advice would you give a decision maker on the basis of this output?
- 11.4 A friend of yours has just learned about Monte Carlo simulation methods and has asked you to do a simulation of a complicated decision problem to help her make a choice. She would be happy to have you solve the problem and then recommend what action she should take. Explain why she needs to be involved in the simulation modeling process and what kind of information you need from her.

QUESTIONS AND PROBLEMS

- 11.5 Find the expected cost for Supplier 2 in Janet Dawes's purchasing problem as diagrammed in Figure 11.3.

Figure 11.15
Three cumulative risk profiles from a simulation.



- 11.6** Simulation is one way to find an expected value for Janet Dawes's problem as diagrammed in Figure 11.3. How could you construct a discrete approximation that would at least provide an approximate expected cost for Supplier 2?
- 11.7** What other real-world situations involve step functions like the one that Janet Dawes faces?
- 11.8** An investor has purchased a call option for 100 shares of Alligator stock and intends to hold it until the day the option expires, at which time he will sell it if he can. The option is worth nothing on its expiration date unless the price of Alligator stock is more than \$45 per share. For values of the stock greater than \$45, the option will be worth $100(\text{Share Price} - \$45)$. The reasoning behind this value is that the call option permits the option's owner to purchase 100 shares at \$45 per share. Thus, if the share price is greater than \$45, then the option owner could buy the shares and immediately resell them at the market price, pocketing the difference. Of course, the investor is uncertain about Alligator's eventual share price on the exercise date, but his uncertainty can be modeled using a normal distribution having mean \$45.50 and standard deviation \$5.00. Construct a simulation to estimate the expected value of the option on the exercise date.
- 11.9** Your boss has asked you to work up a simulation model to examine the uncertainty regarding the success or failure of five different investment projects. He provides probabilities for the success of each project individually: $p_1 = 0.50$, $p_2 = 0.35$, $p_3 = 0.65$, $p_4 = 0.58$, $p_5 = 0.45$. Because the projects are run by different people in different segments of the investment market, you both agree that it is reasonable to believe that, given these probabilities, the outcomes of the projects are independent. He points out, however, that he really is not fully confident in these probabilities and that he could be off by as much as 0.05 in either direction on any given probability.
- a** How can you incorporate his uncertainty about the probabilities into your simulation?
- b** Now suppose he says that if he is optimistic about the success of one project, he is likely to be optimistic about the others as well. For your simulation, this means that if one of the probabilities increases, the others also are likely to increase. How might you incorporate this information into your simulation?

- 11.10** A decision maker is working on a problem that requires her to study the uncertainty surrounding the payoff of an investment. There are three possible levels of payoff — \$1000, \$5000, and \$10,000. As a rough approximation, the decision maker believes that each possible payoff is equally likely. But she is not fully comfortable with the assessment that each probability is exactly $1/3$, and so would like to conduct a sensitivity analysis. In fact, she believes that each probability could range from 0 to $1/2$.
- a Show how a Monte Carlo simulation could facilitate a sensitivity analysis of the probabilities of the payoffs.
 - b Suppose the decision maker is willing to say that each of the three probabilities could be chosen from a uniform distribution between 0 and 1. Could you incorporate this information into your simulation? If so, how? If not, explain why not, or what additional information you would need.
- 11.11** Perform a simulation of the Eagle Airlines problem:
- a Treat all of the variables in Table 5.1 as uncertain, and create triangular distributions for each of them in @RISK or Crystal Ball. For each variable, use the base value as the most likely value, and the high and low values as the upper and lower bounds of the triangular distribution. Run the simulation to obtain a risk profile for Profit. What are the 0.05, 0.25, 0.50, 0.75, and 0.95 fractiles of this risk profile?
 - b Now eliminate the uncertainty on all variables *except* Operating Cost, Hours Flown, Capacity, and Charter Price, leaving all of the other variables at their base values. Now rerun the simulation. Obtain the risk profile for Profit, being sure to note the 0.05, 0.25, 0.50, 0.75, and 0.95 fractiles. Compare the risk profile obtained from this restricted model with the results from part a.
 - c In Chapter 5, we indicated that Capacity and Hours Flown were probabilistically related. How could you model this in your simulation? If you are using @RISK or Crystal Ball, specify positive correlation (about 0.7) for these two variables, rerun the model, and compare the results with those from parts a and b.

C A S E S T U D I E S

CHOOSING A MANUFACTURING PROCESS

AJS, Ltd., is a manufacturing company that performs contract work for a wide variety of firms. It primarily manufactures and assembles metal items, and so most of its equipment is designed for precision machining tasks. The executives of AJS currently are trying to decide between two processes for manufacturing a product. Their main criterion for measuring the value of a manufacturing process is net present value (NPV). The contractor will pay AJS \$8 per unit. AJS is using a three-year horizon for its evaluation (the current year and the next two years).

Process 1

Under the first process, AJS's current machinery is used to make the product. The following inputs are used:

Demand Demand for each of the three years is unknown. These three quantities are modeled as discrete random variables denoted D_0 , D_1 , and D_2 with the following probability distributions:

D_0	$P(D_0)$	D_1	$P(D_1)$	D_2	$P(D_2)$
11K	0.2	8K	0.2	4K	0.1
16K	0.6	19K	0.4	21K	0.5
21K	0.2	27K	0.4	37K	0.4

Variable Cost Variable cost per unit changes each year, depending on the costs for materials and labor. Let V_0 , V_1 , and V_2 represent the three variable costs. The uncertainty surrounding each variable is represented by a normal distribution with mean \$4 and standard deviation \$0.40.

Machine Failure Each year, AJS's machines fail occasionally, but obviously it is impossible to predict when or how many failures will occur during the year. Each time a machine fails, it costs the firm \$8000. Let Z_0 , Z_1 , and Z_2 represent the number of machine failures in each of the three years, and assume that each is a Poisson random variable with parameter $m = 4$.

Fixed Cost Each year a fixed cost of \$12,000 is incurred.

Process 2

The second process involves scrapping the current equipment (it has no salvage value) and purchasing new equipment to make the product at a cost of \$60,000. Assume that the firm pays cash for the new machine, and ignore tax effects.

Demand Because of the new machine, the final product is slightly altered and improved, and consequently the demands are likely to be higher than before, although more uncertain. The new demand distributions are:

D_0	$P(D_0)$	D_1	$P(D_1)$	D_2	$P(D_2)$
14K	0.3	12K	0.36	9K	0.4
19K	0.4	23K	0.36	26K	0.1
24K	0.3	31K	0.28	42K	0.5

Variable Cost Variable cost still changes each year, but this time V_0 , V_1 , and V_2 are each judged to be normal with mean \$3.50 and standard deviation \$1.00.

Machine Failures Equipment failures are less likely with the new equipment, occurring each year according to a Poisson distribution with parameter $m = 3$. They also tend to be less serious, costing only \$6000.

Fixed Cost The fixed cost of \$ 12,000 is unchanged.

Questions

- 1 Draw an influence diagram for this decision problem. Do you think it would be feasible to solve this problem with an influence diagram? Explain.
- 2 Write out the formula for the NPV for both processes described above. Use the variable names as specified, and assume a 10% interest rate.
- 3 For Process 1, construct a model and perform 1000 simulation trials. Estimate the mean and standard deviation of NPV for this process. Print a histogram of the results, and estimate the probability of a negative NPV occurring.
- 4 Repeat Question 3 for Process 2.
- 5 Compare the distribution of NPV for each of the two alternatives. Which process would be better for AJS? Why?

Source: This case was provided by Tom Mc Williams.

ORGANIC FARMING

Jane Keller surveyed the freshly plowed field on her farm. She and her husband, Tim, had taken over the farm from her parents 10 years ago. Since that time, she and Tim had worked hard to improve the profitability of the farm. Even though the work was hard, the lifestyle was rewarding. And she found that common sense combined with basics from some business courses had helped her in making difficult decisions.

She faced one such decision now. Over the years, she had noticed more and more of her neighbors adopting a variety of organic farming methods. Many of the techniques were easy to adopt and made good sense. For example, companion planting and promoting a balanced ecology on the farm helped to create an environment in which plants were less susceptible to insect and disease damage. Last year she had grown some produce using only organic methods and had sold it locally. She learned a lot on that small-scale project, and in particular she learned that she did not know much about natural methods for preventing specific diseases and controlling insects.

Still, Jane was intrigued by the possibility of organic farming. Growers who could label their produce as "Organically Grown" commanded premium prices from specialty stores and at the local farmer's market. The Organic Farmers Association

(OFA) provided the necessary certification. Meeting OFA's requirements involved (1) documenting the use of the land over a period of years to ensure little or no contamination from nonorganic pesticides, herbicides, and fertilizers; and (2) adhering to the farming methods that OFA deemed "organic."

This year Tim wanted to expand the operation by planting in a new field that had not been farmed in 20 years. Because the new field would be some distance from the other planted areas, it would be an ideal location to grow organic produce. There would be no problem getting the new field certified by the OFA. When she made this suggestion to Tim, he agreed with the idea in principle but wanted to think seriously about the project. They had several questions to answer. Would they really be able to make money from such a project? What kinds of prices could they anticipate for organic produce? What risks were involved?

Jane agreed to do some research. After visiting with her neighbors and spending a lot of time at the local organic gardening store, she was beginning to develop a plan. Although she could plant a large variety of different vegetables and herbs, most of the area would be devoted to tomatoes, green beans, and potatoes. She and Tim had plenty of experience with these crops, and she felt most comfortable experimenting with them. They would plant enough area so that, if the weather cooperated, they could expect to harvest 250 bushels of each crop. She was very uncertain about the exact yield, however, because of the variety of possible diseases and insects that could cause trouble and because of her own lack of experience with organic methods. For the tomatoes, she judged that there was a 68% chance that the yield would be between 235 and 265 bushels, and she was "almost sure" (95% chance) that the yield would be between 220 and 280 bushels. The green beans and potatoes she judged were somewhat less sensitive; for both of these crops, she estimated a 68% chance that the yield would fall between 242 and 258 bushels, and she was 95% sure that the yield would fall between 234 and 266 bushels.

The uncertain yield was complicated further by the effects of weather. Jane knew that the weather could be either too dry or too rainy for the crops, and that this would reduce her yield somewhat. At the same time, adverse weather would affect all growers in her region, leading to a smaller supply of produce and hence higher prices. The worst possible scenario for the Kellers would be for the weather to be perfect, resulting in a bumper crop around the region, while the Keller's venture into organics resulted in a low yield because of insects, disease, and their own inexperience. Prices would be low and they would have relatively little to sell.

Jane realized that she was facing a difficult judgmental task; ideally she would have to assess her uncertainty for the weather, the yield for each crop under different weather conditions, and the possible prices under the various conditions. To simplify matters, she decided to think about two scenarios. Under the first scenario, the weather would not affect the yield adversely. Under these conditions, she judged that the bushel price of tomatoes could range from \$5.00 to \$5.80. For potatoes, the range was \$4.15 to \$4.60; and for green beans, between \$5.90 and \$6.80. In each case, she figured that all prices in the specified range were equally likely.

Under the second scenario, the weather would have adverse effects on the region's crops. Jane estimated a probability of 0.15 for adverse weather. Because

tomatoes were the most sensitive in this regard, the crop size would be reduced by some 20%. With smaller yields in the region, prices would be higher, ranging from \$5.50 to \$6.00 per bushel. Potatoes were the least sensitive, with only a 7% crop-size reduction and prices between \$4.50 and \$4.80. Green beans were not terribly sensitive to the weather in terms of quantity, so Jane estimated a 4% reduction. Under adverse weather conditions, however, the quality of the beans could be highly variable, and so she estimated their price range under this scenario to be between \$5.50 and \$7.00. Again, she judged that all prices in the specified ranges were equally likely.

Costs also had to be factored in. Organic methods were more labor-intensive than conventional methods, but because Jane and Tim did the work themselves, this did not really affect profits. Even so, Jane estimated that costs for the field would be approximately \$800. With the uncertainty about the methods, however, she decided to represent costs with a normal distribution having mean \$800 and standard deviation \$50.

When she discussed this with Tim, he asked what this meant in terms of a "bottom line." What could they expect their gross sales to be? Could she give him some idea of how uncertain their profit would be? If they went with conventional farming methods, they could expect profits to be approximately \$2700, with a standard deviation of some \$100. If Jane could develop a probability distribution for profits under organic methods, perhaps the two distributions could be compared.

Tim added another wrinkle. He showed her an advertisement and a related story in the newspaper about a newly developed, genetically engineered bacterial pesticide named VegeTech. The ad claimed (and the story confirmed) that tests with VegeTech in their own geographical area had resulted in an average 10% increase in yield for all kinds of produce. The cost for this increase was approximately 20 cents per bushel. Given anticipated produce prices, this would appear to be cost-effective. The complication was that, while VegeTech had been fully approved by the appropriate federal agencies, OFA had not yet decided whether to approve its use as an organic substance. Proponents argued that it consisted of bacteria that attacked insects, while opponents argued that the bacteria had been synthesized and were not "organic" in the classic sense of the word. OFA had guaranteed that it would run tests and make a decision late this summer about the use of the substance; unfortunately, farmers would have to decide whether or not to use VegeTech before learning of OFA's decision. If the Kellers decided to use VegeTech, and if the OFA failed to certify it as an accepted organic substance, then the Kellers would not be able to label their produce "Organically Grown." The net effect would be a 15% reduction in the prices they could charge. By all indications, there was a 50-50 chance that the OFA would certify VegeTech.

Question

- 1 Construct a simulation model to address the issues that the Kellers face. Do you think that they should stick with the conventional methods or try organic agriculture? If they go organic, should they try VegeTech this year? Support your conclusions with appropriate simulation outputs (graphs, expected values, and so on).

OVERBOOKING, PART III

Consider again Mockingbird Airlines' problem as described in the overbooking case study in Chapter 9 (pages 345-346).

Questions

- 1 Construct a simulation model of the system, and use it to find Mockingbird's optimal policy regarding overbooking. Compare this answer with the one based on the analysis done in Chapter 9.
- 2 Suppose that you are uncertain about the no-show rate. It could be as low as 0.02 or it could be as high as 0.06, and all values in between are equally likely. Furthermore, the cost of satisfying the bumped passengers may not be constant. That is, the airline may in some cases be able to entice a passenger or two to relinquish their seats in exchange for compensation that would be less than a refund and another free ticket. Alternatively, in some cases the total cost, including loss of goodwill, might be construed as considerably higher. Suppose, for example, that the cost of satisfying an excess customer is normally distributed with mean \$300 and standard deviation \$40.

Modify the simulation model constructed in Question 1 to include the uncertainties about the no-show rate and the cost. Do these sources of uncertainty affect the optimal overbooking policy?

- 3 How else might Mockingbird's analysts address the uncertainty about the no-show rate and the cost?

REFERENCES

Hertz's (1964) article in *Harvard Business Review* extolled the virtues of simulation for the decision-analysis community early on. Hertz and Thomas (1983, 1984) provide discussion and examples of the use of simulation for decision analysis. Other texts that include introductory material on simulation are Holloway (1979) and Samson (1988). Vatter et al. (1978) contains several interesting simulation case studies in decision making. More technical introductions to Monte Carlo simulation at a moderate level are provided by Law and Kelton (1991) and Watson (1989).

Not only can analysts use programs such as @RISK and Crystal Ball, but mainstream decision-analysis programs also provide simulation capability. For example, the professional version of DPL includes the option of using simulation to evaluate a decision model.

Hertz, D. B. (1964) "Risk Analysis in Capital Investment." *Harvard Business Review*. Reprinted in *Harvard Business Review*, September-October 1979, 169-181.

Hertz, D. B., and H. Thomas (1983) *Risk Analysis and Its Applications*. New York: Wiley.

- Hertz, D. B., and H. Thomas (1984) *Practical Risk Analysis*. New York: Wiley.
- Holloway, C. A. (1979) *Decision Making under Uncertainty: Models and Choices*. Englewood Cliffs, NJ: Prentice-Hall.
- Law, A. M., and D. Kelton (1991) *Simulation Modeling and Analysis*, 2nd ed. New York: McGraw-Hill.
- Samson, D. (1988) *Managerial Decision Analysis*. Homewood, IL: Irwin.
- Vatter, P., S. Bradley, S. Frey, and B. Jackson (1978) *Quantitative Methods in Management: Text and Cases*. Homewood, IL: Irwin.
- Watson, G. (1989) *Computer Simulation*, 2nd ed. New York: Wiley.

CHAPTER 12. Value of Information

Decision makers who face uncertain prospects often gather information with the intention of reducing uncertainty. Information gathering includes consulting experts, conducting surveys, performing mathematical or statistical analyses, doing research, or simply reading books, journals, and newspapers. The intuitive reason for gathering information is straightforward; to the extent that we can reduce uncertainty about future outcomes, we can make choices that give us a better chance at a good outcome.

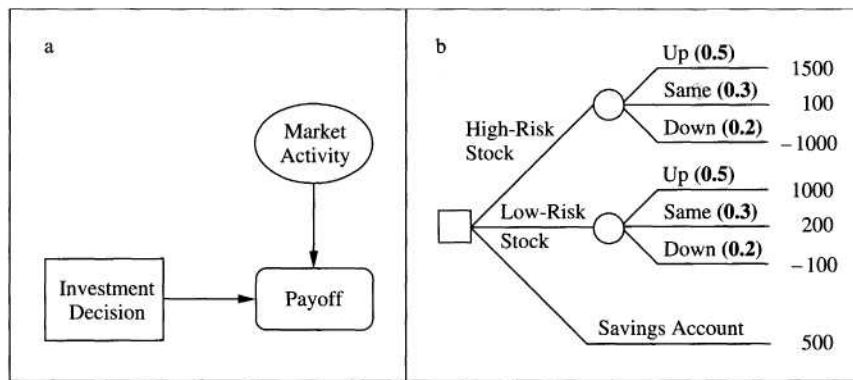
In this chapter, we will work a few examples that should help you understand the principles behind information valuation. Naturally, the examples also will demonstrate the techniques used to calculate information value. At the end of the chapter we will consider a variety of issues, including information in complex decisions, the use of information evaluation as an integral part of the decision-analysis process, what to do in the case of multiple nonmonetary objectives, and the problem of evaluating and selecting experts for the information they can provide.

The main example for this chapter is the stock market example that we introduced in Chapter 5 in our discussion of sensitivity analysis (pages 169-170). For convenience, the details are repeated here.

INVESTING IN THE STOCK MARKET

An investor has some funds available to invest in one of three choices: a high-risk stock, a low-risk stock, or a savings account that pays a sure \$500. If he invests in the stocks, he must pay a brokerage fee of \$200.

Figure 12.1
 (a) Influence-diagram
 and (b) decision-tree
 representations of the
 investor's problem.



His payoff for the two stocks depends on what happens to the market. If the market goes up, he will earn \$1700 from the high-risk stock and \$1200 from the low-risk stock. If the market stays at the same level, his payoffs for the high- and low-risk stocks will be \$300 and \$400, respectively. Finally, if the stock market goes down, he will lose \$800 with the high-risk stock but still gain \$100 with the low-risk stock.

The investor's problem can be modeled with either an influence diagram or a decision tree. These two representations are shown in Figure 12.1.

Value of Information: Some Basic Ideas

Before we begin an in-depth study of information from a decision-analysis perspective, let us consider certain fundamental notions. What does it mean for an expert to provide perfect information? How does probability relate to the idea of information? What is an appropriate basis on which to evaluate information in a decision situation? This section addresses these questions and thus sets the stage for a complete development of the value of information in the rest of the chapter.

Probability and Perfect Information

An expert's information is said to be perfect if it is always correct. We can use conditional probabilities to model perfect information. Suppose that when state S will occur, the expert always says so (and never says that some other state will occur). In our stock market example, imagine an expert who always correctly identifies a situation in which the market will increase:

$$P(\text{Expert Says "Market Up"} \mid \text{Market Really Does Go Up}) = 1$$

Because the probabilities must add to 1, we also must have

$$P(\text{Expert Says "Market Will Stay the Same or Fall" | Market Really Will Go Up}) = 0$$

But this is only half the story. The expert also must never say that state S will occur if any other state (S^c) will occur. There must be no chance of our expert saying that the market will rise when it really will not:

$$P(\text{Expert Says "Market Will Go Up" | Market Really Will Stay the Same or Fall}) = 0$$

Notice the difference between this probability statement and the preceding. Both are conditional probabilities, but the conditions are different.

If the expert's information is perfect, then upon hearing the expert's report no doubt about the future remains; if the expert says the market will rise, then we know that the market really will rise. Having used conditional probabilities to model the expert's perfect information, we can use Bayes' theorem to "flip" the probabilities as we did in Chapter 7 and show that there is no uncertainty after we have heard the expert. We want to know $P(\text{Market Really Will Go Up | Expert Says "Market Will Go Up"})$. Some notation will make our lives easier:

Market Up = The market really goes up
 Market Down = The market really stays flat or goes down
 Exp Says "Up" = The expert says the market will go up
 Exp Says "Down" = The expert says the market will stay flat or go down

Now we can apply Bayes' theorem:

$$\begin{aligned} P(\text{Market Up | Exp Says "Up"}) &= \frac{P(\text{Exp Says "Up" | Market Up})P(\text{Market Up})}{[P(\text{Exp Says "Up" | Market Up})P(\text{Market Up}) + P(\text{Exp Says "Up" | Market Down})P(\text{Market Down})]} \\ &= \frac{1 P(\text{Market Up})}{1 P(\text{Market Up}) + 0 P(\text{Market Down})} \\ &= 1 \end{aligned}$$

Observe that the posterior probability $P(\text{Market Up | Expert Says "Up"})$ is equal to 1 regardless of the prior probability $P(\text{Market Up})$. This is because of the conditional probabilities that we used to represent the expert's perfect performance. Of course, this situation is not typical of the real world. In real problems we rarely can eliminate uncertainty altogether. If the expert sometimes makes mistakes, these conditional probabilities would not be 1's and 0's and the posterior probability would not be 1 or 0; there still would be some uncertainty about what would actually happen.

This exercise may seem a bit arcane. Its purpose is to introduce the idea of thinking about information in a probabilistic way. We can use conditional probabilities and Bayes' theorem to evaluate all kinds of information in virtually any decision setting.

The Expected Value of Information

How can we place a value on information in a decision problem? For example, how could we decide whether to hire the expert described in the last section? Does it depend on what the expert says? In the investment decision, the optimal choice is to invest in the high-risk stock. Now imagine what could happen. If the expert says that the market will rise, the investor still would choose the high-risk stock. In this case, the information appears to have no value in the sense that the investor would have taken the same action regardless of the expert's information. On the other hand, the expert might say that the market will fall or remain the same, in which case the investor would be better off with the savings account. In this second case, the information has value because it leads to a different action, one with a higher expected value than what would have been experienced without the expert's information.

We can think about information value after the fact, as we have done in the preceding paragraph, but it is much more useful to consider it before the fact—that is, before we actually get the information or before we hire the expert. What effects do we anticipate the information will have on our decision? We will talk about the *expected value of information*. By considering the expected value, we can decide whether an expert is worth consulting, whether a test is worth performing, or which of several information sources would be the best to consult.

The worst possible case would be that, regardless of the information we hear, we still would make the same choice that we would have made in the first place. In this case, the information has zero expected value! If we would take the same action regardless of what an expert tells us, then why hire the expert in the first place? We are just as well off as we would have been without the expert. Thus, at the worst, the expected value of information is zero. But if there are certain cases — things an expert might say or outcomes of an experiment — on the basis of which we would change our minds and make a different choice, then the expected value of the information must be positive; in those cases, the information leads to a greater expected value. The expected value of information can be zero or positive, but never negative.

At the other extreme, perfect information is the best possible situation. Nothing could be better than resolving all of the uncertainty in a problem. When all uncertainty is resolved, we no longer have to worry about unlucky outcomes; for every choice, we know exactly what the outcome will be. Thus, the expected value of perfect information provides an upper bound for the expected value of information in general. Putting this together with the argument in the previous paragraph, the expected value of any information source must be somewhere between zero and the expected value of perfect information.

Finally, you might have noticed that we continue to consider the expected value of information in terms of the particular choices faced. Indeed, the expected value of information is critically dependent on the particular decision problem at hand. For this reason, different people in different situations may place different values on the same information. For example, General Motors may find that economic forecasts from an expensive forecaster may be a bargain in helping the company refine its production plans. The same economic forecasts may be an extravagant waste of money for a restaurateur in a tourist town.

Expected Value of Perfect Information

Now we will see how to calculate the expected value of perfect information (EVPI) in the investment problem. For an expected value-maximizing investor, the optimal choice is the high-risk stock because it has the highest EMV (\$580); however, this is partly because the investor is optimistic about what the market will do. How much would he be willing to pay for information about whether the market will move up, down, or sideways?

Suppose he could consult an expert with perfect information — a clairvoyant — who could reveal exactly what the market would do. By including an arrow from "Market Activity" to "Investment Decision," the influence diagram in Figure 12.2 represents the decision situation in which the investor has access to perfect information. Remember, an arrow leading from an uncertainty node to a decision node means that the decision is made knowing the outcome of the uncertainty node. This is exactly what we want to represent in the case of perfect information; the investor knows what the market will do before he invests his money.

With a representation of the decision problem including access to perfect information, how can we find the EVPI? Easy. Solve each influence diagram, Figures 12.1a and 12.2. Find the EMV of each situation. Now subtract the EMV for Figure 12.1a (\$580) from the EMV for Figure 12.2 (\$1000). The difference (\$420) is the EVPI. We can interpret this quantity as the maximum amount that the investor should be willing to pay the clairvoyant for perfect information.

It also is useful to look at the decision-tree representation. To do this, draw a decision tree that includes the opportunity to obtain perfect information (Figure 12.3). As in the influence-diagram representation, the EMV for consulting the clairvoyant is \$1000. This is \$420 better than the EMV obtained by acting without the information. As before, EVPI is the difference, \$420.

Recall that in a decision tree the order of the nodes conforms to a chronological ordering of the events. Is this what happens in the perfect-information branch in Figure 12.3? Yes and no. Yes in the sense that the uncertainty regarding the stock market's activity is resolved before the investment decision is made. This is the important part. But once the decision is made, the market still must go through

Figure 12.2
Perfect information in the investor's problem.


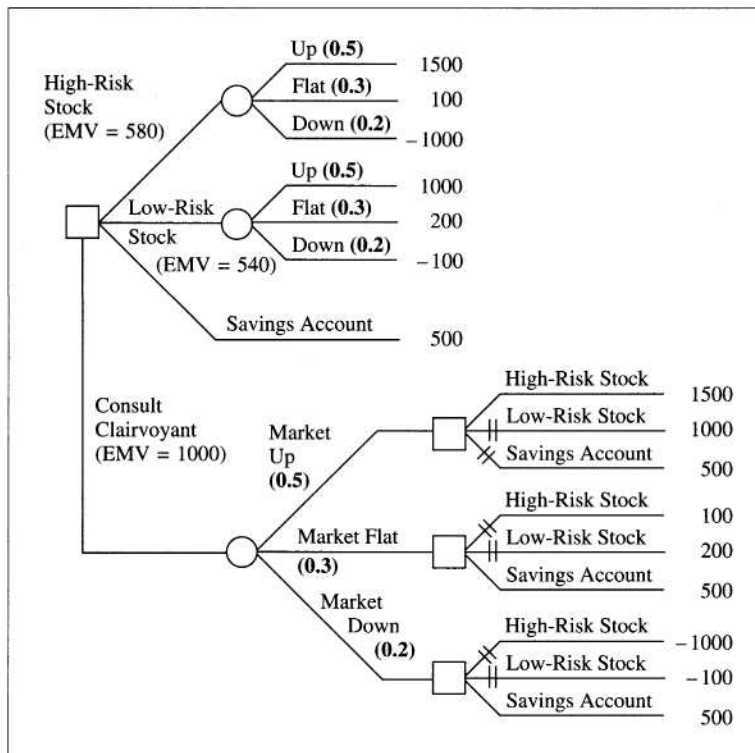


Figure 12.3
Investment decision tree with the perfect-information alternative.



its performance. It is simply that the investor knows exactly what that performance will be.

This points out a useful aspect of expected-value-of-information analysis with decision trees. If a decision maker faces some uncertainty in a decision, which is represented by those uncertainty nodes that come after his decision in a decision tree, redrawing the tree to capture the idea of perfect information is easy. Simply reorder the decision and uncertainty nodes! That is, redraw the tree so that the uncertainty nodes for which perfect information is available come before the decision node. This is exactly what we did in the perfect-information branch in Figure 12.3; in this branch, the "Market Activity" and "Investment Decision" nodes are reversed relative to their original positions.

It is worth reiterating that the way we are thinking about the value of information is in a strictly *a priori* sense. The decision-tree representation reinforces this notion because we actually include the decision branch that represents the possibility of consulting the clairvoyant. The investor has not yet consulted the clairvoyant; rather, he is considering whether to consult the clairvoyant in the first place. That action increases the expected value of the decision. Specifically, in this case there is a 50% chance that the clairvoyant will say that the market is not going up, in which case the appropriate decision would be the savings account rather than the high-risk stock.

Expected Value of Imperfect Information

We rarely have access to perfect information. In fact, our information sources usually are subject to considerable error. Thus, we must extend our analysis to deal with imperfect information.

The analysis of imperfect information parallels that of perfect information. We still consider the expected value of the information before obtaining it, and we will call it the *expected value of imperfect information* (EVII) to express the notion of collecting some information from a sample.


In the investment example, suppose that the investor hires an economist who specializes in forecasting stock market trends. Because he can make mistakes, however, he is not a clairvoyant, and his information is imperfect. For example, suppose his track record shows that if the market actually will rise, he says "up" 80% of the time, "flat" 10%, and "down" 10%. We construct a table (Table 12.1) to characterize his performance in probabilistic terms. The probabilities therein are conditional; for example, $P(\text{Economist Says "Flat"} \mid \text{Flat}) = 0.70$. The table shows that he is better when times are good (market up) and worse when times are bad (market down); he is somewhat more likely to make mistakes when times are bad.

How should the investor use the economist's information? Figure 12.4 shows an influence diagram that includes an uncertainty node representing the economist's forecast. The structure of this influence diagram should be familiar from Chapter 3; the economist's information is an example of imperfect information. The arrow from

Table 12.1
Conditional probabilities characterizing economist's forecasting ability.

Economist's Prediction	True Market State		
	Up	Flat	Down
"Up"	0.80	0.15	0.20
"Flat"*	0.10	0.70	0.20
"Down"	<u>0.10</u>	<u>0.15</u>	<u>0.60</u>
	1.00	1.00	1.00

Figure 12.4
Imperfect information in the investor's problem.



"Market Activity" to "Economic Forecast" means that the probability distribution for the particular forecast is conditioned on what the market will do. This is reflected in the distributions in Table 12.1. In fact, the distributions contained in the "Economic Forecast" node are simply the conditional probabilities from that table.

Solving the influence diagram in Figure 12.4 gives the EMV associated with obtaining the economist's imperfect information before action is taken. The EMV turns out to be \$822. As we did in the case of perfect information, we calculate EVII as the difference between the EMVs from Figures 12.4 and 12.1a, or the situation with no information. Thus, EVII equals $\$822 - \$580 = \$242$.

The influence-diagram approach is easy to discuss because we actually do not see the detail calculations. On the other hand, the decision-tree approach shows the calculation of EVII in its full glory. Figure 12.5 shows the decision-tree representation of the situation, with a branch that represents the alternative of consulting the economist. Look at the way in which the nodes are ordered in the "Consult Economist" alternative. The first event is the economist's forecast. Thus, we need probabilities $P(\text{Economist Says "Up"})$, $P(\text{Economist Says "Flat"})$, and $P(\text{Economist Says "Down"})$. Then the investor decides what to do with his money. Finally, the market goes up, down, or sideways. Because the "Market Activity" node follows the "Economists Forecast" node in the decision tree, we must have conditional probabilities for the market such as $P(\text{Market Up} \mid \text{Economist Says "Up"})$ or $P(\text{Market Flat} \mid \text{Economist Says "Down"})$. What we have, however, is the opposite. We have probabilities such as $P(\text{Market Up})$ and conditional probabilities such as $P(\text{Economist Says "Up"} \mid \text{Market Up})$.

As we did when we first introduced the notion of the value of an expert's information at the beginning of this chapter, we must use Bayes' theorem to find the posterior probabilities for the actual market outcome. For example, what is $P(\text{Market Up} \mid \text{Economist Says "Up"})$? It stands to reason that after we hear him say "up," we should think it more likely that the market actually will go up than we might have thought before. We used Bayes' theorem to "flip" probabilities in Chapter 7. There are several ways to think about this situation. First, applying Bayes' theorem is tantamount to reversing the arrow between the nodes "Market Activity" and "Economic Forecast" in Figure 12.4. In fact, reversing this arrow is the first thing that must be done when solving the influence diagram (Figure 12.6). Or we can think in terms of flipping a probability tree as we did in Chapter 7. Figure 12.7a represents the situation we have, and Figure 12.7b represents what we need.

Whether we think of the task as flipping a probability tree or reversing an arrow in an influence diagram, we still must use Bayes' theorem to find the probabilities we need. For example,

$$\begin{aligned} &P(\text{Market Up} \mid \text{Economist Says "Up"}) \\ &= P(\text{Up} \mid \text{"Up"}) \\ &= \frac{P(\text{"Up"} \mid \text{Up}) P(\text{Up})}{P(\text{"Up"} \mid \text{Up})P(\text{UP}) + P(\text{"Up"} \mid \text{Flat}) P(\text{Flat}) + P(\text{"Up"} \mid \text{Down}) P(\text{Down})} \end{aligned}$$

Figure 12.5
Incomplete decision tree for the investment example, including the alternative for consulting the economist.

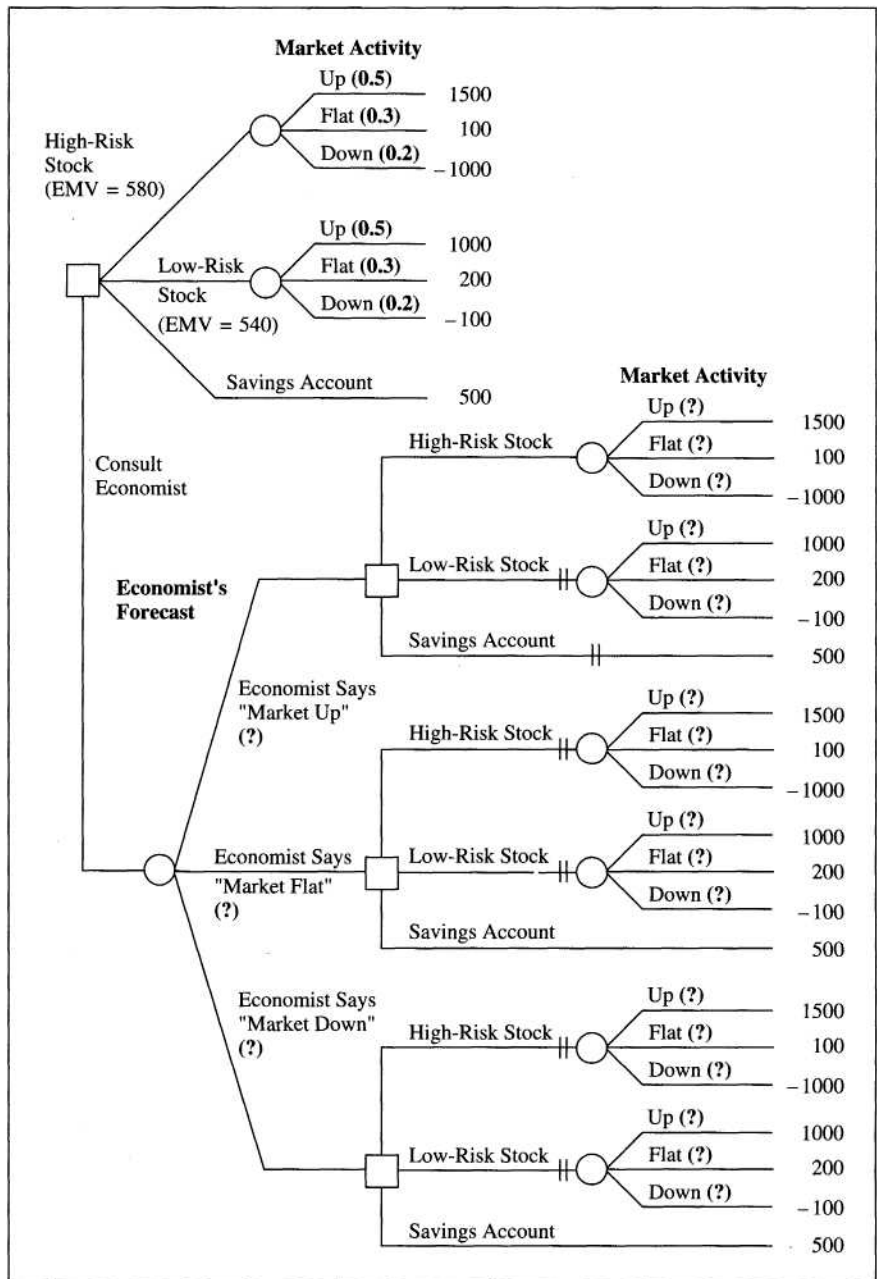


Figure 12.6

First step in solving the influence diagram. We reverse the arrow between "Economic Forecast" and "Market Activity."

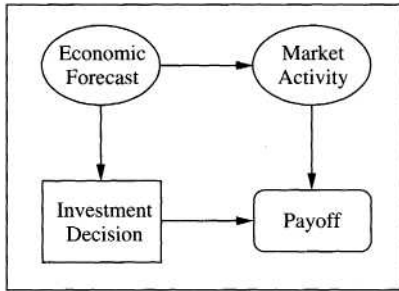
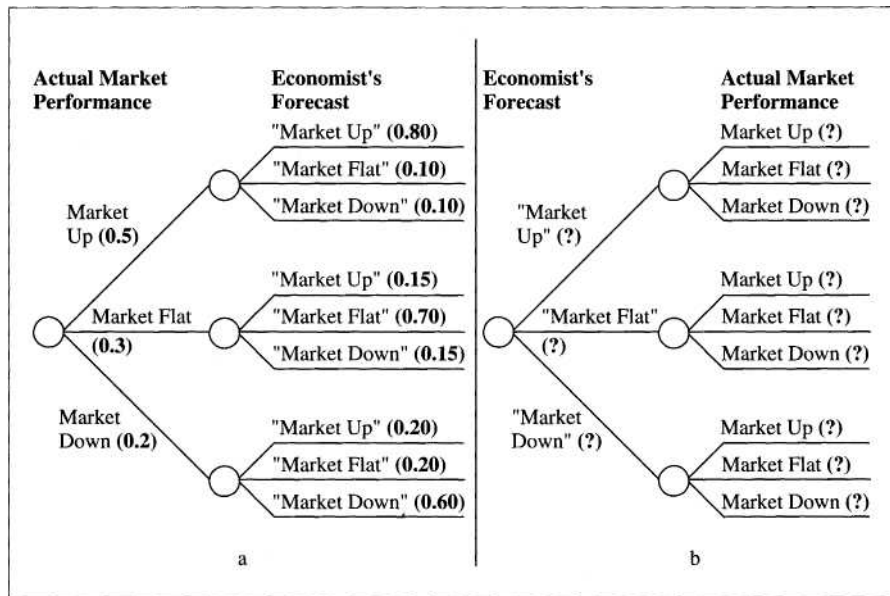


Figure 12.7

Flipping the probability tree to find posterior probabilities required for value-of-information analysis. In (a) we see what we have; in (b) we see what we need.



$P(\text{Up})$, $P(\text{Flat})$, and $P(\text{Down})$ are the investor's prior probabilities, while $P(\text{Economist Says "Up" | Up})$, and so on, are the conditional probabilities shown in Table 12.1. From the principle of total probability, the denominator is $P(\text{Economist Says "Up"})$. Substituting in values for the conditional probabilities and priors,

$$\begin{aligned}
 P(\text{Market Up} \mid \text{Economist Says "Up"}) &= \frac{0.8(0.5)}{0.8(0.5) + 0.15(0.3) + 0.2(0.2)} \\
 &= \frac{0.400}{0.485} \\
 &= 0.8247
 \end{aligned}$$

$P(\text{Economist Says "Up"})$ is given by the denominator and is equal to 0.485.

Of course, we need to use Bayes' theorem to calculate nine different posterior probabilities to fill in the gaps in the decision tree in Figure 12.5. Table 12.2 shows the results of these calculations; these probabilities are included on the appropriate branches in the completed decision tree (Figure 12.8).

Table 12.2
 Posterior probabilities
 for market trends
 depending on
 economist's
 information.

Economist's Prediction	Posterior Probability for:		
	Market Up	Market Flat	Market Down
"Up"	0.8247	0.0928	0.0825
"Flat"	0.1667	0.7000	0.1333
"Down"	0.2325	0.2093	0.5581

We also noted that we needed the marginal probabilities P("Up"), P("Flat"), and P("Down"). These probabilities are $P(\text{"Up"}) = 0.485$, $P(\text{"Flat"}) = 0.300$, and $P(\text{"Down"}) = 0.215$; they also are included in Figure 12.8 to represent our uncertainty about what the economist will say. As usual, the marginal probabilities can be found in the process of calculating the posterior probabilities because they simply come from the denominator in Bayes' theorem.

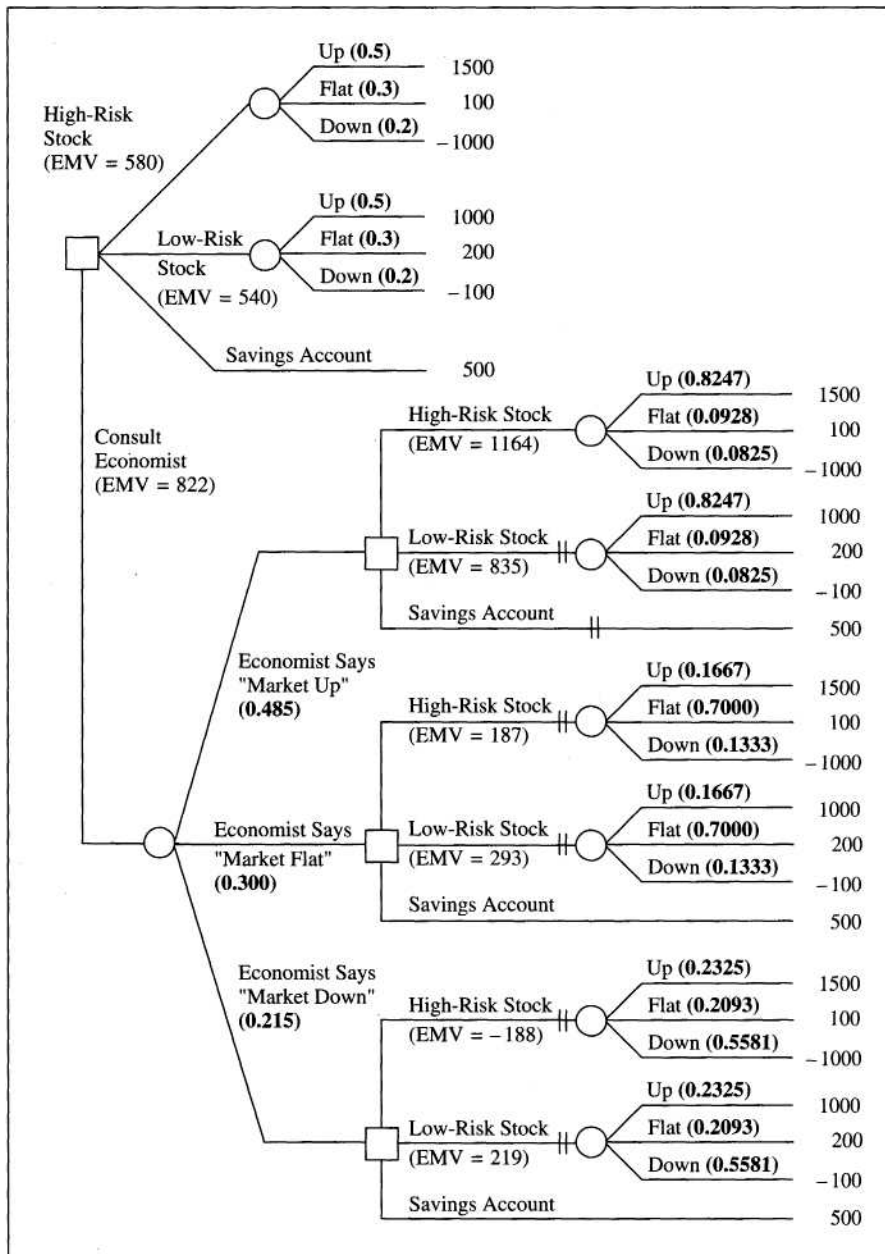
From the completed decision tree in Figure 12.8 we can tell that the EMV for consulting the economist is \$822, while the EMV for acting without consulting him is (as before) only \$580. The EVII is the difference between the two EMVs. Thus, EVSI is \$242 in this example, just as it was when we solved the problem using influence diagrams. Given this particular decision situation, the investor would never want to pay more than \$242 for the economic forecast.

As with perfect information, \$242 is the value of the information only in an expected-value sense. If the economist says that the market will go up, then we would invest in the high-risk stock, just as we would if we did not consult him. Thus, if he does tell us that the market will go up, the information turns out to do us no good. But if he tells us that the market will be flat or go down, we would put our money in the savings account and avoid the relatively low expected value associated with the high-risk stock. In those two cases, we would "save" $500 - 187 = 313$ and $500 - (-188) = 688$, respectively, with the savings in terms of expected value. Thus, EVSI also can be calculated as the "expected incremental savings," which is $0(0.485) + 313(0.300) + 688(0.215) = 242$.

Such probability calculations can make value-of-information analysis tedious and time-consuming. This is where computers can play a role. Influence-diagram programs can perform all of the necessary probability calculations and thus make finding EVSI a simple matter. Some decision-tree programs do the same thing, and it always is possible to construct an electronic spreadsheet model that will perform the calculations.

One last note regarding value-of-information calculations. If you have used your calculator to work through the examples in this chapter, you may have found that your answers differed slightly from those in this book. This is because calculations involving Bayes' theorem and value of information tend to be highly sensitive to rounding error. The appropriate strategy is to carry calculations out to many decimal places throughout a given problem, rounding off to dollars or cents only at the end of the problem. Rounding off intermediate results that are used in later calculations sometimes can make a large difference in the final answer.

Figure 12.8
Completed decision tree for the investment example.



Value of Information in Complex Problems

The example we have looked at in this section has been a fairly simple one. There was only one uncertain event, the behavior of the market, and we modeled the uncertainty with a simple discrete distribution. As we know, however, most real-world problems involve considerably more complex uncertainty models. In particular, we need to consider two specific situations. First, how can we handle continuous probability distributions? Second, what happens when there are many uncertain events and information is available about some or all of them?

The answer to the first question is straightforward conceptually, but in practice the calculation of EVPI or EVII may be difficult when dealing with continuous probability distributions. The principle is the same. Evaluate decision options with and without the information, and find the difference in the EMVs, as we have done in the discrete case. The problem, of course, is calculating the EMVs. Obviously, it always is possible to construct a discrete approximation as discussed in Chapter 8. Another possibility is to construct a Monte Carlo simulation model. Finally, for some theoretical probability models, analytical results are possible. The mathematics for such analysis, however, tend to be somewhat complicated and are beyond the scope of this introductory textbook. References for interested readers are included at the end of the chapter.

The second question asks how we handle value-of-information problems when there are many uncertain events. Again, in principle, the answer is easy, and is most transparent if we think in terms of influence diagrams. Perfect information about any particular event simply implies the presence of an informational arc from the event to the decision node. Naturally, it is possible to include such arcs for a subset of events. The only requirement is that the event not be downstream in the diagram from the decision node, because the inclusion of the information arc would lead to a cycle in the influence diagram. Solving the influence diagram with the informational arcs in place provides the EMV with the information, and this then may be compared to the EMV without the information to obtain the EVPI for the particular information sought.

Consider the same problem when the model is in decision-tree form. In the decision tree in Figure 12.3, the information branch was constructed by reversing the event node and the decision node. The same principle can apply if there are many sources of uncertainty; simply move those chance nodes for which information is to be obtained so that they precede the decision node. Now calculating EMV for the information branch will give the EMV in the case of perfect information for those events that precede the decision node in the decision tree.

For imperfect information, the same general principles apply. For influence diagrams, include an imperfect-information node that provides information to the decision maker. An excellent example is the toxic-chemicals influence diagram in Figure 3.20. Two sources of imperfect information are included in that model, the exposure survey and the lab test. In a decision-tree model, it would be a matter of constructing a tree having the appropriate informational chance nodes preceding the decision node. Unfortunately, if there are more than one or two such chance nodes, the decision tree

can become extremely unwieldy. Moreover, it may be necessary to calculate and track the marginal and posterior probabilities for the decision tree; these calculations can be done automatically in the influence diagram.

Value of Information, Sensitivity Analysis, and Structuring

Our motivation for studying value of information has been to examine situations in which information is available and to show how decisions can be made systematically regarding what source of information to select and how much an expert's information might be worth. Strictly speaking, this is precisely what value-of-information analysis can do. But it also can play an elegant and subtle role in the structuring of decisions and in the entire decision-analysis process of developing a requisite decision model. Recall the ideas of sensitivity analysis from Chapter 5. In that chapter we talked about a process of building a decision structure. The first step is to find out, using a tornado diagram, those variables to which the decision was sensitive; these variables require probabilistic modeling. The second step, after constructing a probabilistic model, may be to perform sensitivity analysis on the probabilities.

A third step in the structuring of a probabilistic model would be to calculate the EVPI for each uncertain event. This analysis would indicate where the analyst or decision maker should focus subsequent efforts in the decision-modeling process. That is, if EVPI is very low for an event, then there is little sense in spending a lot of effort in reducing the uncertainty by collecting information. But if EVPI for an event is relatively high, it may indeed be worthwhile to put considerable effort into the collection of information that relates to the event. Such information can have a relatively large payoff by reducing uncertainty and improving the decision maker's EMV. In this way, EVPI analysis can provide guidance to the decision analyst as to what issues should be tackled next in the development of a requisite decision model.

We will end this chapter with a short description of a rather unusual application of value-of-information analysis. Although few applications are as elaborate as this one, this example does provide an idea of how the idea of information value, as developed in decision analysis, can be used to address real-world concerns.

SEEDING HURRICANES

Hurricanes pack tremendous power in their high winds and tides. Recent storms such as Camille (1969) and Hugo (1989) caused damage in excess of a billion dollars and amply demonstrated the destructive potential of large hurricanes. In the 1960s, the U.S. government experimented with the seeding of hurricanes, or the practice of dropping silver iodide into the storm to reduce peak winds by forcing precipitation. After early limited experiments, Hurricane Debbie was seeded in 1969 with massive amounts of

silver iodide on two separate occasions. Each seeding was followed by substantial drops in peak wind speed. Given these results, should hurricanes that threaten highly populated areas be seeded on the basis of current knowledge? Should the government pursue a serious research program on the effects of hurricane seeding?

Howard, Matheson, and North addressed these specific questions about hurricane seeding. They asked whether it would be appropriate to seed hurricanes, or would a research program be appropriate, or should the federal government simply not pursue this type of weather modification? To answer these questions, they adopted a decision-analysis framework. On the basis of a relatively simple probabilistic model of hurricane winds, along with the relationships among wind speed, damage, and the effect of seeding, they were able to calculate expected dollar losses for two decision alternatives — seeding and not seeding a typical threatening hurricane. On the basis of their model, they concluded that seeding would be the preferred alternative if the federal government wanted to reduce expected damage.

The authors realized that the government might be interested in more than the matter of reducing property damage. What would happen, for example, if the decision was made to seed a hurricane, the wind speed subsequently increased, and increased damage resulted? The government most likely would become the target of many lawsuits for having taken action that appeared to have adverse effects. Thus, the government also faced an issue of responsibility if seeding were undertaken. On the basis of the authors' analysis, however, the government's "responsibility cost" would have to have been relatively high in order for the optimal decision to change.

To address the issue of whether further research on seeding would be appropriate, the authors used a value-of-information approach. They considered the possibility of repeating the seeding experiment that had been performed on Hurricane Debbie and the potential effects of such an experiment. By modeling the possible outcomes of this experiment and considering the effect on future seeding decisions, the authors were able to calculate the expected value of the research. Including reasonable costs for government responsibility and anticipated damage over a single hurricane season, the expected value of the experiment was determined to be approximately \$10.2 million. The authors also extended their analysis to include all future hurricane seasons, discounting future costs by 7% per year. In this case, the expected value of the research was \$146 million. To put these numbers in perspective, the cost of the seeding experiment would have been approximately \$500,000. [Source: R. A. Howard, J. E. Matheson, and D. W. North (1972) "The Decision to Seed Hurricanes." *Science*, 176, 1191-1202.]

Value of Information and Nonmonetary Objectives

Throughout this chapter we have calculated the expected value of information on the basis of EMVs, implicitly assuming that the only objective that matters is making money. This has been a convenient fiction; as you know, in many decision

situations there are multiple objectives. For example, consider the Federal Aviation Administration (FAA) bomb-detection example again (Chapter 3, Figure 3.8). Recall that the FAA was interested in maximizing the detection effectiveness and passenger acceptance of the system while at the same time minimizing the cost and time to implementation. For any given system, there may be considerable uncertainty about the level of passenger acceptance, for example, but this uncertainty could be reduced in many ways. A less expensive, but highly imperfect, option would be to run a survey. A more expensive test would be to install some of the systems at a few airports and try them out. But how much should be spent on such efforts?

With the FAA example, the answer is relatively straightforward, because minimizing cost happens to be one of the objectives. The answer would be to find the additional cost that makes the net expected value of getting the information equal to the expected value without the information (and, naturally, without the cost). If there is a clear trade-off rate that can be established for, say, dollars of cost per additional point on the passenger-acceptance scale, then the increase in expected passenger-acceptance level that results from additional information can be translated directly into dollars of cost. (Any increase in the expected value of the other objectives—which may happen due to choosing different options under different information scenarios—would also have to be included in the calculations. Doing so can also be accomplished through the specification of similar trade-offs between cost and the other objectives. Understanding such trade-offs is not complex and is treated in more detail in Chapter 15.)

When a decision situation does not involve a monetary objective, the same techniques that we have developed in this chapter can still be used. Suppose one objective is to minimize the decision maker's time; different choices and different outcomes require different amounts of time from the decision maker. Information can be valued in terms of time; the expected value of perfect information might turn out to be, say, five days of work. If resolving some uncertainty would take more time than that, it would not be worth doing!

Value of Information and Experts

In Chapter 8 we discussed the role of experts in decision analysis. One of the issues that analysts face in using experts is how to value them and how to decide how many and which experts to consult. In general, the valuation part is not a practical concern; the high stakes involved in typical public-policy risk analyses typically warrant the use of experts who may charge several thousand dollars per day in consulting fees. What is less obvious, however, is that experts typically provide information that is somewhat interrelated. Because experts tend to read the same journals, go to the same conferences, use the same techniques in their studies, and even communicate with each other, it comes as no surprise that the infor-

mation they provide can be highly redundant. The real challenge in expert use is to recruit experts who look at the same problem from very different perspectives. Recruiting experts from different fields, for example, can be worthwhile if the information provided is less redundant. It can even be the case that a highly diverse set of less knowledgeable (and less expensive) experts can be much more valuable than the same number of experts who are more knowledgeable (and cost more) but give redundant information!

Calculating EVPI: DATA and DPL

As we have seen in this chapter, performing EVPI calculations involves nothing more than reordering nodes in a decision tree. This is easy to do in any decision-tree program, and it can be done for any combination of timing and event information simply by restructuring the tree. DATA provides a menu command specifically for calculating EVPI for a specific chance node or set of nodes. DPL has no specific EVPI command but does provide a menu command called "Reorder Node," which facilitates restructuring the tree for EVPI calculations. One final comment about DPL: Once the decision tree has been created, EVPI calculations must be done by restructuring the tree; adding arcs in the influence-diagram view will not do the trick, because DPL uses the decision tree to specify sequencing of the nodes.

SUMMARY

By considering the expected value of information, we can make better decisions about whether to obtain information or which information source to consult. We saw that the expected value of any bit of information must be zero or greater, and it cannot be more than the expected value of perfect information. Both influence diagrams and decision trees can be used as frameworks for calculating expected values. Influence diagrams provide the neatest representation because information available for a decision can be represented through appropriate use of arcs and, if necessary, additional uncertainty nodes representing imperfect information. In contrast, representing the expected value of imperfect information with decision trees is more complicated, requiring the calculation of posterior and marginal probabilities. The expected value of information is simply the difference between the EMV calculated both with and without the information.

The final sections in the chapter discussed generally how to solve value-of-information problems in more complex situations. We concluded with discussions of the role that value-of-information analysis can play in the decision-analysis process of developing a requisite decision model, how to value information related to nonmonetary objectives, and evaluation and selection of experts.

EXERCISES

- 12.1 Explain why in decision analysis we are concerned with the *expected* value of information.
- 12.2 Calculate the EVPI for the decision shown in Figure 12.9.
- 12.3 What is the EVPI for the decision shown in Figure 12.10? Must you perform any calculations? Can you draw any conclusions regarding the relationship between value of information and deterministic dominance?
- 12.4 For the decision tree in Figure 12.11, assume Chance Events E and F are independent.
 - a Draw the appropriate decision tree and calculate the EVPI for Chance Event E only.
 - b Draw the appropriate decision tree and calculate the EVPI for Chance Event F only.
 - c Draw the appropriate decision tree and calculate the EVPI for both Chance Events E and F: that is, perfect information for both E and F is available before a decision is made.
- 12.5 Draw the influence diagram that corresponds to the decision tree for Exercise 12.4. How would this influence diagram be changed in order to answer parts a, b, and c in Exercise 12.4?

Figure 12.9

Generic decision tree for Exercise 12.2.




Figure 12.10

Generic decision tree for Exercise 12.3.





Figure 12.11
Generic decision tree
for Exercise 12.4.



QUESTIONS AND PROBLEMS

- 12.6** The claim was made in the chapter that information always has positive value. What do you think of this? Can you imagine any situation in which you would prefer *not* to have some unknown information revealed?
- a** Suppose you have just visited your physician because of a pain in your abdomen. The doctor has indicated some concern and has ordered some tests whose results the two of you are expecting in a few days. A positive test result will suggest that you may have a life-threatening disease, but even if the test is positive, the doctor would want to confirm it with further tests. Would you want the doctor to tell you the outcome of the test? Why or why not?
 - b** Suppose you are selling your house. Real-estate transaction laws require that you disclose what you know about significant structural defects. Although you know of no such defects, a buyer insists on having a qualified engineer inspect the house. Would you want to know the outcome of the inspection? Why or why not?
 - c** Suppose you are negotiating to purchase an office building. You have a lot of experience negotiating commercial real-estate deals, and your agent has explained this to the seller, who is relatively new to the business. As a result, you expect to do very well in this negotiation. Because of the unique circumstance of the building, your agent has suggested obtaining an appraisal of the property by a local expert. You know exactly how you would use this information; it would provide an upper bound on what you are willing to pay. This fact is also clear to the seller, who will know whether you obtained the appraisal but will only find out the appraised value if you elect to reveal it. Would you obtain the appraisal? Why or why not?
- 12.7** Consider another oil-wildcatting problem. You have mineral rights on a piece of land that you believe may have oil underground. There is only a 10% chance that you will strike oil if you drill, but the payoff is \$200,000. It costs \$10,000 to drill. The alternative is not to drill at all, in which case your profit is zero.
- a** Draw a decision tree to represent your problem. Should you drill?

- b Draw an influence diagram to represent your problem. How could you use the influence diagram to find EVPI?
 - c Calculate EVPI. Use either the decision tree or the influence diagram.
 - d Before you drill you might consult a geologist who can assess the promise of the piece of land. She can tell you whether your prospects are "good" or "poor." But she is not a perfect predictor. If there is oil, the conditional probability is 0.95 that she will say prospects are good. If there is no oil, the conditional probability is 0.85 that she will say poor. Draw a decision tree that includes the "Consult Geologist" alternative. Be careful to calculate the appropriate probabilities to include in the decision tree. Finally, calculate the EVII for this geologist. If she charges \$7000, what should you do?
- 12.8** In Problem 4.16, find:
- a EVPI for the cost of making the processor.
 - b EVPI for the cost of subcontracting the processor.
 - c EVPI for both uncertain events.
- 12.9** Look again at Problem 8.11, which concerns whether you should drop your decision-analysis course. Estimate EVPI for your score in the course.
- 12.10** In Problem 9.33, the issue is whether or not to close the plant, and if your boss knew exactly how many machines would fail during your absence, he would be able to decide what to do without the fear of making a mistake.
- a Find his EVPI concerning the number of machine failures during your absence over the next two weeks.
 - b Suppose that the cost of repairing the broken machines is \$30,000. Then what is EVPI?
 - c Suppose that the cost of repairing the machines is \$15,000 but the cost of closing the plant is \$20,000. Now calculate EVPI.
- 12.11** Consider the Texaco-Pennzoil example from Chapter 4 (pages 101-102).
- a What is EVPI to Hugh Liedtke regarding Texaco's reaction to a counteroffer of \$5 billion? Can you explain this result intuitively?
 - b The timing of information acquisition may make a difference.
 - (i) For example, suppose that Liedtke could obtain information about the final court decision before making his current decision (take the \$2 billion or counteroffer \$5 billion). What would be EVPI of this information?
 - (ii) Suppose that Liedtke knew he would be able to obtain perfect information only after he has made his current decision but before he would have to respond to a potential Texaco counteroffer of \$3 billion. What would be EVPI in this case?
 - c In part b, EVPI for (ii) should be less than EVPI calculated in (i). Can you explain why? (Incidentally, if your results disagree with this, you should check your calculations!)
- 12.12** In the Texaco-Pennzoil case, what is EVPI if Liedtke can learn both Texaco's reaction and the final court decision before he makes up his mind about the current \$2 billion offer? (*Hint*: Your answer should be more than the sum of the EVPIs for Texaco's reaction and the court decision calculated separately in Problem 12.11.) Can you explain why the interaction of the two bits of information should have this effect?

- 12.13** In Problem 5.9, assume that the grower's loss incurred with the burners would be \$17,500 and that the loss incurred with the sprinklers would be \$27,500.
- Find EVPI for the weather conditions (freeze or not).
 - Now assume that the loss incurred with the burners is uniformly distributed between \$15,000 and \$20,000. Also assume that the loss incurred with the sprinklers is uniformly distributed between \$25,000 and \$30,000. Now estimate EVPI regarding these losses, under the assumption that a better weather forecast cannot be obtained.
 - Do you think the farmer should put more effort into learning his costs more precisely or should he concentrate on obtaining better weather forecasts?
- 12.14** Reconsider the Eagle Airlines example from Chapter 5, particularly as diagrammed in Figure 5.7. Assume that q and r are both 0.5. Calculate EVPI for each of the three uncertain events individually. What can you conclude from your analysis?
- 12.15** In Exercise 12.4, is it necessary to assume that the events are independent? What other assumption could be made?

C A S E S T U D I E S

TEXACO PENNZOIL REVISITED

Often when we face uncertainty, we would like more than simply to know the outcome; it would be nice to control the outcome and bring about the best possible result! A king might consult his wizard as well as his clairvoyant, asking the wizard to cast a spell to cause the desired outcome to occur. How much should such a wizard be paid for these services? The expected value of his wizardry naturally depends on the decision problem at hand, just as the expected value of information does. But the way to calculate the "expected value of wizardry" (to use Ron Howard's term) is very similar to solving the calculations for the expected value of perfect information. To demonstrate this idea, we again will examine Hugh Liedtke's decision situation as diagrammed in Figure 4.2. Now consider the entirely hypothetical possibility that Liedtke could pay someone to influence Texaco's CEO, James Kinnear.

Questions

- What would be the most desirable outcome from the "Texaco Reaction" chance node?
- Construct the decision tree now with three alternatives: "Accept \$2 Billion," "Counteroffer \$5 Billion," and "Counteroffer \$5 Billion and Influence Kinnear."
- Solve your decision tree from Question 2. What is the maximum amount that Liedtke could afford to pay in order to influence Kinnear?

MEDICAL TESTS

One of the principles that arises from a decision-analysis approach to valuing information is that information is worthless if no possible informational outcome will change the decision. For example, suppose that you are considering whether to make a particular investment. You are tempted to hire a consultant recommended by your Uncle Jake (who just went bankrupt last year) to help you analyze the decision. If, however, you think carefully about the things that the consultant might say and conclude that you would (or would not) make the investment regardless of the consultant's recommendation, then you should not hire the consultant. This principle makes perfectly good sense in the light of our approach; do not pay for information that cannot possibly change your mind.

In medicine, however, it is standard practice for physicians to order extensive batteries of tests for patients. Although different kinds of patients may be subjected to different overall sets of tests, it is nevertheless the case that many of these tests provide information that is worthless in a decision-analysis sense; the doctor's prescription would be the same regardless of the outcome of a particular test.

Questions

- 1 As a patient, would you be willing to pay for such tests? Why or why not?
- 2 What incentives do you think the doctor might have for ordering such tests, assuming he realizes that his prescription would not change?
- 3 How do his incentives compare to yours?

DUMOND INTERNATIONAL, PART II

[Refer back to the DuMond International case study at the end of Chapter 5 (pages 182-183).] Nancy Milnor had returned to her office, still concerned about the decision. Yes, she had persuaded the directors that their disagreements did not affect her analysis; her analysis still showed the new product to be the appropriate choice. The members of the board, however, had not been entirely satisfied. The major complaint was that there was still too much uncertainty. Could she find out more about the likelihood of a ban, or could she get a better assessment from engineering regarding the delay? What about a more accurate sales forecast for the new product?

Nancy gazed at her decision tree (Figure 5.21). Yes, she could address each of those questions, but where should she start?

Question

- 1 Calculate EVPI for the three uncertain events in DuMond's decision as diagrammed in Figure 5.21. Where should Nancy Milnor begin her investigation?

REFERENCES

This chapter has focused primarily on the technical details of calculating the value of information. The most complete, and most highly technical, reference for this kind of analysis is Raiffa and Schlaifer (1961). Winkler (1972) provides an easily readable discussion of value of information. Both texts contain considerably more material than what is included here, including discussion of EVPI and EVII for continuous distributions. In particular, they delve into questions of the value of sample information for updating natural conjugate prior distributions for model parameters (Chapter 10 in this text) and the selection of appropriate sample sizes, a decision problem intimately related to value of information. Another succinct and complete reference is Chapter 8 in LaValle (1978).

Three articles by LaValle (1968a, b, c) contain many results regarding the expected value of information, including the result that the expected value of information about multiple events need not equal the sum of the values for individual events. Ponsard (1976) and LaValle (1980) show that expected value of information can be negative in some competitive situations. Clemen and Winkler (1985) discuss the value of information from experts, especially focusing on the effect of probabilistic dependence among experts.

At the end of the chapter, we discussed the way in which the value of information can be related to sensitivity analysis and the decision-analysis process. These ideas are Ron Howard's and have been part of his description of decision analysis since the early 1960s. Many of the articles in Howard and Matheson (1983) explain this process, and illustrative applications show how the process has been applied in a variety of real-world decision problems.

- Clemen, R., and R. Winkler (1985) "Limits for the Precision and Value of Information from Dependent Sources." *Operations Research*, 33, 427–442.
- Howard, R.A., and J. E. Matheson (eds.) (1983) *The Principles and Applications of Decision Analysis* (2 volumes). Palo Alto, CA: Strategic Decisions Group.
- LaValle, I. (1968a) "On Cash Equivalents and Information Evaluation in Decisions under Uncertainty; Part I. Basic Theory." *Journal of the American Statistical Association*, 63, 252-276.
- LaValle, I. (1968b) "On Cash Equivalents and Information Evaluation in Decisions under Uncertainty; Part II. Incremental Information Decisions." *Journal of the American Statistical Association*, 63, 277-284.
- LaValle, I. (1968c) "On Cash Equivalents and Information Evaluation in Decisions under Uncertainty; Part III. Exchanging Partition-*J* for Partition-*K* Information." *Journal of the American Statistical Association*, 63, 285-290.
- LaValle, I. (1978) *Fundamentals of Decision Analysis*. New York: Holt.

LaValle, I. (1980) "On Value and Strategic Role of Information in Semi-Normalized Decisions." *Operations Research*, 28, 129-138.

Ponssard, J. (1976) "On the Concept of the Value of Information in Competitive Situations." *Management Science*, 22, 737-749.

Raiffa, H., and R. Schlaifer (1961) *Applied Statistical Decision Theory*. Cambridge, MA: Harvard University Press.

Winkler, R. L. (1972) *Introduction to Bayesian Inference and Decision*. New York: Holt.

CHAPTER 13. Risk Attitudes

This chapter marks the beginning of our in-depth study of preferences. Before we begin, let us review where we have been and think about where we are going. The first six chapters provided an introduction to the process of structuring decision problems for decision analysis and an overview of the role that probability and utility theory play in making choices. Chapters 7 through 12 have focused on probability concerns: using probability in a variety of ways to model uncertainty in decision problems, including the modeling of information sources in value-of-information problems.

At this point, we change directions and look at the preference side of decision analysis. How can we model a decision maker's preferences? This chapter looks at the problems associated with risk and return trade-offs. Chapter 14 briefly explores the axiomatic foundations of utility theory and discusses certain paradoxes from cognitive psychology. These paradoxes generally indicate that people do not make choices that are perfectly consistent with the axioms, even though they may agree that the axioms are reasonable! Although such inconsistencies generally do not have serious implications for most decisions, there are certain occasions when they can cause difficulty.

The primary motivating example for this chapter comes from the history of railroads in the United States. Imagine what might have gone through E. H. Harriman's mind as he considered his strategy for acquiring the Northern Pacific Railroad in March 1901.

E. H. HARRIMAN FIGHTS FOR THE NORTHERN PACIFIC RAILROAD

"How could they do it?" E. H. Harriman asked, still angry over the fact that James Hill and J. P. Morgan had bought the Burlington Railroad out from under his nose. "Every U.S. industrialist knows I control the railroads in the West. I have the Illinois Central, the Union Pacific, the Central and Southern Pacific, not to mention the Oregon Railroad and Navigation Company. Isn't that true?"

"Yes, sir," replied his assistant.

"Well, we will put the pressure on Messrs. Hill and Morgan. They will be surprised indeed to find out that I have acquired a controlling interest in their own railroad, the Northern Pacific. I may even be able to persuade them to let me have the Burlington. By the way, how are the stock purchases going?"

"Sir, we have completed all of the purchases that you authorized so far. You may have noticed that our transactions have driven the price of Northern Pacific stock up to more than \$100 per share."

Harriman considered this information. If he bought too fast, he could force the stock price up high enough and fast enough that Hill might begin to suspect that Harriman was up to something. Of course, if Harriman could acquire the shares quickly enough, there would be no problem. On the other hand, if he bought the shares slowly, he would pay lower prices, and Hill might not notice the acquisition until it was too late. His assistant's information, however, suggested that his situation was somewhat risky. If Harriman's plan were discovered, Hill could persuade Morgan to purchase enough additional Northern Pacific shares to enable them to retain control. In that case, Harriman would have paid premium prices for the stock for nothing! On the other hand, if Hill did not make the discovery immediately, the triumph would be that much sweeter.

"How many more shares do we need to have control?" asked Harriman.

"If you could purchase another 40,000 shares, sir, you would own 51 percent of the company."

Another 40,000 shares. Harriman thought about giving Hill and Morgan orders on how to run their own railroad. How enjoyable that would be! Yes, he would gladly increase his investment by that much.

"Of course," his assistant continued, "if we try to purchase these shares immediately, the price will rise very quickly. You will probably end up paying an additional \$15 per share above what you would pay if we were to proceed more slowly."

"Well, \$600,000 is a lot of money, and I certainly would not want to pay more. But it would be worth the money to be sure that we would be able to watch Hill and Morgan squirm! Send a telegram to my broker in New York right away to place the order. And be quick! It's already Friday. If we are going to do this, we need to do it today. I don't want Hill to have the chance to think about this over the weekend."

Risk

Basing decisions on expected monetary values (EMVs) is convenient, but it can lead to decisions that may not seem intuitively appealing. For example, consider the following two games. Imagine that you have the opportunity to play one game or the other, but only one time. Which one would you prefer to play? Your choice also is drawn in decision-tree form in Figure 13.1.

Game 1 Win \$30 with probability 0.5

Lose \$1 with probability 0.5

Game 2 Win \$2000 with probability 0.5

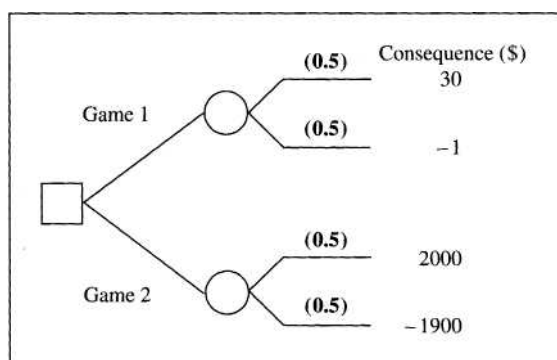
Lose \$1900 with probability 0.5

Game 1 has an expected value of \$14.50. Game 2, on the other hand, has an expected value of \$50.00. If you were to make your choice on the basis of expected value, then you would choose Game 2. Most of us, however, would consider Game 2 to be riskier than Game 1, and it seems reasonable to suspect that most people actually would prefer Game 1.

Using expected values to make decisions means that the decision maker is considering only the average or expected payoff. If we take a long-run frequency approach, the expected value is the average amount we would be likely to win over many plays of the game. But this ignores the range of possible values. After all, if we play each game 10 times, the worst we could do in Game 1 is to lose \$10. On the other hand, the worst we could do in Game 2 is lose \$19,000!

Many of the examples and problems that we have considered so far have been analyzed in terms of expected monetary value (EMV). EMV, however, does not capture risk attitudes. For example, consider the Texaco-Pennzoil example in Chapter 4 (pages 101-102). If Hugh Liedtke were afraid of the prospect that Pennzoil could end

Figure 13.1
Two lottery games.
Which game would you choose?



up with nothing at the end of the court case, he might be willing to take the \$2 billion that Texaco offered. To consider the Eagle Airlines case (Chapter 5, pages 155-156), purchasing the airplane is a much riskier alternative than leaving the money in the bank. If Carothers were sensitive to risk, he might prefer to leave the money in the bank. Even someone like E. H. Harriman considered the riskiness of the situations in which he found himself. In our example, Harriman weighed the value of a riskless alternative (immediately purchasing the 40,000 shares that were required to gain control) against the risky alternative of not purchasing the shares and the possible outcomes that might then follow. Even though all of the dollar amounts were not specified, it is clear that Harriman was not thinking in terms of EMV.

Individuals who are afraid of risk or are sensitive to risk are called *risk-averse*. We could explain risk aversion if we think in terms of a *utility function* (Figure 13.2) that is curved and opening downward (the technical term for a curve with this shape is *concave*). This utility function represents a way to translate dollars into "utility units." That is, if we take some dollar amount (x), we can locate that amount on the horizontal axis. Read up to the curve and then horizontally across to the vertical axis. From that point we can read off the utility value $U(x)$ for the dollars we started with.

A utility function might be specified in terms of a graph, as in Figure 13.2, or given as a table, as in Table 13.1. A third form is a mathematical expression. If graphed, for example, all of the following expressions would have the same general concave shape (opening downward) as the utility function graphed in Figure 13.2:

$$U(x) = \log(x)$$

$$U(x) = +\sqrt{x} \quad [\text{or } U(x) = x^{0.5}]$$

Of course, the utility and dollar values in Table 13.1 also could be graphed, as could the functional forms shown above. Likewise, the graph in Figure 13.2 could be con-

Figure 13.2
A utility function that displays risk aversion.




Table 13.1
A utility function in
tabular form.

Wealth	Utility Value
2500	1.50
1500	1.24
1000	0.93
600	0.65
400	0.47
0	0.15

verted into a table of values. The point is that the utility function makes the translation from dollars to utility regardless of its displayed form.

Risk Attitudes

We think of a typical utility curve as (1) upward sloping and (2) concave (the curve opens downward). An upward-sloping utility curve makes fine sense; it means that more wealth is better than less wealth, everything else being equal. Few people will argue with this. Concavity in a utility curve implies that an individual is risk-averse. Imagine that you are forced to play the following game:


Win \$500 with probability 0.5

Lose \$500 with probability 0.5

Would you pay to get out of this situation? How much? The game has a zero expected value, so if you would pay something to get out, you are avoiding a risky situation with zero expected value. Generally, if you would trade a gamble for a sure amount that is less than the expected value of the gamble, you are risk-averse. Purchasing insurance is an example of risk-averse behavior. Insurance companies analyze a lot of data in order to understand the probability distributions associated with claims for different kinds of policies. Of course, this work is costly. To make up these costs and still have an expected profit, an insurance company must charge more for its insurance policy than the policy can be expected to produce in claims. Thus, unless you have some reason to believe that you are more likely than others in your risk group to make a claim, you probably are paying more in insurance premiums than the expected amount you would claim.

Not everyone displays risk-averse behavior all the time, and so utility curves need not be concave. A convex (opening upward) utility curve indicates risk-seeking behavior (Figure 13.3). The risk seeker might be eager to enter into a gamble; for example, he or she might pay to play the game just described. An individual who plays a state lottery exhibits risk-seeking behavior. State lottery tickets typically cost \$1.00 and have an expected value of approximately 50 cents.

Figure 13.3
*Three different shapes
 for utility functions.*



Finally, an individual can be risk-neutral. Risk neutrality is reflected by a utility curve that is simply a straight line. For this type of person, maximizing EMV is the same as maximizing expected utility. This makes sense; someone who is risk-neutral does not care about risk and can ignore risk aspects of the alternatives that he or she faces. Thus, EMV is a fine criterion for choosing among alternatives, because it also ignores risk.

Although most of us are not risk-neutral, it often is reasonable for a decision maker to assume that his or her utility curve is nearly linear in the range of dollar amounts for a particular decision. This is especially true for large corporations that make decisions involving amounts which are small relative to their total assets. In many cases, it may be worthwhile to use EMV in a first-cut analysis, and then check to see whether the decision would be sensitive to changes in risk attitude. If the decision turns out to be fairly sensitive (that is, if the decision would change for a slightly risk-averse or slightly risk-seeking person), then the decision maker may want to consider modeling his or her risk attitude carefully.

This discussion makes it sound as though individuals can be reduced to their utility functions, and those utility functions can reveal whether the individual is risk-averse or risk-seeking. Keep in mind, however, that the utility function is only a model of an individual's attitude toward risk. Moreover, our development of utility functions in this chapter is intended to help with the modeling of risk attitudes at a fundamental level, and our model may not be able to capture certain complicated psychological aspects. For example, some individuals may be extremely frightened by risk. Others may find that small wagers greatly increase their enjoyment in watching a sporting event, for example. Still others may find that waiting for the uncertainty to be resolved is a source of excitement and exhilaration, although concern about losing money is a source of anxiety. For some people, figuring out exactly what their feelings are toward risky alternatives may be extremely complicated and may depend on the amount at stake, the context of the risk, and the time horizon.

Investing in the Stock Market, Revisited

If we have a utility function that translates from dollars to utility, how should we use it? The whole idea of a utility function is that it should help to choose from among alternatives that have uncertain payoffs. Instead of maximizing expected value, the decision maker should maximize expected utility. In a decision tree or influence-diagram payoff table, the net dollar payoffs would be replaced by the corresponding utility values and the analysis performed using those values. The best choice then should be the action with the highest expected utility.

As an example, let us reconsider the stock market-investment example from Chapters 5 and 12. You will recall that an investor has funds that he wishes to invest. He has three choices: a high-risk stock, a low-risk stock, or a savings account that would pay \$500. If he invests in the stocks, he must pay a \$200 brokerage fee.

With the two stocks his payoff depends on what happens to the market. If the market goes up, he will earn \$1700 from the high-risk stock and \$1200 from the low-risk stock. If the market stays at the same level, his payoffs for the high- and low-risk stocks will be \$300 and \$400, respectively. Finally, if the stock market goes down, he will lose \$800 with the high-risk stock but still earn \$100 from the low-risk stock. The probabilities that the market will go up, stay the same, or go down are 0.5, 0.3, and 0.2, respectively.

Figure 13.4 shows his decision tree, including the brokerage fee and the payoffs for the two stocks under different market conditions. Note that the values at the ends of the branches are the *net* payoffs, taking into account both the brokerage fee and the investment payoff. Table 13.2 gives his utility function.

We already calculated the expected values of the three investments in Chapter 12. They are

$$\text{EMV (High-Risk Stock)} = 580$$

$$\text{EMV (Low-Risk Stock)} = 540$$

$$\text{EMV (Savings Account)} = 500$$

Figure 13.4
Decision tree for the stock market investor.

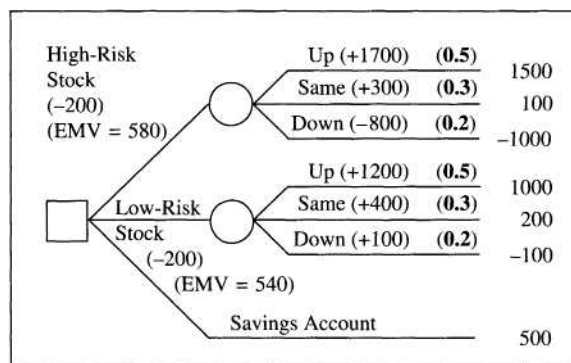


Table 13.2
Utility function for the investment problem.

Dollar Value	Utility Value
1500	1.00
1000	0.86
500	0.65
200	0.52
100	0.46
-100	0.33
-1000	0.00

As a result, an expected-value maximizer would choose the high-risk stock.

Figure 13.5 shows the investor's decision tree with the utility values instead of the payoffs. Solving this decision tree, we calculate the expected utility (EU) for the three investments:

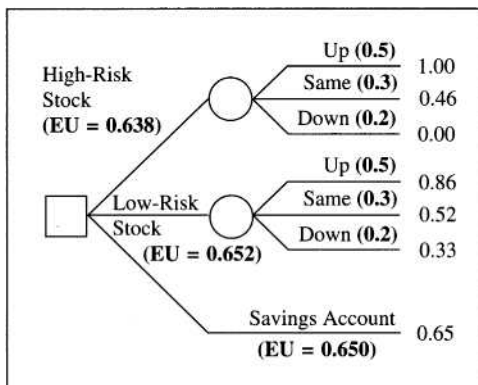
$$EU (\text{High-Risk Stock}) = 0.638$$

$$EU (\text{Low-Risk Stock}) = 0.652$$

$$EU (\text{Savings Account}) = 0.650$$

Now the preferred action is to invest in the low-risk stock because it provides the highest expected utility, although it does not differ much from that for the savings account. You can see how the expected utilities make it possible to rank these investments in order of preference. According to the utility function we are using, this investor dislikes risk enough to find the high-risk stock the least preferred of his three alternatives.

Figure 13.5
Decision tree for stock market investor — utility values instead of dollars.



Expected Utility, Certainty Equivalents, and Risk Premiums

Two concepts are closely linked to the idea of expected utility. One is that of a *certainty equivalent*, or the amount of money that is equivalent in your mind to a given situation that involves uncertainty. For example, suppose you face the following gamble:

Win \$2000 with probability 0.50

Lose \$20 with probability 0.50

Now imagine that one of your friends is interested in taking your place. "Sure," you reply, "I'll sell it to you." After thought and discussion, you conclude that the least you would sell your position for is \$300. If your friend cannot pay that much, then you would rather keep the gamble. (Of course, if your friend were to offer more, you would take it!)

Your certainty equivalent for the gamble is \$300. This is a sure thing; no risk is involved. From this, the meaning of certainty equivalent becomes clear. If \$300 is the least that you would accept for the gamble, then the gamble must be equivalent in your mind to a sure \$300.

In the example at the beginning of the chapter, Harriman decided that he would pay the additional \$600,000 simply to avoid the riskiness of the situation. His thinking at the time was that committing the additional money would ensure his control of the Northern Pacific Railroad. He indicated that he did not want to pay more, and so we can think of \$600,000 as his certainty equivalent for the gamble of purchasing the shares more slowly and risking detection.

Let us again consider the stock market investor. We can make certain inferences about his certainty equivalent for the gambles represented by the low-risk and high-risk stocks because we have information about his utility function. For example, his expected utility for the low-risk stock is 0.652, which is just a shade more than $U(\$500) = 0.650$. Thus, his certainty equivalent for the low-risk stock must be only a little more than \$500. Likewise, his expected utility for the high-risk stock is 0.638, which is somewhat less than 0.650. Therefore, his certainty equivalent for the high-risk stock must be less than \$500 but not as little as \$200, which has a utility of 0.520.

You can see also that we can rank the investments by their certainty equivalents. The high-risk stock, having the lowest certainty equivalent, is the least preferred. The low-risk stock, on the other hand, has the highest certainty equivalent, and so is the most preferred. Ranking alternatives by their certainty equivalents is the same as ranking them by their expected utilities. If two alternatives have the same certainty equivalent, then they must have the same expected utility, and the decision maker would be indifferent to a choice between the two.

Closely related to the idea of a certainty equivalent is the notion of *risk premium*. The risk premium is defined as the difference between the EMV and the certainty equivalent:

$$\text{Risk Premium} = \text{EMV} - \text{Certainty Equivalent}$$

Consider the gamble between winning \$2000 and losing \$20, each with probability 0.50. The EMV of this gamble is \$990. On reflection, you assessed your certainty equivalent to be \$300, and so your risk premium is

$$\begin{aligned} \text{Risk Premium} &= \$990 - \$300 = \\ & \$690 \end{aligned}$$


Because you were willing to trade the gamble for \$300, you were willing to "give up" \$690 in expected value in order to avoid the risk inherent in the gamble. You can think of the risk premium as the premium you pay (in the sense of a lost opportunity) to avoid the risk.

Figure 13.6 graphically ties together utility functions, certainty equivalents, and risk premiums. Notice that the certainty equivalent and the expected utility of a gamble are points that are "matched up" by the utility function. That is,

$$EU(\text{Gamble}) = U(\text{Certainty Equivalent})$$

In words, the utility of the certainty equivalent is equal to the expected utility of the gamble. Because these two quantities are equal, the decision maker must be indifferent to the choice between them. After all, that is the meaning of certainty equivalent. Now we can put all of the pieces together in Figure 13.6. Imagine a gamble has expected utility Y . The value Y is in utility units, so we must first locate Y on the vertical axis. Trace a horizontal line from the expected utility point until the line intersects the utility curve. Now drop down to the horizontal axis to find the certainty equivalent. The difference between the expected value and the certainty equivalent is the risk premium.

Figure 13.6
Graphical representation of risk premium.



For a risk-averse individual, the horizontal EU line reaches the concave utility curve before it reaches the vertical line that corresponds to the expected value. Thus, for a risk-averse individual the risk premium must be positive. If the utility function were convex, the horizontal EU line would reach the expected value before the utility curve. The certainty equivalent would be greater than the expected value, and so the risk premium would be negative. This would imply that the decision maker would have to be paid to give up an opportunity to gamble.

In any given situation, the certainty equivalent, expected value, and risk premium all depend on two factors: the decision maker's utility function and the probability distribution for the payoffs. The values that the payoff can take combine with the probabilities to determine the EMV. The utility function, coupled with the probability distribution, determines the expected utility and hence the certainty equivalent. The degree to which the utility curve is nonlinear determines the distance between the certainty equivalent and the expected payoff.

If the certainty equivalent for a gamble is assessed directly, then finding the risk premium is straightforward — simply calculate the EMV of the gamble and subtract the assessed certainty equivalent. In other cases, the decision maker may have assessed a utility function and now faces a particular gamble that he or she wishes to analyze. If so, there are four steps in finding the gamble's risk premium:

- 1 Find the EU for the gamble.
- 2 Find the certainty equivalent, or the sure amount that has the utility value equal to the EU that was found in Step 1.
- 3 Calculate the EMV for the gamble.
- 4 Subtract the certainty equivalent from the expected payoff to find the risk premium. This is the difference between the expected value of the risky situation and the sure amount for which the risky situation would be traded.

Here is a simple example. Using the hypothetical utility function given in Figure 13.7, we will find the risk premium for the following gamble:

Win \$4000 with probability	0.40
Win \$2000 with probability	0.20
Win \$0 with probability	0.15
Lose \$2000 with probability	0.25

The first step is to find the expected utility:

$$\begin{aligned}
 EU &= 0.40 U(\$4000) + 0.20 U(\$2000) + 0.15 U(\$0) + 0.25 U(-\$2000) \\
 &= 0.40(0.90) + 0.20(0.82) + 0.15(0.67) + 0.25(0.38) = 0.72
 \end{aligned}$$

The second line is simply a matter of estimating the utilities from Figure 13.7 and substituting them into the equation.

Figure 13.7
Utility function for a risk-averse decision maker.





Figure 13.8
Finding a certainty equivalent



For Step 2, the certainty equivalent is the sure amount that gives the same utility as the expected utility of the gamble. Figure 13.8 shows the process of finding the certainty equivalent for the gamble that has $EU = 0.72$. We start at the vertical axis with the utility value of 0.72, read across to the utility curve, and then drop down to the horizontal axis. From Figure 13.8, we can see that the certainty equivalent is approximately \$400.

Step 3 calculates the expected payoff or EMV:

$$EMV = 0.40(\$4000) + 0.20(\$2000) + 0.15(\$0) + 0.25(-\$2000) = \$1500$$

Finally, in Step 4, we calculate the risk premium by subtracting the certainty equivalent from the EMV:

$$\text{Risk Premium} = \$1500 - \$400 = \$1100$$

Keeping Terms Straight

One problem that students often have is that of confusion about the terms we use in utility theory. The basic idea, remember, is to use a utility function to translate dollars into utility units. If we compare two risky projects on the basis of expected utility (EU), we are working in utility units. When we calculate certainty equivalents or risk premiums, however, we are working in dollars. Thus, a certainty equivalent is not the same as the expected utility of a gamble. The two measurements provide equivalent information, but only in the sense that the certainty equivalent for a gamble is the sure amount that gives the same utility as the expected utility of the gamble. The translation from certainty equivalent to expected utility and back again is through the utility function, as depicted in Figures 13.6 and 13.8. Again, a certainty equivalent is a dollar amount, whereas expected utility is in utility units. Be careful to use these terms consistently.

Utility Function Assessment

Different people have different risk attitudes and thus are willing to accept different levels of risk. Some are more prone to taking risks, while others are more conservative and avoid risk. Thus, assessing a utility function is a matter of subjective judgment, just like assessing subjective probabilities. In this section we will look at two utility-assessment approaches that are based on the idea of certainty equivalents. The following section introduces an alternative approach.

It is worth repeating at this point our credo about modeling and decision making. Remember that the objective of the decision-analysis exercise is to help you make a better decision. To do this, we construct a model, or representation, of the decision. When we assess a utility function, we are constructing a mathematical model or representation of preferences. This representation then is included in the overall model of the decision problem and is used to analyze the situation at hand. The objective is to find a way to represent preferences that incorporates risk attitudes. A perfect representation is not necessary. All that is required is a model that represents feelings about risk well enough to understand and analyze the current decision.

Assessment Using Certainty Equivalents

The first assessment method requires the decision maker to assess several certainty equivalents. Suppose you face an uncertain situation in which you may have \$10 in the worst case, \$100 in the best case, or possibly something in between. You have a variety of options, each of which leads to some uncertain payoff between \$10 and \$100. To evaluate the alternatives, you must assess your utility for payoffs from \$10 to \$100.

We can get the first two points of your utility function by arbitrarily setting $U(100) = 1$ and $U(10) = 0$. This may seem a bit strange but is easily explained. The idea of the utility function, remember, is to rank-order risky situations. We can always take any utility function and rescale it — add a constant and multiply by a positive constant — so that the best outcome has a utility of 1 and the worst has a utility of 0. The rank ordering of risky situations in terms of expected utility will be the same for both the original and rescaled utility functions. What we are doing here is taking advantage of this ability to rescale. We are beginning the assessment process by setting two utility points. The remaining assessments then will be consistent with the scale set by these points. (We could just as well set the endpoints at 100 and 0, or 100 and -50, say. We are using 1 and 0 because this choice of endpoints turns out to be particularly convenient. But we will have to be careful not to confuse these utilities with probabilities!)

Now imagine that you have the opportunity to play the following lottery, which we will call a *reference lottery*:


Win \$100 with probability 0.5

Win \$10 with probability 0.5

What is the minimum amount for which you would be willing to sell your opportunity to play this game? \$25? \$30? Your job is to find your certainty equivalent (CE) for this reference gamble. A decision tree for your choice is shown in Figure 13.9.

Finding your certainty equivalent is where your subjective judgment comes into play. The CE undoubtedly will vary from person to person. Suppose that for this reference gamble your certainty equivalent is \$30. That is, for \$31 you would take the money, but for \$29 you would rather play the lottery; \$30 must be your true indifference point.

Figure 13.9
A "reference gamble"
for assessing a utility
function. Your job is to
find the certainty
equivalent (CE) so that
you are indifferent to
options A and B.



The key to the rest of the analysis is this: *Because you are indifferent between \$30 and the risky gamble, the utility of \$30 must equal the expected utility of the gamble.* We know the utilities of \$10 and \$100, so we can figure out the expected utility of the gamble:

$$\begin{aligned} U(30) &= 0.5 U(100) + 0.5 U(10) \\ &= 0.5(1) + 0.5(0) = 0.5 \end{aligned}$$

We have found a third point on your utility curve. To find another, take a different reference lottery:

Win \$100 with probability 0.5
Win \$30 with probability 0.5

Now find your certainty equivalent for this new gamble. Again, the certainty equivalent will vary from person to person, but suppose that you settle on \$50. We can do exactly what we did before, but with the new gamble. In this case, we can find the utility of \$50 because we know $U(100)$ and $U(30)$ from the previous assessment.

$$\begin{aligned} U(50) &= 0.5 U(100) + 0.5 U(30) \\ &= 0.5(1) + 0.5(0.5) = \\ &0.75 \end{aligned}$$

This is the fourth point on your utility curve.

Now consider the reference lottery:

Win \$30 with probability 0.5
Win \$10 with probability 0.5

Again you must assess your CE for this gamble. Suppose it turns out to be \$18. Now we can do the familiar calculations:

$$\begin{aligned} U(18) &= 0.5 U(30) + 0.5 U(10) = \\ &0.5(0.5) + 0.5(0) = 0.25 \end{aligned}$$

We now have five points on your utility curve, and we can graph and draw a curve through them. The graph is shown in Figure 13.10. A smooth curve drawn through the assessed points should be an adequate representation of your utility function for use in solving your decision problem.

Assessment Using Probabilities

The CE approach requires that you find a dollar amount that makes you indifferent between the gamble and the sure thing in Figure 13.9. Another approach involves setting the sure amount in Alternative B and adjusting the probability in the reference gamble to achieve indifference. We will call this the *probability-equivalent* (PE) assessment technique.

For example, suppose you want to know your utility for \$65. This is not one of the certainty equivalents that you assessed, and thus $U(65)$ is unknown. You could make an educated guess. Based on the previous assessments and the graph in Figure 13.10, $U(65)$ must be between 0.75 and 1.00; it probably is around 0.85. But rather than guess, you can assess the value directly. Consider the reference lottery:

Win \$100 with probability p Win
 \$10 with probability $(1 - p)$

This gamble is shown in Figure 13.11.

To find your utility value for \$65, adjust p until you are indifferent between the sure \$65 and the reference gamble. That is, think about various probabilities that make the chance of winning \$100 greater or less until you are indifferent between Alternatives C and D in Figure 13.11. Now you can find $U(65)$ because you know that $U(100) = 1$ and $U(10) = 0$:

$$\begin{aligned} U(65) &= p U(100) + (1 - p) U(10) \\ &= p(1) + (1 - p)(0) \\ &= p \end{aligned}$$

Figure 13.10
 Graph of the utility function assessed using the certainty-equivalent approach.





Figure 13.11
 A reference gamble for assessing the utility of \$65 using the probability-equivalent method.



The probability that makes you indifferent just happens to be your utility value for \$65. For example, if you chose $p = 0.87$ to achieve indifference, then $U(65) = 0.87$.

Gambles, Lotteries, and Investments

As with the assessment of subjective probabilities, we have framed utility assessment in terms of gambles or lotteries. For many individuals this evokes images of carnival games or gambling in a casino, images that may seem irrelevant to the decision at hand or even distasteful. An alternative is to think in terms of investments that are risky. Instead of considering if you would accept or reject a particular gamble, think about whether you would make a particular investment. Would you agree to invest in a security with a p chance of yielding \$1000 and a $1 - p$ chance of yielding nothing? How much would you be willing to pay for such a security? Framing your utility-assessment questions in this way may help you to think more clearly about your utility, especially for investment decisions.

Risk Tolerance and the Exponential Utility Function

The assessment process just described works well for assessing a utility function subjectively, and it can be used in any situation, although it can involve a fair number of assessments. An alternative approach is to base the assessment on a particular mathematical function, such as one of those that we introduced early in the chapter. In particular, let us consider the *exponential utility function*:

$$U(x) = 1 - e^{-x/R}$$


This utility function is based on the constant $e = 2.71828 \dots$, the base of natural logarithms. This function is concave and thus can be used to represent risk-averse preferences. As x becomes large, $U(x)$ approaches 1. The utility of zero, $U(0)$, is equal to 0, and the utility for negative x (being in debt) is negative.

In the exponential utility function, R is a parameter that determines how risk-averse the utility function is. In particular, R is called the *risk tolerance*. Larger values of R make the exponential utility function flatter, while smaller values make it more concave or more risk-averse. Thus, if you are less risk-averse — if you can tolerate more risk — you would assess a larger value for R to obtain a flatter utility function. If you are less tolerant of risk, then you would assess a smaller R and have a more curved utility function.

How can R be determined? A variety of ways exist, but it turns out that R has a very intuitive interpretation that makes its assessment relatively easy. Consider the gamble

Win $\$Y$ with probability 0.5 Lose
 $\$Y/2$ with probability 0.5

Figure 13.12
Assessing your risk tolerance. Find the largest value of Y for which you would prefer Alternative E.



Would you be willing to take this gamble if Y were \$100? \$2000? \$35,000? Or, framing it as an investment, how much would you be willing to risk ($\$Y/2$) in order to have a 50% chance of tripling your money (winning $\$Y$ and keeping your $\$Y/2$)? At what point would the risk become intolerable? The decision tree is shown in Figure 13.12.

The largest value of Y for which you would prefer to take the gamble rather than not take it is approximately equal to your risk tolerance. This is the value that you can use for R in your exponential utility function. For example, suppose that after considering the decision tree in Figure 13.12 you conclude that the largest Y for which you would take the gamble is $Y = \$900$. Hence, $R = \$900$. Using this assessment in the exponential utility function would result in the utility function

$$U(x) = 1 - e^{-x/900}$$

This exponential utility function provides the translation from dollars to utility units. Once you have your R value and your exponential utility function, it is fairly easy to find certainty equivalents. For example, suppose that you face the following gamble:

- Win \$2000 with probability 0.4
- Win \$1000 with probability 0.4
- Win \$500 with probability 0.2

The expected utility for this gamble is

$$\begin{aligned} EU &= 0.4 U(\$2000) + 0.4 U(\$1000) + 0.2 U(\$500) \\ &= 0.4(0.8916) + 0.4(0.6708) + 0.2(0.4262) = \\ &0.7102 \end{aligned}$$

To find the CE we must work backward through the utility function. We want to find the value x such that $U(x) = 0.7102$. Set up the equation

$$0.7102 = 1 - e^{-x/900}$$

Subtract 1 from each side to get

$$-0.2898 = -e^{-x/900}$$

Multiply through to eliminate the minus signs:

$$0.2898 = e^{-x/900}$$

Now we can take natural logs of both sides to eliminate the exponential term:

$$\ln(0.2898) = \ln(e^{-x/900}) = \frac{-x}{900}$$

The rule (from algebra) is that $\ln(e^y) = y$. Now we simply solve for x :

$$\begin{aligned} \ln(0.2898) &= \frac{-x}{900} \\ x &= -900[\ln(0.2898)] \\ &= \$1114.71 \end{aligned}$$

The procedure above requires that you use the exponential utility function to translate the dollar outcomes into utilities, find the expected utility, and finally convert to dollars to find the exact certainty equivalent. That can be a lot of work, especially if there are many outcomes to consider. Fortunately, an approximation is available from Pratt (1964) and also discussed in McNamee and Celona (1987). Suppose you can figure out the expected value and variance of the payoffs. Then the CE is approximately

$$\text{Certainty Equivalent} \approx \text{Expected Value} - \frac{0.5(\text{Variance})}{\text{Risk Tolerance}}$$

In symbols,

$$\text{CE} \approx \mu - \frac{0.5\sigma^2}{R}$$

where μ and σ^2 are the expected value and variance, respectively. For example, in the gamble above, the expected value (EMV or μ) equals \$1300, and the standard deviation (σ) equals \$600. Thus, the approximation gives

$$\begin{aligned} \text{CE} &\approx \$1300 - \frac{0.5(\$600)^2}{900} \\ &\approx \$1100 \end{aligned}$$

The approximation is within \$15. That's pretty good! This approximation is especially useful for continuous random variables or problems where the expected value and variance are relatively easy to estimate or assess compared to assessing the entire probability distribution. The approximation will be closest to the actual value when the outcome's probability distribution is a symmetric, bell-shaped curve.

What are reasonable R values? For an individual's utility function, the appropriate value for R clearly depends on the individual's risk attitude. As indicated, the less risk-averse a person is, the larger R is. Suppose, however, that an individual or a group (a board of directors, say) has to make a decision on behalf of a corporation. It is important that these decision makers adopt a decision-making attitude based on corporate goals and acceptable risk levels for the corporation. This can be quite different from an individual's personal risk attitude; the individual director may be

unwilling to risk \$10 million, even though the corporation can afford such a loss. Howard (1988) suggests certain guidelines for determining a corporation's risk tolerance in terms of total sales, net income, or equity. Reasonable values of R appear to be approximately 6.4% of total sales, 1.24 times net income, or 15.7% of equity. These figures are based on observations that Howard has made in the course of consulting with various companies. More research may refine these figures, and it may turn out that different industries have different ratios for determining reasonable R 's.

Using the exponential utility function seems like magic, doesn't it? One assessment, and we are finished! Why bother with all of those certainty equivalents that we discussed above? You know, however, that you never get something for nothing, and that definitely is the case here. The exponential utility function has a specific kind of curvature and implies a certain kind of risk attitude. This risk attitude is called *constant risk aversion*. Essentially it means that no matter how much wealth you have — how much money in your pocket or bank account — you would view a particular gamble in the same way. The gamble's risk premium would be the same no matter how much money you have. Is constant risk aversion reasonable? Maybe it is for some people. Many individuals might be less risk-averse if they had more wealth.

In later sections of this chapter we will study the exponential utility function in more detail, especially with regard to constant risk aversion. The message here is that the exponential utility function is most appropriate for people who really believe that they would view gambles the same way regardless of their wealth level. But even if this is not true for you, the exponential utility function can be a useful tool for modeling preferences. The next section shows how sensitivity analysis can be performed in terms of risk tolerance.

Risk Tolerance and Sensitivity Analysis: Eagle Airlines, Revisited

In Chapter 5 we learned how to use sensitivity analysis to help with probability assessment. We also can use sensitivity analysis to deal with preferences. The idea of risk tolerance suggests a tidy way to perform the analysis. Essentially, we will vary the risk tolerance in the exponential utility function to determine at what point the decision changes.

Let us reconsider the Eagle Airlines case that we used for sensitivity analysis in Chapter 5 (pages 155-156). Recall that the question was whether Dick Carothers should purchase an additional aircraft and expand Eagle Airlines' operations. After considerable modeling, we worked the problem down to the consideration of three uncertain variables: the capacity of the scheduled flights (proportion of seats sold), the operating cost, and the total number of hours flown during a year. Figure 13.13 shows the decision tree with specific probabilities on the branches. Using these probabilities, we calculate the EMV for purchasing the airplane as \$9644, which is considerably more than \$4200. It also is true, however, that purchasing the airplane is

substantially riskier than investing in the money market. If Carothers is risk-averse, just how risk-averse can he be and still be willing to purchase the airplane? Or in terms of his risk tolerance, how low could his risk tolerance be before he would rather invest in the money market?

Answering this question is conceptually simple, but the calculations are tedious. We will use the exponential utility function. Conceptually, all we want to do is find the critical value for R such that the expected utility of purchasing the airplane equals the expected utility of investing in the money market. Then the question is whether Carothers's risk tolerance is above or below the critical value. If it is above the critical value, then he should buy the airplane; if it is less, then he should go with the less risky money market.

What is the easiest way to find the critical value? The simplest is to use a computer and search for the R value that makes the two expected utilities equal. Decision-analysis programs such as DATA and DPL allow the user to specify a value for R , the risk tolerance. In this case, it is a simple matter to try different values for R until the critical one is found.

An electronic spreadsheet provides a convenient environment for sensitivity analysis with respect to risk tolerance. The expected utility calculations can be programmed easily into the cells of the spreadsheet. Figure 13.14 shows a Microsoft Excel spreadsheet with a table created to perform a sensitivity analysis on the risk tolerance. The decision tree, which is not shown but is in the upper part of the spreadsheet, has been created using macro commands to enter the decision and

Figure 13.13
Eagle Airlines'
decision tree.

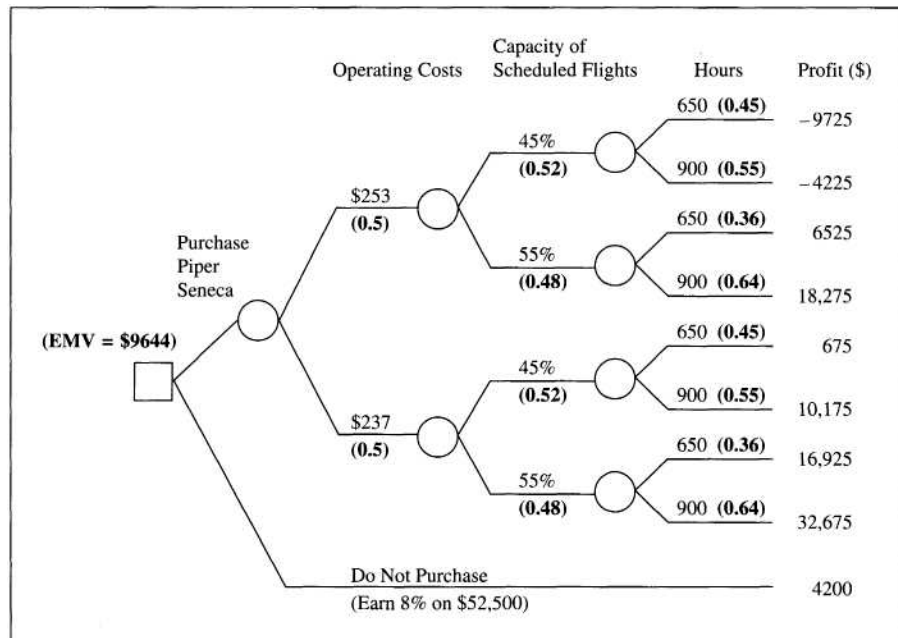


Figure 13.14
Using a spreadsheet to find the critical value for R in the Eagle Airlines problem.

	A	B	C	D
16	EAGLE AIRLINES RISK SENSITIVITY			
17				
18		R= 14240		
19				
20		M		
21	14200	M		
22	14210	M		
23	14220	M		
24	14230	M		
25	14240	M		
26	14250	S		
27	14260	S		
28	14270	S		
29	14280	S		
30	14290	S		
31	14300	S		
32	14310	S		
33	14320	S		

chance nodes shown in Figure 13.13. Cells B5 and B6 (not shown) contain the final calculations of expected utility for purchasing the Piper Seneca and the utility for earning \$4200 in the money market, each case using the R value specified in Cell B18. In Cell B20, a formula reports which of the two alternatives has the greater expected utility. A table in Cells A20-B33 has been created in which the R values along the right-hand side are used in the expected utility calculations. This table is created by selecting the block of cells, choosing the "Table" command, and specifying B18 as the column-input cell. It may be worthwhile to recalculate the table several times, each time using a narrower range of R values. In the first calculation of the table, using R values over a very wide range helps to find the approximate location of the critical R . Then the table can be calculated several times, each using a finer grid to home in on the precise critical value. As you can see in Figure 13.14, the table checks R values that differ by only \$10.

As Figure 13.14 also shows, the critical R value is approximately \$14,240. Now Carothers must ask whether he would be willing to accept an investment in which he would have the same chance of winning \$14,240 or losing \$7120. If he concludes that he would not take this gamble, then his risk tolerance must be smaller than \$14,240, and he should not purchase the airplane. But if he would participate gladly in this gamble, then his risk tolerance must be greater than \$14,240, and he should buy the airplane.

Decreasing and Constant Risk Aversion (Optional)

In this section we will consider how individuals might deal with risky investments. Suppose you had \$1000 to invest. How would you feel about investing \$500 in an

extremely risky venture in which you might lose the entire \$500? Now suppose you have saved more money and have \$20,000. Now how would you feel about that extremely risky venture? Is it more or less attractive to you? How do you think a person's degree of risk aversion changes with wealth?

Decreasing Risk Aversion

If an individual's preferences show *decreasing risk aversion*, then the risk premium decreases if a constant amount is added to all payoffs in a gamble. Expressed informally, decreasing risk aversion means the more money you have, the less nervous you are about a particular bet.

For example, suppose an individual's utility curve can be described by a logarithmic function:

$$U(x) = \ln(x)$$

where x is the wealth or payoff and $\ln(x)$ is the natural logarithm of x . Using this logarithmic utility function, consider the gamble

Win \$10 with probability 0.5

Win \$40 with probability 0.5

To find the certainty equivalent, we first find the expected utility. The utility values for \$10 and \$40 are

$$U(\$10) = \ln(10) = 2.3026$$

$$U(\$40) = \ln(40) = 3.6889$$

Calculating expected utility:

$$EU = 0.5(2.3026) + 0.5(3.6889) = 2.9957$$

To find the certainty equivalent, you must find the certain value x that has $U(x) = 2.9957$; thus, set the utility function equal to 2.9957:

$$2.9957 = \ln(x) \text{ Now solve}$$

for x . To remove the logarithm, we take antilogs:

$$e^{2.9957} = e^{\ln(x)} = x$$

The rule here corresponds to what we did with the exponential function. Here we have $e^{\ln(y)} = y$. Finally, we simply calculate $e^{2.9957}$:

$$x = e^{2.9957} = \$20 = CE$$

To find the risk premium, we need the expected payoff, which is

$$EMV = 0.5(\$10) + 0.5(\$40) = \$25$$

Thus, the risk premium is $EMV - CE = \$25 - \$20 = \$5$.

Using the same procedure, we can find risk premiums for the lotteries as shown in Table 13.3. Notice that the sequence of lotteries is constructed so that each is like having the previous one plus \$10. For example, the \$20-\$50 lottery is like having the \$10-\$40 lottery plus a \$10 bill. The risk premium decreases with each \$10 addition. The decreasing risk premium reflects decreasing risk aversion, which is a property of the logarithmic utility function.

An Investment Example

For another example, suppose that an entrepreneur is considering a new business investment. To participate, the entrepreneur must invest \$5000. There is a 25% chance that the investment will earn back the \$5000, leaving her just as well off as if she had not made the investment. But there is also a 45% chance that she will lose the \$5000 altogether, although this is counterbalanced by a 30% chance that the investment will return the original \$5000 plus an additional \$10,000. Figure 13.15 shows the entrepreneur's decision tree.

We will assume that this entrepreneur's preferences can be modeled with the logarithmic utility function $U(x) = \ln(x)$, where x is interpreted as total wealth. Suppose that the investor now has \$10,000. Should she make the investment or avoid it?

The easiest way to solve this problem is to calculate the expected utility of the investment and compare it with the expected utility of the alternative, which is to do nothing. The expected utility of doing nothing simply is the utility of the current wealth, or $U(10,000)$, which is

$$U(10,000) = \ln(10,000) = 9.2103$$

The expected utility of the investment is easy to calculate:


$$\begin{aligned} EU &= 0.30 U(20,000) + 0.25 U(10,000) + 0.45 U(5000) = \\ &0.30(9.9035) + 0.25(9.2103) + 0.45(8.5172) = 9.1064 \end{aligned}$$

Because the expected utility of the investment is less than the utility of not investing, the investment should not be made.

Table 13.3
Risk premiums from
logarithmic utility
function.

50-50 Gamble Between (\$)	Expected Value (\$)	Certainty Equivalent (\$)	Risk Premium (\$)
10, 40	25	20.00	5.00
20, 50	35	31.62	3.38
30, 60	45	42.43	2.57
40, 70	55	52.92	2.08

Figure 13.15
Entrepreneur's investment decision. Current wealth is denoted by x .



Now, suppose that several years have passed. The utility function has not changed, but other investments have paid off handsomely, and she currently has \$70,000. Should she undertake the project now? Recalculating with a base wealth of \$70,000 rather than \$10,000, we find that the utility of doing nothing is $U(70,000) = 11.1563$, and the EU for the investment is 11.1630. Now the expected utility of the investment is greater than the utility of doing nothing, and so she should invest.

The point of these examples is to show how decreasing risk aversion determines the way in which a decision maker views risky prospects. As indicated, the wealthier a decreasingly risk-averse decision maker is, the less anxious he or she will be about taking a particular gamble. Generally speaking, decreasing risk aversion makes sense when we think about risk attitudes and the way that many people appear to deal with risky situations. Many would feel better about investing money in the stock market if they were wealthier to begin with. For such reasons, the logarithmic utility function is commonly used by economists and decision theorists as a model of typical risk attitudes.

Constant Risk Aversion

An individual displays constant risk aversion if the risk premium for a gamble does not depend on the initial amount of wealth held by the decision maker. Intuitively, the idea is that a constantly risk-averse person would be just as anxious about taking a bet regardless of the amount of money available.

If an individual is constantly risk-averse, the utility function is exponential. It would have the following form:

$$U(x) = 1 - e^{-x/R}$$

For example, suppose that the decision maker has assessed a risk tolerance of \$35:

$$U(x) = 1 - e^{-x/35}$$

We can perform the same kind of analysis that we did with the logarithmic utility function above. Consider the gamble

Win \$10 with probability 0.5

Win \$40 with probability 0.5

As before, the expected payoff is \$25. To find the CE, we must find the expected utility, which requires plugging the amounts \$10 and \$40 into the utility function:

$$U(10) = 1 - e^{-10/35} = 0.2485$$

$$U(40) = 1 - e^{-40/35} = 0.6811$$

Thus, $EU = 0.5 (0.2485) + 0.5 (0.6811) = 0.4648$. To find the certainty equivalent, set the utility function to 0.4648. The value for x that gives the utility of 0.4648 is the gamble's CE:

$$0.4648 = 1 - e^{-x/35}$$

Now we can solve for x as we did earlier when working with the exponential utility function:

$$0.5352 = e^{-x/35}$$

$$\ln(0.5352) = -0.6251 = -x/35$$

$$x = 0.6251(35) = \$21.88 = CE$$

Finally, the expected payoff (EMV) is \$25, and so the risk premium is

$$\text{Risk Premium} = EMV - CE = \$25 - \$21.88 = \$3.12$$

Using the same procedure, we can find the risk premium for each gamble in Table 13.4. The risk premium stays the same as long as the difference between the payoffs does not change. Adding a constant amount to both sides of the gamble does not change the decision maker's attitude toward the gamble.

Alternatively, you can think about this as a situation where you have a bet in which you may win \$15 or lose \$15. In the first gamble above, you face this bet with \$25 in your pocket. In the constant-risk-aversion situation, the way you feel about the bet (as reflected in the risk premium) is the same regardless of how much money is added to your pocket. In the decreasing-risk-aversion situation, adding something to your pocket made you less risk-averse toward the bet, thus resulting in a lower risk premium.

Table 13.4 Risk premiums from exponential utility function.	50-50 Gamble Between (\$)	Expected Value (\$)	Certainty Equivalent (\$)	Risk Premium (\$)
	10, 40	25	21.88	3.12
	20, 50	35	31.88	3.12
	30, 60	45	41.88	3.12
	40, 70	55	51.88	3.12


Figure 13.16 plots the two utility functions on the same graph. They have been rescaled so that $U(10) = 0$ and $U(100) = 1$ in each case. Note their similarity. It does not take a large change in the utility curve's shape to alter the nature of the individual's risk attitude.

Is constant risk aversion appropriate? Consider the kinds of risks that railroad barons such as E. H. Harriman undertook. Would a person like Harriman have been willing to risk millions of dollars on the takeover of another railroad had he not already been fairly wealthy? The argument easily can be made that the more wealth one has, the easier it is to take larger risks. Thus, decreasing risk aversion appears to provide a more appropriate model of preferences than does constant risk aversion. This is an important point to keep in mind if you decide to use the exponential utility function and the risk-tolerance parameter; this utility function displays constant risk aversion.

After all is said and done, while the concepts of decreasing or constant risk aversion may be intriguing from the point of view of a decision maker who is interested in modeling his or her risk attitude, precise determination of a decision maker's utility function is not yet possible. Decision theorists still are learning how to elicit and measure utility functions. Many unusual effects arise from human nature; we will study some of these in the next chapter. It would be an overstatement to suggest that it is possible to determine precisely the degree of an individual's risk aversion or whether he or she is decreasingly risk-averse. It is a difficult enough problem just to determine whether someone is risk-averse or risk-seeking!

Thus, it may be reasonable to use the exponential utility function as an approximation in modeling preferences and risk attitudes. A quick assessment of risk tolerance,

Figure 13.16
Logarithmic and exponential utility functions plotted over the range of \$10 to \$100.



and you are on your way. And if even that seems a bit strained, then it is always possible to use the sensitivity-analysis approach; it may be that a precise assessment of the risk tolerance is not necessary.

Some Caveats

A few things remain to be said about utilities. These are thoughts to keep in mind as you work through utility assessments and use utilities in decision problems.

- 1 Utilities do not add up. That is, $U(A + B) \neq U(A) + U(B)$. This actually is the whole point of having a nonlinear utility function. Thus, when using utilities in a decision analysis, you must calculate net payoffs or net contributions at the end-points of the decision tree before transforming to utility values.
- 2 Utility differences do not express strength of preferences. Suppose that $U(A_1) - U(A_2) > U(A_3) - U(A_4)$. This does not necessarily mean that you would rather go from A_2 to A_1 instead of from A_4 to A_3 . Utility only provides a numerical scale for ordering preferences, not a measure of their strengths. Whether this is reasonable is a matter of some debate. For example, von Winterfeldt and Edwards (1986) give the following example: You are first told that you will receive \$100, and then told that you actually will receive \$500, and then finally told that the actual payment will be \$10,000. It would indeed be a pleasant surprise to go from \$100 to \$500, but for most of us, the delight we would experience from \$500 to \$10,000 would eclipse the difference between \$100 and \$500. Von Winterfeldt and Edwards argue that we can make judgments of just this sort. Whether one agrees or not, it is necessary to interpret utility carefully in this regard.
- 3 Utilities are not comparable from person to person. A utility function is a subjective personal statement of an individual's preferences and so provides no basis for comparing utilities among individuals.

SUMMARY

In this chapter we have explored some basic concepts that underlie risk and return trade-offs, with the aim of being able to understand how to model a decision maker's risk preferences. We discussed the notion of a risk premium ($EMV - CE$), which can be thought of as a measure of how risk-averse a decision maker is in regard to a particular risky situation. The basic procedure for assessing a utility function requires comparison of lotteries with riskless payoffs. Once a utility function has been determined, the procedure is to replace dollar payoffs in a decision tree or influence diagram with utility values and solve the problem to find the option with the greatest

expected utility. We also studied the exponential utility function and the notion of risk tolerance. Because of its nature, the exponential utility function is particularly useful for modeling preferences in decision analysis. The concepts of decreasing and constant risk aversion also were discussed.

EXERCISES

- 13.1** Why is it important for decision makers to consider their attitudes toward risk?
- 13.2** We have not given a specific definition of risk. How would you define it? Give examples of lotteries that vary in riskiness in terms of your definition of risk.
- 13.3** Explain in your own words the idea of a certainty equivalent.
- 13.4** Explain what is meant by the term *risk premium*.
- 13.5** Explain in your own words the idea of risk tolerance. How would it apply to utility functions other than the exponential utility function?
- 13.6** Suppose a decision maker has the utility function shown in Table 13.1. An investment opportunity has EMV \$1236 and expected utility 0.93. Find the certainty equivalent for this investment and the risk premium.
- 13.7** A decision maker's assessed risk tolerance is \$1210. Assume that this individual's preferences can be modeled with an exponential utility function.
- Find $U(\$1000)$, $U(\$800)$, $U(\$0)$, and $U(-\$1250)$.
 - Find the expected utility for an investment that has the following payoff distribution:

$$P(\$1000) = 0.33$$

$$P(\$800) = 0.21$$

$$P(\$0) = 0.33 \quad P(-\$1250) = 0.13$$
 - Find the exact certainty equivalent for the investment and the risk premium.
 - Find the approximate certainty equivalent using the expected value and variance of the payoffs.
 - Another investment possibility has expected value \$2400 and standard deviation \$300. Find the approximate certainty equivalent for this investment.
- 13.8** Many firms evaluate investment projects individually on the basis of expected value and at the same time maintain diversified holdings in order to reduce risk. Does this make sense in light of our discussion of risk attitudes in this chapter?
- 13.9** A friend of yours, who lives in Reno, has life insurance, homeowner's insurance, and automobile insurance and also regularly plays the quarter slot machines in the casinos. What kind of a utility function might explain this kind of behavior? How else might you explain such behavior?
- 13.10** Two risky gambles were proposed at the beginning of the chapter:

Game 1 Win \$30 with probability 0.5 Lose
\$1 with probability 0.5

Game 2 Win \$2000 with probability 0.5 Lose
\$1900 with probability 0.5

Many of us would probably pay to play Game 1 but would have to be paid to participate in Game 2. Is this true for you? How much would you pay (or have to be paid) to take part in either game?

QUESTIONS AND PROBLEMS

13.11 St. Petersburg Paradox. Consider the following game that you have been invited to play by an acquaintance who always pays his debts. Your acquaintance will flip a fair coin. If it comes up heads, you win \$2. If it comes up tails, he flips the coin again. If heads occurs on the second toss, you win \$4. If tails, he flips again. If heads occurs on the third toss, you win \$8, and if tails, he flips again, and so on. Your payoff is an uncertain amount with the following probabilities:

Payoff	Probability	
2	0.50	
4	0.25	
8	0.125	
•	•	
•	•	
•	•	(where n is the number
2^n	0.5^n	of the toss when the
•	•	first head occurs)
•	•	
•	•	

This is a good game to play because you are bound to come out ahead. There is no possible outcome from which you can lose. How much would you pay to play this game? \$10? \$20? What is the expected value of the game? Would you be indifferent between playing the game and having the expected value for sure?

13.12 Assess your utility function in two different ways.

- a Use the certainty-equivalent approach to assess your utility function for wealth over a range of \$100 to \$20,000.

- b** Use the probability-equivalent approach to assess $U(\$1500)$, $U(\$5600)$, $U(\$9050)$, and $U(\$13,700)$. Are these assessments consistent with the assessments made in part a? Plot these assessments and those from part a on the same graph and compare them.
- 13.13** Assess your risk tolerance (R). Now rescale your exponential utility function — the one you obtain by substituting your R value into the exponential utility function — so that $U(\$100) = 0$ and $U(\$20,000) = 1$. [That is, find constants a and b so that $a + b(1 - e^{-100/R}) = 0$ and $a + b(1 - e^{-20,000/R}) = 1$.] Now plot the rescaled utility function on the same graph with the utility assessments from Problem 13.12. How do your assessments compare?
- 13.14** Let us return to the Texaco-Pennzoil example from Chapter 4 and think about Liedtke's risk attitude. Suppose that Liedtke's utility function is given by the utility function in Table 13.5.
- a** Graph this utility function. Based on this graph, how would you classify Liedtke's attitude toward risk?
- b** Use the utility function in conjunction with the decision tree sketched in Figure 4.2 to solve Liedtke's problem. With these utilities, what strategy should he pursue? Should he still counteroffer \$5 billion? What if Texaco counteroffers \$3 billion? Is your answer consistent with your response to part a?
- c** Based on this utility function, what is the least amount (approximately) that Liedtke should agree to in a settlement? (*Hint:* Find a sure amount that gives him the same expected utility that he gets for going to court.) What does this suggest regarding plausible counteroffers that Liedtke might make?
- 13.15** Of course, Liedtke is not operating by himself in the Texaco-Pennzoil case; he must report to a board of directors. Table 13.6 gives utility functions for three different directors. Draw graphs of these. How would you classify each director in terms of his or her attitude toward risk? What would be the strategies of each? (That is, what would each one do with respect to Texaco's current offer, and how would each react to a Texaco counteroffer of \$3 billion? To answer this question, you must solve the decision tree — calculate expected utilities — for *each* director.)
- 13.16** How do you think Liedtke (Problem 13.14) and the directors in Problem 13.15 will be able to reconcile their differences?
- 13.17** Rescale the utility function for Director A in Problem 13.15 so that it ranges between 0 and 1. That is, find constants a and b so that when you multiply the utility function by a

Table 13.5
Utility function
for Liedtke.

Payoff (Billions)	Utility
10.3	1.00
5.0	0.75
3.0	0.60
2.0	0.45
0,0	0.00

Table 13.6
Utility functions for
three Pennzoil
directors.

Payoff	Utility		
	Director A	Director B	Director C
10.3	3.0	100	42.05
5.0	2.9	30	23.50
3.0	2.8	15	16.50
2.0	2.6	8	13.00
0.0	1.0	0	6.00

and then add b , the utility for \$10.30 is 1 and the utility for \$0 is 0. Graph the rescaled utility function and compare it to the graph of the original utility function. Use the rescaled utility function to solve the Texaco-Pennzoil decision tree. Is the optimal choice consistent with the one you found in Problem 13.15?

- 13.18** What if Hugh Liedtke were risk-averse? Based on Figure 4.2, find a critical value for Hugh Liedtke's risk tolerance. If his risk tolerance is low enough (very risk-averse), he would accept the \$2 billion offer. How small would his risk tolerance have to be for $EU(\text{Accept } \$2 \text{ Billion})$ to be greater than $EU(\text{Counteroffer } \$5 \text{ Billion})$?
- 13.19** The idea of dominance criteria and risk aversion come together in an interesting way, leading to a different kind of dominance. If two risky gambles have the same expected payoff, on what basis might a risk-averse individual choose between them without performing a complete utility analysis?
- 13.20** This problem is related to the ideas of dominance that we discussed in Chapters 4 and 8. Investment D below is said to show "second-order stochastic dominance" over Investment C. In this problem, it is up to you to explain why D dominates C.

You are contemplating two alternative uncertain investments, whose distributions for payoffs are as below.

Payoff	Probabilities	
	Investment C	Investment D
50	1/3	1/4
100	1/3	1/2
150	1/3	1/4

If your preference function is given by $U(x) = 1 - e^{-x/100}$, calculate EU for both C and D. Which would you choose?

Plot the CDFs for C and D on the same graph. How do they compare? Use the graph to explain intuitively why any risk-averse decision maker would prefer D. (*Hint:* Think about the concave shape of a risk-averse utility function.)

13.21 Utility functions need not relate to dollar values. Here is a problem in which we know little about five abstract outcomes. What is important, however, is that a person who does know what A to E represent should be able to compare the outcomes using the lottery procedures we have studied.

A decision maker faces a risky gamble in which she may obtain one of five outcomes. Label the outcomes A , B , C , D , and E . A is the most preferred, and E is least preferred. She has made the following three assessments.

- She is indifferent between having C for sure or a lottery in which she wins A with probability 0.5 or E with probability 0.5.
- She is indifferent between having B for sure or a lottery in which she wins A with probability 0.4 or C with probability 0.6.
- She is indifferent between these two lotteries:
 - 1: A 50% chance at B and a 50% chance at D
 - 2: A 50% chance at A and a 50% chance at E

What are $U(A)$, $U(B)$, $U(C)$, $U(D)$, and $U(E)$?

13.22 You have considered insuring a particular item of property (such as an expensive camera, your computer, or your Stradivarius violin), but after considering the risks and the insurance premium quoted, you have no clear preference for either purchasing the insurance or taking the risk. The insurance company then tells you about a new scheme called "probabilistic insurance." You pay half the above premium but have coverage only in the sense that in the case of a claim there is a probability of one-half that you will be asked to pay the other half of the premium and will be completely covered, or that you will not be covered and will have your premium returned. The insurance company can be relied on to be fair in flipping the coin to determine whether or not you are covered.

- a Do you consider yourself to be risk-averse?
- b Would you purchase probabilistic insurance?
- c Draw a decision tree for this problem.
- d Show that a risk-averse individual always should prefer the probabilistic insurance.

(Hint: This is a difficult problem. To solve it you must be sure to consider that you are indifferent between the regular insurance and no insurance. Write out the equation relating these two alternatives and see what it implies. Another strategy is to select a specific utility function — the log utility function $U(x) = \log(x)$, say — and then find values for the probability of a claim, your wealth, the insurance premium, and the value of your piece of property so that the utility of paying the insurance premium is equal to the expected utility of no insurance. Now use these values to calculate the expected utility of the probabilistic insurance. What is the result?)

13.23 An investor with assets of \$10,000 has an opportunity to invest \$5000 in a venture that is equally likely to pay either \$15,000 or nothing. The investor's utility function can be described by the log utility function $U(x) = \ln(x)$, where x is his total wealth.

- a What should the investor do?
- b Suppose the investor places a bet with a friend before making the investment decision. The bet is for \$1000; if a fair coin lands heads up, the investor wins \$1000, but

if it lands tails up, the investor pays \$1000 to his friend. Only after the bet has been resolved will the investor decide whether or not to invest in the venture. What is an appropriate strategy for the investor? If he wins the bet, should he invest? What if he loses the bet?

- c Describe a real-life situation in which an individual might find it appropriate to gamble before deciding on a course of action.

Source: D. E. Bell (1988) "Value of Pre-Decision Side Bets for Utility Maximizers." *Management Science*, 34, 797-800.

- 13.24** A bettor with utility function $U(x) = \ln(x)$, where x is total wealth, has a choice between the following two alternatives:

A Win \$ 10,000 with probability 0.2

Win \$1000 with probability 0.8

B Win \$3000 with probability 0.9 Lose

\$2000 with probability 0.1

- a** If the bettor currently has \$2500, should he choose A or B? **b**

Repeat a, assuming the bettor has \$5000. **c** Repeat a,

assuming the bettor has \$10,000.

- d** Do you think that this pattern of choices between A and B is reasonable? Why or why not?

Source: D. E. Bell (1988) "One-Switch Utility Functions and a Measure of Risk." *Management Science*, 34, 1416-1424.

- 13.25** Repeat Problem 13.24 with $U(x) = 0.0003x - 8.48 e^{-x/2775}$. A utility function of this form is called *linear-plus-exponential*, because it contains both linear ($0.0003x$) and exponential terms. It has a number of interesting and useful properties, including the fact that it only switches once among any pair of lotteries (such as those in Problem 13.24) as wealth increases (see Bell 1995a, b).

- 13.26** Show that the linear-plus-exponential utility function in Problem 13.25 has decreasing risk aversion. (*Hint:* You can use this utility function to evaluate a series of lotteries like those described in the text and analyzed in Tables 13.3 and 13.4. Show that the risk premium for a gamble decreases as wealth increases.)

- 13.27** Buying and selling prices for risky investments obviously are related to certainty equivalents. This problem, however, shows that the prices depend on exactly what is owned in the first place!

Suppose that Peter Brown's utility for total wealth (A) can be represented by the utility function $U(A) = \ln(A)$. He currently has \$1000 in cash. A business deal of interest to him yields a reward of \$100 with probability 0.5 and \$0 with probability 0.5.

- a** If he owns this business deal in addition to the \$1000, what is the smallest amount for which he would sell the deal?
b Suppose he does not own the deal. What equation must be solved to find the largest amount he would be willing to pay for the deal?

- c For part b, it turns out that the most he would pay is \$48.75, which is not exactly the same as the answer in part a. Can you explain why the amounts are different?
- d *{Extra credit for algebra hotshots.}* Solve your equation in part b to verify the answer (\$48.75) given in part c.

Source: This problem was suggested by R. L. Winkler.

13.28 We discussed decreasing and constant risk aversion. Are there other possibilities? Think about this as you work through this problem.

Suppose that a person's utility function for total wealth is

$$U(A) = 200A - A^2 \quad \text{for } 0 \leq A \leq 100$$

where A represents total wealth in thousands of dollars.

- a Graph this preference function. How would you classify this person with regard to her attitude toward risk?
- b If the person's total assets are currently \$10K, should she take a bet in which she will win \$10K with probability 0.6 and lose \$10K with probability 0.4?
- c If the person's total assets are currently \$90K, should she take the bet given in part b?
- d Compare your answers to parts b and c. Does the person's betting behavior seem reasonable to you? How could you intuitively explain such behavior?

Source: R. L. Winkler (1972).

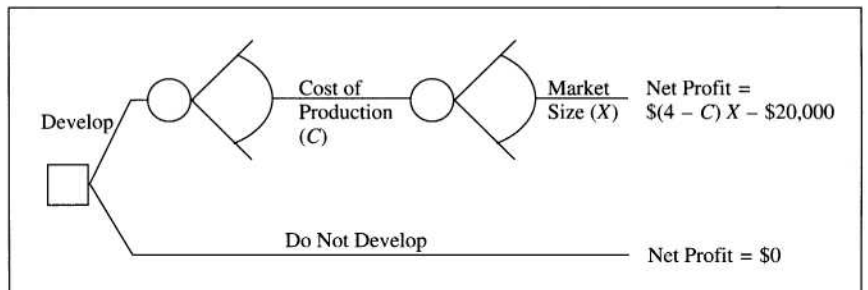
13.29 Suppose that a decision maker has the following utility function:

$$U(x) = -0.000156x^2 + 0.028125x - 0.265625$$

Use this utility function to calculate risk premiums for the gambles shown in Tables 13.3 and 13.4; create a similar table but based on this quadratic utility function. How would you classify the risk attitude of a decision maker with this utility function? Does such a risk attitude seem reasonable to you?

13.30 The CEO of a chemicals firm must decide whether to develop a new process that has been suggested by the research division. His decision tree is shown in Figure 13.17. There are two sources of uncertainty. The production cost is viewed as a continuous random variable, uniformly distributed between \$1.75 and \$2.25, and the size of the market

Figure 13.17
Decision tree for Problem 13.30.



(units sold) for the product is normally distributed with mean 10,300 units and standard deviation 2200 units.

The firm's CEO is slightly risk-averse. His utility function is given by

$$U(Z) = 1 - e^{-Z/20,000} \quad \text{where } Z \text{ is the net profit}$$

Should the CEO develop the new process? Answer this question by running a computer simulation, using 200 trials. Should the decision maker be concerned about the fact that if he develops the new process, the utility could be less than or greater than the utility for \$0? On what basis should he make his decision?

- 13.31** The year is 2020, and you are in the supercomputer business. Your firm currently produces a machine that is relatively expensive (list price \$6 million) and relatively slow (for supercomputers in the twenty-first century). Speed of supercomputers is measured in gigaflops per second (gps), where one "flop" is one calculation. Thus, one 1 gps = 1 billion calculations per second. Your current machine is capable of 150 gps. If you could do it, you would prefer to develop a supercomputer that costs less (to beat the competition) and is faster.

You have a research-and-development (R&D) decision to make based on two alternatives. You can choose one or the other of the following projects, or neither, but budget constraints prevent you from engaging in both projects.

A The super-supercomputer. This project involves the development of a machine that is extremely fast (800 gps) and relatively inexpensive (\$5 million). But this is a fairly risky project. The engineers who have been involved in the early stages estimate that there is only a 50% chance that this project would succeed. If it fails, you will be stuck with your current machine.

B The better supercomputer. This project would pursue the development of an \$8 million machine capable of 500 gps. This project also is somewhat risky. The engineers believe that there is only a 40% chance that this project will achieve its goal. They quickly point out, however, that even if the \$8 million, 500 gps machine does not materialize, the technology involved is such that they would at least be able to produce a \$5 million machine capable of 350 gps.

The decision tree is shown in Figure 13.18. To decide between the two alternatives, you have made the following assessments:

- I** The best possible outcome is the \$5 million, 800 gps machine, and the worst outcome is the status quo \$6 million, 150 gps machine.
 - II** If you had the choice, you would be indifferent between Alternatives X and Y shown in Figure 13.19a.
 - III** If you had the choice, you would be indifferent between Alternatives X' and Y' in Figure 13.19b.
 - a** Using assessments I, II, and III, decide between Projects A and B. Justify your decision.
 - b** Explain why Project A appears to be riskier than Project B. Given that A is riskier than B, would you change your answer to part a? Why or why not?
- 13.32** Show that the value of Y that yields indifference between the two alternatives in Figure 13.12 is within about 4% of the risk tolerance R .

Figure 13.18
 Choosing between two risky supercomputer development projects.





Figure 13.19
 Assessments to assist your choice from among the supercomputer R&D projects in Problem 13.31.



CASE STUDIES

INTERPLANTS, INC.

Don Newcomb was perplexed. It had been five years since he had founded Interplants, Inc., a research-and-development firm that developed genetically engineered plants for interplanetary space flight. During that five years, he and his scientists had made dramatic advances in developing special plants that could be used for food and air-purification systems in space stations and transports. In fact, he mused, their scientific success had been far greater than he had ever expected.

Five years ago, after the world superpowers had agreed to share space-travel technology, the outlook had been quite rosy. Everyone had been optimistic. Indeed, he was one of many investors who had jumped at the chance to be involved in the development of such technology. But now, after five tumultuous years, the prospects were less exciting.

First, there had been the disappointing realization that none of the superpowers had made substantial success on an ion engine to power space vehicles. Such an engine

was absolutely crucial to the success of interplanetary space flight, because — theoretically, at least — it would make travel 10 times as fast as conventionally powered ships. When the importance of such an engine became obvious, the superpowers had generously funded a huge multinational research project. The project had made substantial progress, but many hurdles remained. Don's risk assessors estimated that there was still a 15% chance that the ion engine would prove an infeasible power source. If this were to happen, of course, Don and the many other investors in space-travel technology would lose out.

Then there was the problem with the settlement policy. The superpowers could not agree on a joint policy for the settlement of interplanetary space, including the deployment of space stations as well as settlements on planets and their satellites. The United American Alliance urged discretion and long-range planning in this matter, suggesting that a multinational commission be established to approve individual settlement projects. Pacificasia and the Allied Slavic Economic Community were demanding that space be divided now. By immediately establishing property rights, they claimed, the superpowers would be able to develop the optimum space economy in which the superpowers could establish their own economic policies within their "colonies" as well as determine trade policies with the other superpowers. Europa favored the idea of a commission, but also was eager to explore other available economic possibilities.

The discussion among the superpowers had been going on since long before the founding of Interplants. Five years ago, progress was being made, and it appeared that an agreement was imminent. But 18 months ago the process stalled. The participants in the negotiations had established positions from which they would not budge. Don had followed the discussions closely and had even provided expert advice to the negotiators regarding the potential for interplanetary agricultural developments. He guessed that there was only a 68% chance that the superpowers would eventually arrive at an agreement. Naturally, until an agreement was reached there would be little demand for space-traveling plants.

Aside from these external matters, Don still faced the difficult issue of developing a full-scale production process for his plants. He and his engineers had some ideas about the costs of the process, but at this point, all they could do was come up with a probability distribution for the costs. In thinking about the distribution, Don had decided to approximate it with a three-point discrete distribution. Thus, he characterized the three branches as "inexpensive," "moderate," and "costly," with probabilities of 0.185, 0.63, and 0.185, respectively. Of course, his eventual profit (or loss) depended on the costs of the final process.

Don also had thought long and hard about the profits that he could anticipate under each of the various scenarios. Essentially, he thought about the uncertainty in two stages. First was the determination of costs, and second was the outcome of the external factors (the ion-engine research and the negotiations regarding settlement policy). If costs turned out to be "inexpensive," then, in the event that the superpowers agreed and the ion engine was successful, he could expect a profit of 125 billion credits. He would lose 15 billion credits if either the engine or the negotiations failed. Likewise, if costs were "moderate," he could anticipate either a profit of 100

billion credits if both of the external factors resulted in a positive outcome, or a loss of 18 billion if either of the external factors were negative. Finally, the corresponding figures in the case of a "costly" production process were profits of 75 billion credits or a loss of 23 billion.

"This is so confusing," complained Don to Paul Fiester, his chief engineer. "I really never expected to be in this position. Five years ago none of these risks were apparent to me, and I guess I just don't tolerate risk well."

After a pause, Paul quietly suggested "Well, maybe you should sell the business."

Don considered that. "Well, that's a possibility. I have no idea how many crazy people out there would want it."

"Some of the other engineers and I might be crazy enough," Paul replied. "Depending on the price, of course. At least we'd be going in with our eyes open. We know what the business is about and what the risks are."

Don gave the matter a lot of thought that night. "What should I sell the company for? I hate to give up the possibility of earning up to 125 billion credits. But I don't like the possibility of losing 23 billion either — no one would!" As he lay awake, he finally decided that he would let the business go — with all its risks — for 20 billion credits. If he could get that much for it, he'd sell. If not, he'd just as soon stick with it, in spite of his frustrations with the risks.

Questions

- 1 Draw a decision tree for Don Newcomb's problem.
- 2 What is the significance of his statement that he would sell the business for 20 billion credits?
- 3 Suppose that Don's risk attitude can be modeled with an exponential utility function. If his certainty equivalent were 15 billion credits, find his risk tolerance. What would his risk tolerance be if his CE were 20 billion?

TEXACO PENNZOIL ONE MORE TIME

In Problem 12.11 (page 454), we made EVPI calculations for Liedtke's decision in the Texaco-Pennzoil case. Calculating EVPI is somewhat different when the decision maker is risk-averse; that is, we cannot simply fold a tree back using expected values and then compare the expected values with and without information. And we cannot fold back the tree in terms of expected utility and compare the expected utilities. In fact, in general we must find a value C such that when we subtract that value from each endpoint of the decision tree on the "Acquire Information" branch, the expected utility of this branch is equal to the expected utility of the "No Information" branch. Figure 13.20 illustrates the principle by including a branch for acquiring information about Texaco's response.

tolerance of \$1 billion. An acquaintance of Liedtke's knows Texaco CEO James Kinnear quite well and has offered to find out how Kinnear would react to a \$5 billion counteroffer. Thus, the third alternative available to Liedtke is to acquire information about Texaco.

Question

- 1 Find the expected value of perfect information regarding the Texaco reaction. Compare your answer to the answer from Problem 12.11a.

STRENLAR, PART III

Consider once again Fred Wallace's decision in the Strenlar case study at the end of Chapter 4 (pages 148-150). What if Fred is risk-averse? Assume that Fred's attitude toward risk in this case can be adequately modeled using an exponential utility function in which the utility is calculated for net present value. Thus,

$$U(\text{NPV}) = 1 - e^{-\text{NPV}/R}$$

Question

- 1 Check the sensitivity of Fred's decision to his risk tolerance, R . What is the critical R value for which his optimal decision changes? What advice can you give to Fred?

REFERENCES

For the most part, the material presented in this chapter has been my own version of familiar material. Good basic treatments of expected utility and risk aversion are available in most decision-analysis textbooks, including Bunn (1984), Holloway (1979), Raiffa (1968), Vatter et al. (1978), von Winterfeldt and Edwards (1986), and Winkler (1972). Keeney and Raiffa (1976) offer a somewhat more advanced treatment that focuses on multiattribute utility models, although it does contain an excellent exposition of the basic material at a somewhat higher mathematical level. The text by von Winterfeldt and Edwards (1986) points out that we tend to think about different kinds of money (pocket money, monthly income, investment capital) in different ways; thus, how utility-assessment questions are framed in terms of these different funds can have a strong effect on the utility function.

The material on the exponential utility function and risk tolerance is based primarily on Holloway (1979) and McNamee and Celona (1987). Both books contain excellent discussion and problems on this material.

For students who wish more detail on decreasing and constant risk aversion, look in the financial economics literature. In financial models, utility functions for market participants

are important for modeling the economic system. An excellent starting point is Copeland and Weston (1979). A classic article that develops the idea of a risk-aversion coefficient (the reciprocal of risk tolerance) is Pratt (1964). Bell (1995a, b) discusses the "linear-plus-exponential" utility function $aw - b e^{-cw}$, where w is total wealth. This utility function has many desirable properties; see Problems 13.25 and 13.26.

Bell, D. (1995a) "Risk, Return, and Utility." *Management Science*, 41, 23-30.

Bell D. (1995b) "A Contextual Uncertainty Condition for Behavior under Risk." *Management Science*, 41, in press.

Bunn, D. (1984) *Applied Decision Analysis*. New York: McGraw-Hill.

Copeland, T. E., and J. F. Weston (1979) *Financial Theory and Corporate Policy*. Reading, MA: Addison-Wesley.

Holloway, C. A. (1979) *Decision Making Under Uncertainty: Models and Choices*. Englewood Cliffs, NJ: Prentice-Hall.

Howard, R. A. (1988) "Decision Analysis: Practice and Promise." *Management Science*, 34, 679-695.

Keeney, R., and H. Raiffa (1976) *Decisions with Multiple Objectives*. New York: Wiley.

McNamee, P., and J. Celona (1987) *Decision Analysis for the Professional with Supertree*. Redwood City, CA: Scientific Press.

Pratt, J. (1964) "Risk Aversion in the Small and in the Large." *Econometrica*, 32, 122-136.

Raiffa, H. (1968) *Decision Analysis*. Reading, MA: Addison-Wesley.

Vatter, P. A., S. P. Bradley, S. C. Frey, Jr., and B. B. Jackson (1978) *Quantitative Methods in Management*. Home wood, IL: Irwin.

Von Winterfeldt, D., and W. Edwards (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press.

Winkler, R. L. (1972) *Introduction to Bayesian Inference and Decision*. New York: Holt.

EPILOGUE

Harriman's broker, Jacob H. Schiff, was at his synagogue when Harriman's order for 40,000 shares was placed; the shares were never purchased. By the following Monday, Hill had cabled Morgan in France, and they had decided to buy as many Northern Pacific shares as they could. The share price went from \$114 on Monday to \$1000 on Thursday. In the aftermath, Hill and Morgan agreed that Harriman should be represented on the board of directors. Harriman, however, had little if any influence; James Hill continued to run the Great Northern, Northern Pacific, and Burlington railroads as he saw fit. [Source: S. H. Holbrook (1958) "The Legend of Jim Hill." *American Heritage*, IX(4), 10-13, 98-101.]

CHAPTER 14. Utility Axioms, Paradoxes, and Implications

In this chapter we will look at several issues. First, we will consider some of the foundations of utility theory. From the basis of a few behavioral axioms, it is possible to establish logically that people who behave according to the axioms should make choices consistent with the maximization of expected utility. But since the early 1950s, cognitive psychologists have noted that people do not always behave according to expected utility theory, and a large literature now covers these behavioral paradoxes. We review part of that literature here. Because decision analysis depends on foundational axioms, it is worthwhile to consider some implications of these behavioral paradoxes, particularly with regard to the assessment of utility functions.

The following example previews some of the issues we will consider. This one is a participatory example. You should think hard about the choices you are asked to make before reading on.

PREPARING FOR AN INFLUENZA OUTBREAK

The United States is preparing for an outbreak of an unusual Asian strain of influenza. Experts expect 600 people to die from the disease. Two programs are available that could be used to combat the disease, but because of limited resources only one can be implemented.

Program A (Tried and True) 400 people will be saved.

Program B (Experimental) There is an 80% chance that 600 people will be saved and a 20% chance that no one will be saved.

Which of these two programs do you prefer? Now consider the following two programs:

Program C 200 people will die.

Program D There is a 20% chance that 600 people will die and an 80% chance that no one will die.

Would you prefer C or D?

Source: Reprinted with permission from Tversky, A., and D. Kahneman, "The Framing of Decisions and the Psychology of Choice," *Science*, 211, 453-458. Copyright 1981 by the American Association for the Advancement of Science.

Axioms for Expected Utility

Our first step in this chapter is to look at the behavioral assumptions that form the basis of expected utility. These assumptions, or *axioms*, relate to the consistency with which an individual expresses preferences from among a series of risky prospects. Instead of axioms, we might call these *rules for clear thinking*. In the following discussion, the axioms are presented at a fairly abstract level. Simple examples are given to clarify their meaning. As we put them to work in the development of the main argument, the importance and intuition behind the axioms should become clearer.


- 1 *Ordering and transitivity.* A decision maker can order (establish preference or in difference) any two alternatives, and the ordering is transitive. For example, given any alternatives A_1 , A_2 , and A_3 , either A_1 is preferred to A_2 (which is sometimes written as " $A_1 > A_2$ "), A_2 is preferred to A_1 or the decision maker is indifferent between A_1 and A_2 ($A_1 \sim A_2$). Transitivity means that if A_1 is preferred to A_2 and A_2 is preferred to A_3 , then A_1 is preferred to A_3 . For example, this axiom says that an individual could express his or her preferences regarding, say, cities in which to reside. If that person preferred Amsterdam to London and London to Paris, then he or she would prefer Amsterdam to Paris.
- 2 *Reduction of compound uncertain events.* A decision maker is indifferent between a compound uncertain event (a complicated mixture of gambles or lotteries) and a simple uncertain event as determined by reduction using standard probability manipulations. This comes into play when we reduce compound events into reference gambles. The assumption says that we can perform the reduction without affecting the decision maker's preferences. We made use of this axiom in Chapter 4 in our discussion of risk profiles. The progression from Figure 4.20 through Figure 4.22 is a matter of reducing to simpler terms the compound uncertain event that is associated with the counteroffer.
- 3 *Continuity.* A decision maker is indifferent between a consequence A (for example, win 100) and some uncertain event involving only two basic consequences A_1 and A_2 , where $A_1 > A > A_2$. This simply says that we can construct a reference gamble

with some probability p , $0 < p < 1$, for which the decision maker will be indifferent between the reference gamble and A. For example, suppose you find yourself as the plaintiff in a court case. You believe that the court will award you either \$5000 or nothing. Now imagine that the defendant offers to pay you \$1500 to drop the charges. According to the continuity axiom, there must be some probability p of winning \$5000 (and the corresponding $1 - p$ probability of winning nothing) for which you would be indifferent between taking or rejecting the settlement offer. Of course, if your subjective probability of winning happens to be lower than p , then you would accept the proposal.

- 4 *Substitutability.* A decision maker is indifferent between any original uncertain event that includes outcome A and one formed by substituting for A an uncertain event that is judged to be its equivalent. Figure 14.1 shows how this works. This axiom allows the substitution of uncertain reference gambles into a decision for their certainty equivalents and is just the reverse of the reduction axiom already stated. For example, suppose you are interested in playing the lottery, and you are just barely willing to pay 50 cents for a ticket. If I owe you 50 cents, then you should be just as willing to accept a lottery ticket as the 50 cents in cash.
- 5 *Monotonicity.* Given two reference gambles with the same possible outcomes, a decision maker prefers the one with the higher probability of winning the preferred outcome. This one is easy to see. Imagine that two different car dealerships each can order the new car that you want. Both dealers offer the same price, delivery, warranty, and financing, but one is more likely to provide good service than the other. To which one would you go? The one that has the better chance of providing good service, of course.
- 6 *Invariance.* All that is needed to determine a decision maker's preferences among uncertain events are the payoffs (or consequences) and the associated probabilities.
- 7 *Finiteness.* No consequences are considered infinitely bad or infinitely good.

Most of us agree that these assumptions are reasonable under almost all circumstances. It is worth noting, however, that many decision theorists find some of the axioms controversial! The reasons for the controversy range from introspection regarding particular decision situations to formal psychological experiments in which human subjects make choices that clearly violate one or more of the axioms. We will discuss some of these experiments in the next section.

Figure 14.1
Two decision trees. If A is equivalent to a lottery with a p chance at C and $1 - p$ chance at D, then Decision Tree I is equivalent to Decision Tree II.



For example, the substitutability axiom is a particular point of debate. For some decision makers, the fact of having to deal with two uncertain events in Decision Tree II of Figure 14.1 can be worse than facing the single one in Decision Tree I. Moreover, individuals might make this judgment and at the same time agree that in a single-stage lottery, A is indeed equivalent to the risky prospect with p chance at C and a $1 - p$ chance at D.

As another example, we can pick on the apparently innocuous transitivity axiom. In Figure 14.2, you have two lotteries from which to choose. Each of the six outcomes has probability $1/6$. One way to look at the situation is that the prize in Game B is better than Game A's in five of the six outcomes, and thus it may be reasonable to prefer B, even though the structure of the lotteries is essentially the same. Now consider the games in Figure 14.3. If B was preferred to A in Figure 14.2, then by the same argument C would be preferred to B, D to C, E to D, F to E, and, finally, A would be preferred to F. Thus, these preferences do not obey the transitivity axiom, because transitivity would never permit A to be preferred to something else that is in turn preferred to A.

The controversy about individual axioms notwithstanding, if you accept axioms 1 through 7, then logically you also must accept the following proposition.

Figure 14.2

A pair of lotteries. Outcomes 1 through 6 each occur with probability $1/6$. Would you prefer to play Game A or Game B?

		Outcome					
		1	2	3	4	5	6
Prize	Game A	100	200	300	400	500	600
	Game B	200	300	400	500	600	100

Figure 14.3

More games to consider. Outcomes 1 through 6 each still have probability $1/6$.

		Outcome					
		1	2	3	4	5	6
Prize	Game B	200	300	400	500	600	100
	Game C	300	400	500	600	100	200
	Game D	400	500	600	100	200	300
	Game E	500	600	100	200	300	400
	Game F	600	100	200	300	400	500
	Game A	100	200	300	400	500	600

Proposition: Given any two uncertain outcomes B_1 and B_2 , if assumptions 1 through 7 hold, there are numbers U_1, U_2, \dots, U_n representing preferences (or utilities) associated with the consequences such that the overall preference between the uncertain events is reflected by the expected values of the U 's for each consequence. In other words, if you accept the axioms, (1) it is possible to find a utility function for you to evaluate the consequences, and (2) you should be making your decisions in a way that is consistent with maximizing expected utility.

In the following pages we will demonstrate how the axioms permit the transformation of uncertain alternatives into reference gambles (gambles between the best and worst alternatives) with different probabilities. It is on the basis of this kind of transformation that the proposition can be proved.

Suppose you face the simple decision problem shown in Figure 14.4. For convenience, assume the payoffs are in dollars. The continuity axiom says that we can find reference gambles that are equivalent to the outcomes at the ends of the branches. Suppose (hypothetically) that you are indifferent between \$15 and the following reference gamble:

- Win \$40 with probability 0.36
- Win \$10 with probability 0.64

Likewise, suppose you are indifferent between \$20 and the next reference gamble:

- Win \$40 with probability 0.60
- Win \$10 with probability 0.40

The substitutability axiom says that we can replace the original outcomes with their corresponding reference gambles, as in Figure 14.5. (We have replaced the outcomes in Lottery B with "trivial" lotteries. The reason for doing so will become apparent.) The substitutability axiom says that you are indifferent between A and A' and also between B and B'. Thus, the problem has not changed.

Now use the reduction-of-compound-events axiom to reduce the decision tree (Figure 14.6). In performing this step, the overall probability of winning 40 in A'' is $0.5(0.60) + 0.5(0.36)$, or 0.48; it is similarly calculated for winning 10. For the lower half of the tree, we just retrieve B again. The monotonicity axiom means we prefer A'' to B'', and so by transitivity (which says that $A'' \sim A' \sim A$ and $B'' \sim B' \sim B$) we must prefer A to B in the original decision.

To finish the demonstration, we must show that it is possible to come up with numbers that represent utilities so that a higher expected utility implies a preferred alternative, and vice versa. In this case, we need utilities that result in a higher expected utility for Alternative A. Use the probabilities assessed above in the reference gambles as those utilities. (Now you can see the reason for the extension of B in Figure 14.5; we need the probabilities.) We can redraw the original decision tree, as in Figure 14.7, with the utilities in place of the original monetary payoffs.

Figure 14.4
A simple decision problem under uncertainty.




Figure 14.5
Decision tree after substituting reference gambles for outcomes.

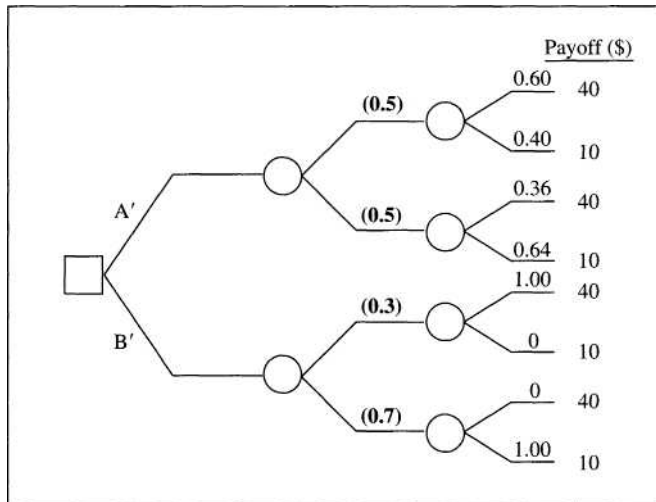



Figure 14.6
Reducing the decision tree.



Calculating expected utilities shows that A has the higher expected utility and hence should be preferred:

$$EU(A) = 0.5(0.60) + 0.5(0.36) = 0.48$$

$$EU(B) = 0.3(1.00) + 0.7(0) = 0.30$$

This exercise may seem both arcane and academic. But put simply, if you think it is reasonable to behave according to the axioms at the beginning of the section, then (1) it is possible to find your utility values, and (2) you should make decisions that would be consistent with the maximization of expected utility.

In our initial Chapter 13 discussions of utility assessment, the claim was made that a utility function could be scaled (multiplied by a positive constant and added to a constant) without changing anything. Now you should see clearly that the whole purpose of a utility function is to rank-order risky prospects. Take Alternatives A and B. We have concluded that A is preferred to B because $EU(A) = 0.48$, which is greater than $EU(B) = 0.30$. Suppose that we take the utility numbers, as shown in Figure 14.7, and scale them. That is, let

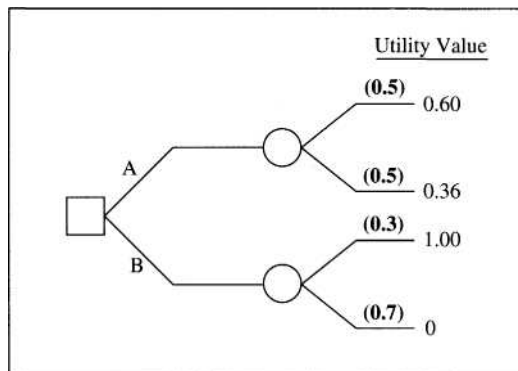
$$U'(x) = a + bU(x)$$

where $b > 0$ and $U(x)$ is our original utility function (from Figure 14.7). We can calculate the expected utilities of A and B on the basis of the following new utility function:

$$EU'(A) = 0.5[a + bU(20)] + 0.5[a + bU(15)] = a + b[0.5U(20) + 0.5U(15)] = a + bEU(A) = a + b(0.48)$$

$$EU'(B) = 0.3[a + bU(40)] + 0.7[a + bU(10)] = a + b[0.3U(40) + 0.7U(10)] = a + bEU(B) = a + b(0.30)$$

Figure 14.7
Original decision tree with utility values replacing monetary values.



It should be clear that $EU'(A)$ will be greater than $EU'(B)$ as long as $b > 0$. As indicated, the implication is that we can scale our utility functions linearly with the constants a and b without changing the rankings of the risky alternatives in terms of expected utility. Specifically, this means that no matter what the scale of a utility function, we always can rescale it so that the largest value is 1 and the smallest is 0. (For that matter, you can rescale it so that the largest and smallest values are whatever you want!)

If you look carefully at our transformations of the original decision problem above, you will see that we never explicitly invoked either the invariance or finiteness axioms. Why are these axioms important? The invariance axiom says that we need nothing but the payoffs and the probabilities; nothing else matters. (In the context of multiattribute utility as discussed in the next two chapters, we would need consequences on all relevant dimensions.)

The finiteness axiom assures us that expected utility will never be infinite, and so we always will be able to make meaningful comparisons. To see the problem with unbounded utility, suppose that you have been approached by an evangelist who has told you that unless you accept his religion, you will burn in Hell for eternity. If you attach an infinitely negative utility to an eternity in Hell (and it is difficult to imagine a worse fate), then no matter how small your subjective probability that the evangelist is right, as long as it is even slightly positive you must be compelled to convert. This is simply because any small positive probability multiplied by the infinitely negative utility will result in an infinitely negative expected utility. Similar problems are encountered if an outcome is accorded infinite positive utility; you would do anything at all if doing so gave you even the slightest chance at achieving some wonderful outcome. Thus, if an outcome in your decision problem has unbounded utility, then the expected utility approach does not help much when it comes to making the decision.

Paradoxes

Even though the axioms of expected utility theory appear to be compelling when we discuss them, people do not necessarily make choices in accordance with them. Research into these behavioral paradoxes began almost as early as the original research into utility theory itself, and now a large literature exists for many aspects of human behavior under uncertainty. Much of this literature is reviewed in von Winterfeldt and Edwards (1986) and Hogarth (1987). We will cover a few high points to indicate the nature of the results.

Framing effects are among the most pervasive paradoxes in choice behavior. Tversky and Kahneman (1981) show how an individual's risk attitude can change depending on the way the decision problem is posed — that is, on the "frame" in which a problem is presented. The difficulty is that the same decision problem usually can be expressed in different frames. A good example is the influenza-outbreak


problem at the beginning of the chapter. You may have noticed that Program A is the same as C and that B is the same as D. It all depends on whether you think in terms of deaths or lives saved. Many people prefer A on one hand, but D on the other.

To a great extent, the reason for the inconsistent choices appears to be that different points of reference are used to frame the problem in two different ways. That is, in Programs A and B the reference point is that 600 people are expected to die, but some may be saved. Thus, we think about gains in terms of numbers of lives saved. On the other hand, in Programs C and D, the reference point is that no people would be expected to die without the disease. In this case, we tend to think about lives lost. One of the important general principles that psychologists Kahneman and Tversky and others have discovered is that people tend to be risk-averse in dealing with gains but risk-seeking in deciding about losses. A typical assessed utility function for changes in wealth is shown in Figure 14.8. These results have been obtained in many different behavioral experiments (for example, see Swalm 1966).

More fundamental than the risk-averse-risk-seeking dichotomy is that the reference point or status quo can be quite flexible in some situations and inflexible in others. For example, the influenza-outbreak example can be viewed with relative ease from either frame. For many people, the financial status quo changes as soon as they file their income-tax return in anticipation of a refund; they "spend" their refund, usually in the form of credit, long before the check arrives in the mail. In other cases, individuals may maintain a particular reference point far longer than they should. A case in point is that gamblers often try to make up their losses; they throw good money after bad. Here "gamblers" can refer to individuals in casinos as well as to stock market investors or even managers who maintain a commitment to a project that has obviously gone sour. Typically, such a gambler will argue that backing out of a failed project amounts to a waste of the resources already spent.

Is a specific axiom being violated when a decision maker's choices exhibit a framing effect? The answer to this question is not exactly clear. Although many possibilities exist, the invariance axiom may be the weak link in this case. It may be that

Figure 14.8
*Typical assessed utility
function for changes
in wealth.*



payoffs (or utilities) and probabilities are not sufficient to determine a decision maker's preferences. Some understanding of the decision maker's frame of reference also may be required.

For another example, consider the problem:

Allais Paradox

You have two decisions to make.

Decision 1: **A** Win \$1 million with probability 1
 B Win \$5 million with probability 0.10
 Win \$1 million with probability 0.89
 Win \$0 with probability 0.01

Before proceeding, choose A or B. Would you give up a sure \$1 million for a small chance at \$5 million and possibly nothing?

Decision 2: **C** Win \$ 1 million with probability 0.11
 Win \$0 with probability 0.89 **D** Win \$5 million with probability 0.10
 Win \$0 million with probability 0.90

Now choose C or D in Decision 2.

This is the well-known Allais Paradox (Allais 1953; Allais and Hagen 1979). The decisions are shown in decision-tree form in Figure 14.9. Experimentally, as many as 82% of subjects prefer A over B and 83% prefer D over C. But we can easily show that choosing A on the one hand and D on the other is contrary to expected utility maximization. Let $U(\$0) = 0$ and $U(\$5,000,000) = 1$; they are the best and worst outcomes. Then,

$$EU(A) = U(\$1 \text{ million})$$

$$EU(B) = 0.10 + 0.89U(\$1 \text{ million})$$

Thus, A is preferred to B if and only if

$$U(\$1 \text{ million}) > 0.10 + 0.89 U(\$1 \text{ million})$$

or

$$U(\$1 \text{ million}) > 0.91$$

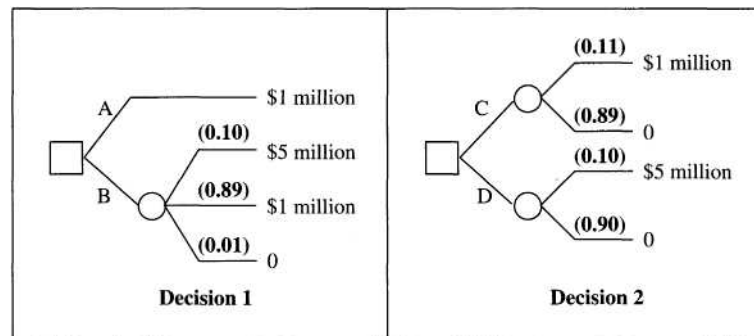
Now for Decision 2,

$$EU(C) = 0.11 U(\$1 \text{ million})$$

$$EU(D) = 0.10 \text{ so D is preferred}$$

to C if and only if $U(\$1 \text{ million}) < 0.91$.

Figure 14.9
Choices in the Allais
Paradox.



$U(\$1 \text{ million})$ cannot be both greater than and less than 0.91 at the same time, so choosing A and D is not consistent with expected utility. Consistent choices are A and C or B and D. Kahneman and Tversky (1981) have attributed this common inconsistency to the *certainty effect*, whereby individuals tend to place too much weight on a certain outcome relative to uncertain outcomes. In the Allais Paradox, the certainty effect would tend to make individuals overvalue A in Decision 1, possibly leading to an inconsistent choice. When confronted with their inconsistency, some individuals revise their choices. Would you revise yours in light of this discussion?

Another way to look at the Allais Paradox is to structure the decision problem using lottery tickets in a hat. Imagine that 100 tickets are numbered sequentially from 1 to 100 and placed in a hat. One ticket will be drawn at random, and you will receive a prize depending on the option you choose and the number on the ticket. Prizes for A and B are described in Table 14.1. For A, you would win \$1 million regardless of the ticket chosen, but for B you would win nothing if Ticket 1 is chosen, \$5 million for Tickets 2 through 11, and \$1 million for Tickets 12 through 100. Note that you win the same (\$1 million) in the two lotteries for Tickets 12 through 100. Thus, your choice should depend only on your preferences regarding the outcomes for Tickets 1 through 12.

The same kind of thing can be done for Options C and D, which also are shown in Table 14.1. If you choose C, you win \$1 million for Tickets 1-11 and nothing for Tickets 12-100. In D you would win nothing for Ticket 1, \$5 million for Tickets 2-11, and nothing for Tickets 12-100. Again, you win exactly the same thing (nothing) in both C and D for Tickets 12-100, and so your preferences between C and D should depend only on your preferences regarding Tickets 1-11. As you can see, the prizes associated with Tickets 1-11 are the same for Options A and C on the one hand and B and D on the other. Thus, if you prefer A to B you also should prefer C to D, and vice versa.

It is intuitively reasonable that your preferences should not depend on Tickets 12-100, because the outcome is the same regardless of the decision made. This is an example of the *sure-thing principle*, which says that our preferences over lotteries or risky prospects should depend only on the parts of the lotteries that can lead to different outcomes. The idea of the sure thing is that, in the choice between A and B,

Table 14.1
Prizes in the Allais Paradox. Tickets 1-100 are placed in a hat, and one ticket is drawn randomly. The dollar amounts shown in the table are the prizes for the four options.

Option	Tickets		
	1	2-11	12-100
A	\$1 million	\$1 million	\$1 million
B	0	\$5 million	\$1 million
C	\$1 million	\$1 million	0
D	0	\$5 million	0

winning \$1 million is a sure thing if one of the tickets from 12 to 100 is drawn, regardless of your choice. If, as in the Allais choices, there are possible outcomes that have the same value to you regardless of the option you choose, and these outcomes occur with the same probability regardless of the option chosen, then you can ignore this part of the lottery. The sure-thing principle can be derived logically from our axioms. For our purposes, it is best to think of it as a behavioral principle that is consistent with our axioms; if you agree to behave in accordance with the axioms, then you also should obey the sure-thing principle. If your choices in the Allais Paradox were inconsistent, then your preferences violated the sure-thing principle.

Implications

As we learned above, people do not always behave according to the behavioral axioms. This fact has distinct implications for the practice of decision analysis. First, there are implications regarding how utility assessments should be made. Second, and perhaps more intriguing, are the implications for managers and policy makers whose jobs depend on how people actually do make decisions. It may be important for such managers to consider some of the above-described behavioral phenomena. In this section, we will look at the issues involved in each of these areas.

Implications for Utility Assessment

We rely on assessments of certainty equivalents and other comparisons to find a utility function. Given the discussion above, it is clear that, in the assessment process, we ask decision makers to perform tasks that we know they do not always perform consistently according to the axioms! Thus, it is no surprise that the behavioral paradoxes discussed above may have some impact on the way that utility functions should be assessed. Our discussion here will be brief, not because the literature is particularly large, but because many research questions remain to be answered.

There are several approaches to the assessment of utilities. In Chapter 13 we introduced the certainty-equivalent approach (find a CE for a specified gamble) and the probability-equivalent approach (find a probability that makes a reference lottery

equivalent to a specific certain amount). If we think of the general case (Figure 14.10) as being indifferent between CE for sure and the reference lottery:


Win G with probability p Win L
with probability $1 - p$

then it is clear that we must preset three out of the four variables p , CE, G, and L. The selection of the fourth, which makes the decision maker indifferent between the two prospects, allows the specification of a utility value. Most practitioners use either the CE or PE approaches. Assessing a certainty equivalent involves fixing G, L, and p and assessing CE; the probability-equivalent method involves fixing CE, G, and L and assessing p .

The relative merits of the possible assessment techniques have been discussed to some degree. Hershey, Kunreuther, and Shoemaker (1982) report that the use of the CE approach tends to result in more risk-averse responses than does the PE approach when the consequences are gains. On the other hand, when the consequences are losses, the CE approach results in more risk-seeking behavior. When using the PE approach, many people appear to exhibit certain forms of probability distortion. Although the evidence is far from conclusive, it indicates that people deal best with 50-50 chances. This empirical result appears to be related to the certainty effect discussed above.

Clearly, these results have an impact on the assessment of utility functions; the nature of the decision maker's responses and hence the deduced risk attitude can depend on the way that questions have been posed in the assessment procedure. McCord and De Neufville (1986) have suggested, in light of the distortion from the certainty effect, that utilities should not be assessed using the CE approach, which requires a decision maker to compare a lottery with a certain quantity. The CE approach, they argue, contains a built-in bias. They suggest that it would be more appropriate to assess utilities by comparing lotteries. For example, suppose that A is the best outcome and C is the worst, and we would like to assess $U(B)$, which is somewhere between A and C. McCord and De Neufville's technique would have the decision maker assess the probability p that produces indifference between the two

Figure 14.10
A general framework for assessing utilities.



If three of the four variables (G, L, CE, and p) are fixed, assessment of the fourth such that the decision maker is indifferent between the two decision branches permits establishment of utility values.

lotteries in the decision tree in Figure 14.11. Because p makes the decision maker indifferent, we can set up the equation

$$0.5 U(B) + 0.5 U(C) = p U(A) + (1 - p) U(C)$$

Substituting $U(A) = 1$ and $U(C) = 0$ and rearranging, this becomes

$$U(B) = 2p$$

Thus, the utility is relatively easy to find once the assessment has been made.

The McCord-De Neufville approach does indeed lead to utility assessments that are less risk-averse than those made with certainty equivalents. Figure 14.12 compares the two methods in terms of hypothetical utility functions that might be assessed. It may be the case, however, that decision makers have a harder time thinking about comparable lotteries than about certainty equivalents.

Managerial and Policy Implications

The idea that people actually make decisions that are sometimes inconsistent with decision-analysis principles is not new. In fact, the premise of a book such as this

Figure 14.11
The McCord-De Neufville utility-assessment procedure.


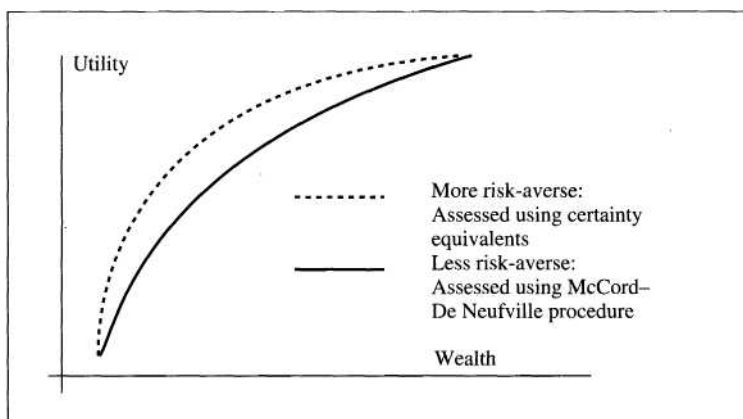


Figure 14.12
Comparing utility functions assessed with different approaches.



one is that it is possible to improve one's decision-making skills. But now we have seen that individuals behave in certain specific and predictable ways. What implications does this have for managers and policy makers?

The most fundamental issue has to do with the reference point or status quo. What is the status quo in a particular decision situation? How do people establish a status quo for decision-making purposes? Can the perception of the status quo be manipulated, and, if so, how? Complete answers to these questions are certainly beyond the scope of our discussion here, but we can discuss important examples in which the status quo plays an important role.

First is the problem of "sunk costs." As briefly mentioned earlier, managers frequently remain committed to a project that obviously has gone bad. In decision-analysis terms, the money that already has been spent on the project no longer should influence current and future decisions. Any decisions, particularly whether to continue or abandon a project, must be forward-looking — that is, consider only future cash flows. To account for the sunk costs, they would have to be accounted for on every branch and hence could not affect the relative ordering of the alternatives.

What do we say to an individual who seems unable to ignore sunk costs? One piece of advice is to make sure that the individual understands exactly what the status quo is. If the individual wants to "throw good money after bad," then he or she may be operating on the basis of a previous status quo; abandoning the project may look like accepting a sure loss of the money already invested. From this perspective, it might seem quite reasonable instead to seek a risky gain by remaining committed to the project. But the real status quo is that the project is unlikely to yield the anticipated benefits. From this new perspective, abandoning the project amounts to the avoidance of a losing venture; funds would be better invested elsewhere.

In other cases, a problem's frame may be specified with equal validity in several ways. Consider the case of seat belts. If people view seat belts as inconvenient and uncomfortable, then they may refuse to wear them; the status quo is the level of comfort when unbuckled. Suppose, however, the status quo were established as how well off people are in general. Who would want to risk the loss of a healthy family, a productive life, and possibly much of one's savings? The use of a seat belt is a form of inexpensive insurance to avoid a possible dramatic loss relative to this new status quo.

A similar argument can be made in the area of environmental protection. People may view environmental programs as anything from inconveniences (being forced to separate recyclable materials from trash) to major economic impediments (adhering to EPA's complex regulations). The implied status quo is the current level of direct costs; the programs represent a sure loss relative to this status quo. If, however, emphasis is given to the current overall condition of the community, nation, or world, then the increased cost induced by the programs can be viewed as a small amount of insurance necessary to avoid a large loss relative to our current overall welfare.

A third situation in which the perception of the status quo can be productively manipulated is in the area of creativity enhancement. In Chapter 6 we discussed removing blocks to creativity. In organizations, one key approach is to make the environment as nonthreatening as possible so that individuals will not be afraid to present

new ideas. In the terminology of this section, the goal is to establish a status quo in the organizational environment such that individuals will view the presentation of a new idea as one in which the loss is small if the idea fails but otherwise the potential gain is large.

In the discussion so far we have focused our attention on issues concerning the status quo. We turn now to the certainty effect, or the fact that we tend to overvalue sure outcomes relative to uncertain ones. This psychological phenomenon can have interesting effects on our decision making, especially under conditions of uncertainty.

In our western culture, we abhor uncertainty. Indeed, it is appropriate to expend resources for information in order to reduce uncertainty. This is true only up to a point, however, as we discussed in Chapter 12. But where may the certainty effect lead? We may try too hard to eliminate uncertainty altogether. For example, in organizations, there may be a tendency to spend far too much effort and far too many resources on tracking down elusive information in an effort to know the "truth." More insidious is a tendency to ignore uncertainty altogether. We often tend to view forecasts, for example, as perfect indicators of the future, ignoring their inherent uncertainty. The use of carefully assessed probabilistic forecasts is an important way to avoid this problem.

As a society, the certainty effect may be a factor in our preoccupation with certain kinds of health and environmental risks. Many individuals — activists in various causes — would prefer to eliminate risk altogether and thus they call for stringent regulations on projects that appear to pose risks to the population or environment. Certainly these individuals play a valuable role in drawing our attention to important issues, but achieving zero risk in our society is not only impractical, but also impossible. A sound approach to the management of risk in our society requires a consideration of both the benefits and costs of reducing risk to various levels.

In this brief section we have considered certain implications of the behavioral paradoxes for the application of decision analysis (especially utility assessment), as well as for organizational decisions and policy making. But we have only scratched the surface here. As research progresses in this fertile area, many other behavioral paradoxes and their implications will be studied. The examples we have considered illustrate the pervasiveness of the effects as well as their importance.

A Final Perspective

We have discussed a variety of inconsistencies in human behavior. Do these inconsistencies invalidate expected utility theory? The argument all along has been that people do not seem to make coherent decisions without some guidance. If a decision maker does wish to make coherent decisions, then a careful decision-analysis approach, including a careful assessment of personal preferences, can help the decision maker in looking for better decisions.

It is easy to get the impression that utility and probability numbers reside in a decision maker's mind and that the assessment procedure simply elicits those numbers.

But the process is more than that. Just as the process of structuring a decision problem helps the decision maker understand the problem better and possibly leads to the recognition of new alternatives for action, so may the assessment process provide a medium in which the decision maker actually can develop his or her preferences or beliefs about uncertainty. The assessment process helps to mold the decision maker's subjective preferences and beliefs. How many of us, for example, have given much thought to assessing a utility function for a cup of coffee (Problem 14.16) or considered the probability distribution for a grade in a decision-analysis course (Problem 8.11)? Certainly the assessed numbers do not already exist inside our heads; they come into being through the assessment procedure. This is exactly why assessment often requires hard thinking and why decisions are hard to make. Thus, perhaps the best way to think about the assessment process is in constructive terms. Reflecting on the decision problem faced and considering our preferences and beliefs about uncertainty provides a basis not only for building a model of the problem and necessary beliefs, but also for constructing those beliefs in the first place. This view may explain some of the behavioral paradoxes that we have discussed; individuals who have not thought long and hard in thoroughly developing their preferences and beliefs might have a tendency to make inconsistent judgments.

This constructive view of the assessment process is a fundamental matter. You may recall that we introduced these ideas way back in Chapter 1 when we suggested that the decision-analysis process actually helps a decision maker develop his or her thoughts about the structure of the problem, beliefs about uncertainty, and preferences. Again, the idea of a requisite model is appropriate. A decision model is requisite in terms of preferences if it captures the essential preference issues that matter to the decision maker for the problem at hand. Thus, this constructive view suggests that decision analysis should provide an environment in which the decision maker can systematically develop his or her understanding of the problem, including preferences and beliefs about uncertainty. But how can we be sure that decision analysis does provide such an environment? More research is required to answer this question completely. The decision-analysis approach, however, does encourage decision makers to think about the issues in a systematic way. Working through the decision-analysis cycle (modeling, analyzing, performing sensitivity analysis, and then modeling again) should help the decision maker to identify and think clearly about the appropriate issues. Thus, the argument is that good initial decision-analysis structuring of a problem will lead to appropriate assessments and careful thought about important issues.

SUMMARY

We have covered a lot of ground in this chapter. We started with the axioms that underlie utility theory and showed how those axioms imply that an individual should make decisions that are consistent with the maximization of expected utility. Then we examined some of the behavioral paradoxes that have been documented. These are situations in which intelligent people make decisions that violate one or more of the axioms, and thus make decisions that are inconsistent with expected utility.

These paradoxes do not invalidate the idea that we should still make decisions according to expected utility; recall that the basic goal of decision analysis is to help people improve their decision-making skills. But the paradoxes do have certain implications. We explored some of these implications, including the possibility of assessing equivalent lotteries rather than certainty equivalents and implications for policy makers. We ended with a constructive perspective of assessment, whereby it provides a medium within which the decision maker can explore and develop preferences and beliefs.

EXERCISES

- 14.1** In your own words explain why the axioms that underlie expected utility are important to decision analysis.
- 14.2** From a decision-analysis perspective, why is it worthwhile to spend time studying the kinds of paradoxes described in this chapter?
- 14.3** Most people are learning constantly about themselves and their environment. Our tastes develop and change as our environment changes. As new technologies arise, new risks are discovered. What are the implications of this dynamic environment for decision analysis?
- 14.4 a** Find the value for p that makes you indifferent between
- Lottery 1** Win \$ 1000 with probability p
 Win \$0 with probability $1 - p$
- and
- Lottery 2** Win \$400 for sure
- b** Now find the q that makes you indifferent between
- Lottery 3** Win \$400 with probability 0.50
 Win \$0 with probability 0.50
- and
- Lottery 4** Win \$ 1000 with probability q
 Win \$0 with probability $1 - q$
- c** According to your assessment in part a, $U(\$400) = p$. In part b, $U(\$400) = 2q$. Explain why this is the case.
- d** To be consistent, your assessments should be such that $p = 2q$. Were your assessments consistent? Would you change them? Which assessment do you feel most confident about? Why?

QUESTIONS AND PROBLEMS

14.5 We used the phrase "rules for clear thinking" to refer to the axioms described in this chapter. However, these "rules" do not cover every aspect of decision making as we have discussed it in this book. In particular, the axioms do not tell us in detail how to structure a problem, how to generate creative alternatives, or how to perform a sensitivity analysis. Can you give some "rules for clear thinking" that would help a decision maker in these aspects of decision making? (See Frisch and Clemen 1994.)

14.6 Imagine that you collect bottles of wine as a hobby. How would you react to these situations?

a You have just learned that one of the bottles (Wine A) that you purchased five years ago has appreciated considerably. A friend has offered to buy the bottle from you for \$100. His offer is a fair price. Would you sell the bottle of wine?

b A second bottle of wine (Wine B) is viewed by experts as being equivalent in value to Wine A. You never purchased a bottle of Wine B, but a casual acquaintance did. In conversation at a party, he offers to sell you his bottle of Wine B for \$100. Would you buy it?

Many people would neither sell Wine A nor buy Wine B. Explain why this pattern of choices is inconsistent with expected utility. What other considerations might be taken into account?

14.7 Consider the following two scenarios:

a You have decided to see a show for which tickets cost \$20. You bought your ticket in advance, but as you enter the theater, you find that you have lost the ticket. You did not make a note of your seat number, and the ticket is not recoverable or refundable in any way. Would you pay another \$20 for a second ticket?

b You have decided to see a show whose tickets cost \$20. As you open your wallet to pay for the ticket, you discover that you have lost a \$20 bill. Would you still buy a ticket to see the show?

Many individuals would not purchase a second ticket under the first scenario, but they would under the second. Explain why this is inconsistent with expected utility. How would you explain this kind of behavior?

14.8 Imagine yourself in the following two situations:

a You are about to go shopping to purchase a stereo system at a local store. Upon reading the newspaper, you find that another stereo store across town has the system you are interested in for \$1089.99. You had been planning to spend \$1099.95, the best price you had found after considerable shopping. Would you drive across town to purchase the stereo system at the other store?

b You are about to go shopping to purchase a popcorn popper at a local hardware store for \$19.95, the best price you have seen yet. Upon reading the paper, you discover that a department store across town has the same popper for \$9.99. Would you drive across town to purchase the popper at the department store?

Many people would drive across town to save money on the popcorn popper, but not on the stereo system. Why is this inconsistent with expected utility? What explanation can you give for this behavior?

- 14.9 A classified advertisement, placed by a car dealer, reads as follows:

NEW
1991 **BMW** 525i
Retail \$38,838.50
WAS \$32,998
NOW
\$32,498

- a What is your reaction to this advertisement?
- b In the same paper, a computer dealer advertised a \$500 discount on a computer system, marked down from \$3495 to \$2995. Which do you think is a better deal, the special offer on the car or the computer? Why?
- 14.10 Consider these two scenarios:
- a You have made a reservation to spend the weekend at the coast. To get the reservation, you had to make a nonrefundable \$50 deposit. As the weekend approaches, you feel a bit out of sorts. On balance, you decide that you would be happier at home than at the coast. Of course, if you stay home, you forfeit the deposit. Would you stay home or go to the coast? What arguments would you use to support your position?
- b You have decided to spend the weekend at the coast and have made a reservation at your favorite resort. While driving over, you discover a new resort. After looking it over, you realize that you would rather spend your weekend here. Your reservation at the coast can be canceled easily at no charge. Staying at the new resort, however, would cost \$50 more. Would you stay at the new resort or continue to the coast?
- Are your decisions consistent in parts a and b? Explain why or why not.
- 14.11 Even without a formal assessment process, it often is possible to learn something about an individual's utility function just through the preferences revealed by choice behavior. Two persons, A and B, make the following bet: A wins \$40 if it rains tomorrow and B wins \$10 if it does not rain tomorrow.
- a If they both agree that the probability of rain tomorrow is 0.10, what can you say about their utility functions?
- b If they both agree that the probability of rain tomorrow is 0.30, what can you say about their utility functions?
- c Given no information about their probabilities, is it possible that their utility functions could be identical?
- d If they both agree that the probability of rain tomorrow is 0.20, could both individuals be risk-averse? Is it possible that their utility functions could be identical? Explain.


Source: R. L. Winkler (1972) *Introduction to Bayesian Inference and Decision*. New York: Holt.

- 14.12** Assess your utility function for money in the bank over a range from \$100 to \$20,000 using the McCord-De Neufville procedure described in this chapter. Plot the results of your assessments on the same graph as the assessments made for Problem 13.12. Discuss the differences in your assessments from one method to the other. Can you explain any inconsistencies among them?
- 14.13** Assume that you are interested in purchasing a new model of a personal computer whose reliability has not yet been perfectly established. Measure reliability in terms of the number of days in the shop over the first three years that you own the machine. (Does this definition of reliability pass the clarity test?) Now assess your utility function for computer reliability over the range from 0 days (best) to 50 days (worst). Use whatever assessment technique with which you feel most comfortable, and use computer assessment aids if they are available.
- 14.14** You are in the market for a new car. An important characteristic is the lifespan of the car. (Define lifespan as the number of miles driven until the car breaks down, requiring such extensive repairs that it would be cheaper to buy an equivalent depreciated machine.) Assess your utility function for automobile life span over the range from 40,000 to 200,000 miles.
- 14.15** Being a student, you probably have well-developed feelings about homework. Given the same amount of material learned, the less the better, right? (I thought so!) Define homework as the number of hours spent outside of class on various assignments that enter into your final grade. Now, assuming that the amount of material learned is the same in all instances, assess your utility function for homework over the range from 0 hours per week (best) to 20 hours per week (worst). *{Hint: You may have to narrow the definition of homework. For example, does it make a difference what kind of course the homework is for? Does it matter whether the homework is term papers, case studies, short written assignments, oral presentations, or something else?}*
- 14.16** We usually think of utility functions as always sloping upward (more is better) or downward (less is better; fewer nuclear power plant disasters, for example). But this is not always the case. In this problem, you must think about your utility function for coffee versus milk.

Imagine that you are about to buy a cup of coffee. Let c ($0 \leq c \leq 1$) represent the proportion of the contents of the cup accounted for by coffee, and $1 - c$ the proportion accounted for by milk.

- a** Assess your utility function for c for $0 \leq c \leq 1$. Note that if you like a little milk in your coffee, the high point on your utility function may be at a value of c somewhere between 0 and 1.
- b** Compare (A) the mixture consisting of proportions c of coffee and $1 - c$ of milk in a cup and (B) the lottery yielding a cup of coffee with probability c and a cup of milk with probability $1 - c$. (The decision tree is shown in Figure 14.13.) Are the expected amounts of milk and coffee the same in A and B? [That is, if you calculate $E(c)$, is it the same in A and B?] Is there any value of c for which you are indifferent between A and B? (How about when c is 0 or 1?) Are you indifferent between A and B for the value of c at the high point of your utility function?
- c** How would you describe your risk attitude with respect to c ? Are you risk-averse or risk-prone, or would some other term be more appropriate?

Figure 14.13
Decision tree
for problem
14.16b.



- 14.17** In a court case, a plaintiff claimed that the defendant should pay her \$3 million for damages. She did not expect the judge to award her this amount; her expected value actually was \$1 million. The defendant also did not believe that the court would award the full \$3 million, and shared the plaintiff's expected value of \$1 million.
- Assuming that the plaintiff is thinking about the award in terms of gains, explain why you might expect her to be risk-averse in this situation. If she is risk-averse, what kind of settlement offer might she accept from the defendant?
 - Assuming that the defendant is thinking about the situation in terms of losses, explain why you might expect him or her to be risk-seeking. What would this imply about settlement offers to which the defendant might agree? (*Hint*: Draw an example of a risk-seeking utility curve over the possible negative payoffs for the defendant. Now find the certainty equivalent.)
 - Discuss your answers to parts a and b. What are the implications for settlements in real-world court cases? What would happen if the defendant's expected value were less than the plaintiff's?

CASE STUDIES

THE LIFE INSURANCE CASE

Peggy Ewen sat back in her chair and listened as Tom Pitman tried to explain. "I don't know what's going on," Tom said. "I have no trouble making the phone calls, and I seem to be able to set up in-home visits. I am making at least as many visits as anyone else in the office. For some reason, though, I cannot talk them into buying the product. I seem to be unlucky enough to have run into a lot of people who just are not interested in life insurance."

Peggy thought about this. Tom had been with the company for five months now. He was bright and energetic. He had gone through the training program easily, and had appeared to hit the ground running. His paperwork was always perfect. For some reason, though, his career selling life insurance was virtually stalled. His sales rate was only one-third that of the next best salesperson. Why?

Peggy asked, "How do you feel about going to the in-home visits?"

"Fine," Tom replied. "Well, I'm starting to feel a little apprehensive about it just because I'm becoming less sure of myself."

"Well, that's something we'll have to work on. But how do the visits go? What do you talk about? Tell me what a typical visit would be like."

"Let's see. Usually I'll come in, sit down, and we'll chat briefly. They'll offer me a cup of coffee. After a short visit, we get right down to business. I go through the presentation material provided by the company. Eventually we get around to talking about the reasons for purchasing life insurance. I really stress the importance of being able to make up for the loss of income. The presentation material stresses the idea of building up savings for sending kids to school or for retirement. You know, the idea of being sure that the extra money will be there down the road. But I really don't think that's why most people buy life insurance. I think they buy it to be sure that their family will be able to make up for a loss. For just a small premium, they can be sure that the loss won't happen, or at least they can minimize the loss."

Peggy seemed interested in Tom's account. "So you really stress the idea that for a little bit of money they can insure against the loss of income."

"Yes," Tom answered. "I'd rather have them look at life insurance as protection against a potential loss, rather than as a savings mechanism that would provide some sure amount in the future. Most of them know that there are better ways to save, anyway."

"And how would you classify your typical client? What kind of income bracket?"

"Mostly young couples just starting out," said Tom. "Maybe they've just had their first child. Not much income yet. Not much savings, either. We usually discuss this early on in the conversation. In general they seem to be quite aware of their financial situation. Occasionally they are even quite sensitive about it."

Peggy looked at Tom and grinned. "Tom, I do believe that there's something you can do right now to improve your sales rate."

Questions

- 1 About what issue is Peggy thinking?
- 2 What are the implications for people interested in selling financial securities such as insurance, annuities, and so on?
- 3 What are the implications for their customers?

NUCLEAR POWER PARANOIA

Ray Kaplan was disgusted. The rally against the power plant had gone poorly. The idea had been to get a lot of people out to show the utility company that the community did not support the idea of nuclear power. It was just too dangerous! Too much chance of an accident. Sure, the officials always pointed out that there had never

been a serious nuclear power plant accident in the United States. They always pointed to the safeguards in the plants, and to the safety features required by federal regulations. "You're safer in our plant than on the freeway," they claimed. Well, the same officials met with the protesters again today and said the same old things. Ray was getting tired of it. He hopped on his motorcycle and rode home. Fast.

He was still ruminating about the rally while he broiled his steak over the charcoal that evening. As he snacked on a bowl of ice cream after dinner, he asked himself, "Can't they see the potential for disaster?"

The next day, Ray decided to mow his lawn and then ride out to the beach. About a mile from his house, he realized that he wasn't wearing his motorcycle helmet. Well, he thought, the beach was only about 20 miles away. Besides, it would be nice to feel the wind in his hair. The ride was nice, even though the traffic was heavier than he had expected. The fumes from one of the trucks really got to him briefly, but fortunately the exit for the beach was right there. As he lay down on his towel to soak up the sunshine, his mind went back to the rally. "Darn," he thought. "I'll probably develop an ulcer just worrying about that silly power plant."

Question

- 1 What do you think about Ray Kaplan's behavior?

THE MANAGERS PERSPECTIVE

Ed Freeman just couldn't understand it. Why were the activists so blind? For years the information had been available showing just how safe nuclear power plants were. Study after study had concluded that the risk of fatalities from a nuclear power plant accident was far less than driving a car, flying in a commercial airliner, and many other commonplace activities in which people freely chose to engage. Sure, there had been some close calls, such as Three Mile Island, and there had been the terrible accident at Chernobyl in the Soviet Union. Still, the overall record of the nuclear power industry in the United States was excellent. No one could deny it if they would only compare the industry to others.

His risk assessors had gone through their own paces, documenting the safety features of his plant for the Nuclear Regulatory Commission; it was up to date and, in fact, one of the safest plants in the country. The experts had estimated the probability of an accident at the plant as nearly zero. Furthermore, even if an accident were to occur, the safety systems that had been built in would minimize the public's exposure.

Given all this, he just could not understand the public opposition to the plant. He knew that these were bright people. They were articulate, well read, and were able to marshal their supporters with great skill. But they seemed to ignore all of the data as well as the experts' reports and conclusions.

"I guess it takes all kinds," he sighed as he prepared to go back to work.

Questions

- 1 This case and "Nuclear Power Paranoia" go together. People often are willing to engage voluntarily in activities that are far more risky (in the sense of the probability of a serious injury or death) than living near a nuclear power plant. Why do you think this is the case?
- 2 What makes new technologies seem risky to you?

REFERENCES

The axioms of expected utility, along with the notion of subjective probability, were first discussed by Ramsey (1931), but the world appears to have ignored him. Most economists refer to "von Neumann-Morgenstern utility functions" because in 1947 von Neumann and Morgenstern published their celebrated *Theory of Games and Economic Behavior* in which they also set forth a set of axioms for choice behavior that leads to maximization of expected utility. The axioms subsequently appeared in a wide variety of forms and in many textbooks. Some examples are Luce and Raiffa (1957), Savage (1954), DeGroot (1970), and, more recently, French (1986). French's text is excellent for those interested in the axiomatic mathematics that underlie decision theory.

The various axioms have been debated widely. Our discussion of the transitivity axiom, for example, was suggested by Dr. Peter Fishburn as part of his acceptance speech for the Ramsey Medal, an award for distinguished contributions to decision analysis. If Peter Fishburn, one of the foremost scholars of decision theory, is willing to concede that intransitive preferences might not be unreasonable, perhaps we should pay attention! Fishburn (1989) summarizes many of the recent developments in the axioms and theory.

The text by von Winterfeldt and Edwards (1986) also contains much intriguing discussion of the axioms from the point of view of behavioral researchers and a discussion of the many paradoxes found in behavioral decision theory. Hogarth (1987) also covers this topic. Tversky and Kahneman (1981) provide an excellent and readable treatment of framing effects, and Kahneman and Tversky (1979) present a theory of behavioral decision making that accounts for many anomalies in individual decision behavior.

The fact that people do not normally follow the axioms perfectly has been the source of much debate about expected utility theory. Many theorists have attempted to relax the axioms in ways that are consistent with the observed patterns of choices that people make (e.g., Fishburn 1988). For the purpose of decision analysis, though, the question is whether we should model what people actually do or whether we should help them to adhere to axioms that are compelling but in some instances difficult to follow because of our own frailties. This debate is taken up by Rex Brown (1992) and Ron Howard (1992) and is a central theme of a collection of articles in Edwards (1992). Luce and von Winterfeldt (1994) make an argument for decision analysis in which the status quo is explicitly used as an essential element in decision models, and subjective assessments are made relative to the status quo.

The debate about axioms is fine as far as it goes, but it is important to realize that many aspects of decision making are not covered by the axioms. In particular, the axioms tell you what to do once you have your problem structured and probabilities and utilities assessed. The axioms do indicate that decisions problems can be decomposed into issues of value (utility) and uncertainty (probability), but no guidance is provided regarding how to determine the important dimensions of value or the important uncertainties in a decision situation. The axioms provide little help in terms of creativity, sensitivity analysis, or even

how to recognize a decision situation. Frisch and Clemen (1994) discuss this issue from the point of view of psychological research on decision making and the prescriptive goals of decision analysis.

The constructionist view of decision analysis that is presented in the last section of this chapter does not appear to be widely discussed, although such a view has substantial and fundamental implications for research in decision theory as well as in decision-analysis practice. For more discussion of this topic with regard to utility assessment, see von Winterfeldt and Edwards (1986, p. 356), Fischer (1979), and Payne, Bettman, and Johnson (1993). On the probability side, Shafer and Tversky (1986) view probability and structuring of inference problems in an interesting "constructive" way. Phillip's notion of the development of a requisite model (1982, 1984) and Watson and Buede's (1987) approach to decision analysis contain many of the elements of the constructionist view.

For many people, risk refers primarily to dangerous circumstances that may cause bodily injury, disease, or even death. The two cases, "Nuclear Power Paranoia" and "The Manager's Perspective," introduce the idea of risk to life and limb instead of risk in monetary gambles. Although one might think that measuring such risks would be a straightforward matter of assessing probability distributions for injuries or deaths — for example, with respect to options associated with building or siting a power plant — it turns out that people are very sensitive to certain kinds of risk. How individuals perceive risk has an important impact on how risks should be managed and how to communicate risk information to the public. For good introductory material, see Slovic (1987) and Morgan (1993).

Allais, M. (1953) "Le Comportement de l'Homme Rationnel Devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine." *Econometrica*, 21, 503-546.

Allais, M., and J. Hagen (eds.) (1979) *Expected Utility Hypotheses and the Allais Paradox*. Dordrecht, The Netherlands: Reidel.

Brown, R. (1992) "The State of the Art of Decision Analysis: A Personal Perspective." *Interfaces*, 22, 5-14.

DeGroot, M. (1970) *Optimal Statistical Decisions*. New York: McGraw-Hill.

Edwards, W. [ed.] (1992) *Utility Theories: Measurements and Applications*. Boston, MA: Kluwer.

Fischer, G. (1979) "Utility Models for Multiple Objective Decisions: Do They Accurately Represent Human Preferences?" *Decision Sciences*, 10, 451-479.

Fishburn, P. C. (1988) *Nonlinear Preference and Utility Theory*. Baltimore: Johns Hopkins.

Fishburn, P. C. (1989) "Foundations of Decision Analysis: Along the Way." *Management Science*, 35, 387-405.

French, S. (1986) *Decision Theory: An Introduction to the Mathematics of Rationality*. London: Wiley.

Frisch, D., and Clemen, R. (1994) "Beyond Expected Utility: Rethinking Behavioral Decision Research." *Psychological Bulletin*, 116, 46-54.

Hershey, J. C, H. C. Kunreuther, and P. J. H. Shoemaker (1982) "Sources of Bias in Assessment Procedures for Utility Functions." *Management Science*, 28, 936-954.

Hogarth, R. (1987) *Judgment and Choice*, 2nd ed. New York: Wiley.

Howard, R. (1992) "Heathens, Heretics, and Cults: The Religious Spectrum of Decision Aiding." *Interfaces*, 22, 15-27.

- Kahneman, D., and A. Tversky (1979) "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, *47*, 263-291.
- Luce, R. D., and H. Raiffa (1957) *Games and Decisions: Introduction and Critical Survey*. New York: Wiley.
- Luce, R. D., and D. von Winterfeldt (1994) "What Common Ground Exists for Descriptive, Prescriptive, and Normative Utility Theories?" *Management Science*, *40*, 263-279.
- McCord, M., and R. De Neufville (1986) "'Lottery Equivalents': Reduction of the Certainty Effect Problem in Utility Assessment." *Management Science*, *32*, 56-60.
- Morgan, G. (1993) "Risk Analysis and Management." *Scientific American*, July, 32-41.
- Payne, J., J. Bettman, and E. Johnson (1993) *The Adaptive Decision Maker*. Cambridge: Cambridge University Press.
- Phillips, L. D. (1982) "Requisite Decision Modelling." *Journal of the Operational Research Society*, *33*, 303-312.
- Phillips, L. D. (1984) "A Theory of Requisite Decision Models." *Acta Psychologica*, *56*, 29-48.
- Ramsey, F. P. (1931) "Truth and Probability." In R. B. Braithwaite (ed.) *The Foundations of Mathematics and Other Logical Essays*. New York: Harcourt, Brace.
- Savage, L. J. (1954) *The Foundations of Statistics*. New York: Wiley.
- Shafer, G., and A. Tversky (1986) "Languages and Designs for Probability Judgment." *Cognitive Science*, *9*, 309-339.
- Slovic, P. (1987) "Perception of Risk." *Science*, *236*, 280-285.
- Swalm, R. (1966) "Utility Theory — Insights into Risk Taking." *Harvard Business Review*, 123-136.
- Tversky, A., and D. Kahneman (1981) "The Framing of Decisions and the Psychology of Choice." *Science*, *211*, 453-458.
- von Neumann, J., and O. Morgenstern (1947) *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- von Winterfeldt, D., and W. Edwards (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press.
- Watson, S. R., and D. M. Buede (1987) *Decision Synthesis: The Principles and Practice of Decision Analysis*. Cambridge: Cambridge University Press.

EPILOGUE

We began the chapter with the Asian influenza example. This kind of study has been done repeatedly by many different experimenters, and the results are always the same; many of the subjects make inconsistent choices that depend on the framing of the problem. Of course, many of these experiments have been done using college students and other individuals who are not used to making this kind of decision. It would be nice to think that individuals who make difficult decisions often would not be susceptible to such inconsistencies. Unfortunately, such is not the case. Tversky and Kahneman (1981) report the same kinds of inconsistencies among decisions made by university faculty and physicians.

CHAPTER 15. Conflicting Objectives I: Fundamental Objectives and the Additive Utility Function

The utility functions for money that we have considered have embodied an important fundamental trade-off: monetary return versus riskiness. We have argued all along that the basic reason for using a utility function as a preference model in decision making is to capture our attitudes about risk and return. Accomplishing high returns and minimizing exposure to risk are two conflicting objectives, and we already have learned how to model our preference trade-offs between these objectives using utility functions. Thus, we already have addressed two of the fundamental conflicting objectives that decision makers face.

There are many other trade-offs that we make in our decisions, however. Some of the examples from Chapters 3 and 4 involved cost versus safety or fun versus salary. The FAA bomb-detection example balanced detection effectiveness, passenger acceptance, time to implement, and cost. Other examples are familiar: When purchasing cars or computers, we consider not only reliability and life span but also price, ease of use, maintenance costs, operating expenses, and so on. When deciding among school courses, you might be interested in factors such as the complementarity of material covered, importance in relation to your major and career goals, time schedule, and quality of the instructor. Still other examples and possible objectives appear in Figure 15.1 (abstracted from Keeney and Raiffa 1976).

As individuals, we usually can do a fair job of assimilating enough information so that we feel comfortable with a decision. In many cases, we end up saying things like, "Well, I can save some money now and buy a new car sooner," "You get what you pay for," and "You can't have everything." These are obvious intuitive statements that reflect the informal trade-offs that we make. Understanding trade-offs in detail, however, may be critical for a company executive who is interested in acquiring hundreds of personal computers or a large fleet of automobiles for the firm.

Figure 15.1
Four examples of decisions involving complicated preference trade-offs. Source: Keeney, R., and Raiffa, H. (1976) *Decisions with Multiple Objectives: Preference and Value Tradeoffs*. New York: John Wiley, pp. 1-4. Reprinted by permission.

<p>1 A mayor must decide whether to approve a major new electric power generating station. The city needs more power capacity, but the new plant would worsen the city's air quality. The mayor might consider the following issues:</p> <ul style="list-style-type: none"> • The health of residents • The economic conditions of the residents • The psychological state of the residents • The economy of the city and the state • Businesses • Local politics <p>2 Imagine the issues involved in the treatment of heroin addicts. A policy maker might like to:</p> <ul style="list-style-type: none"> • Reduce the size of the addict pool • Reduce costs to the city and its residents • Reduce crimes against property and persons • Improve the "quality of life" of addicts • Improve the "quality of life" of nonaddicts • Curb organized crime • Live up to the high ideals of civil rights and civil liberties • Decrease the alienation of youth • Get elected to higher political office 	<p>3 In choosing a site for a new airport near Mexico City, the head of the Ministry of Public Works had to balance such objectives as:</p> <ul style="list-style-type: none"> • Minimize the costs to the federal government • Raise the capacity of airport facilities • Improve the safety of the system • Reduce noise levels • Reduce access time to users • Minimize displacement of people for expansion • Improve regional development (roads, for instance) • Achieve political aims <p>4 A doctor prescribing medical treatment must consider a variety of issues:</p> <ul style="list-style-type: none"> • Potential health complications for the patient (perhaps death) • Money cost to the patient • Patient's time spent being treated • Cost to insurance companies • Payments to the doctor • Utilization of resources (nurses, hospital space, equipment) • Information gained in treating this patient (may be helpful in treating others)
---	---

In this chapter we will present a relatively straightforward way of dealing with conflicting objectives. Essentially, we will create an *additive* preference model; that is, we will calculate a utility score for each objective and then add the scores, weighting them appropriately according to the relative importance of the various objectives. The procedure is easy to use and intuitive. Computer programs are available that make the required assessment process fairly simple. But with the simple additive form comes limitations. Some of those limitations will be exposed in the problems at the end of the chapter. In Chapter 16, we will see how to construct more complicated preference models that are less limiting.

Where are we going? The first part of the chapter reviews some of the issues regarding identifying objectives, constructing objective hierarchies, and creating useful attribute scales. With objectives and attribute scales specified, we move on to the matter of understanding trade-offs. In the section titled "Trading Off Conflicting Objectives: The Basics," we will look at an example that offers a relatively simple choice involving three automobiles and two objectives. In this initial discussion, we will develop intuitive ways to trade off two conflicting objectives. The purpose of

this discussion is to introduce ideas, help you focus on the primary issues involved, and provide a framework for thinking clearly.

The next section formally introduces the additive utility function and explores some of its properties. In particular, we introduce indifference curves and the marginal rate of substitution between objectives. We also show that the simple multiobjective approach described in Chapters 3 and 4 is consistent with the additive utility function.

The additive utility function is composed of two different kinds of elements, scores on individual attribute scales and weights for the corresponding objectives. Many different methods exist for assessing the scores and the weights, and these are the topics of the following sections.

In the last section of the chapter, we consider an actual example in which the city of Eugene, Oregon, evaluates four candidate sites for a new public library. The example shows the process of defining objectives and attribute scales, rating the alternatives on each attribute scale, assessing weights, and then putting all of the assessments together to obtain an overall comparison of the four sites.

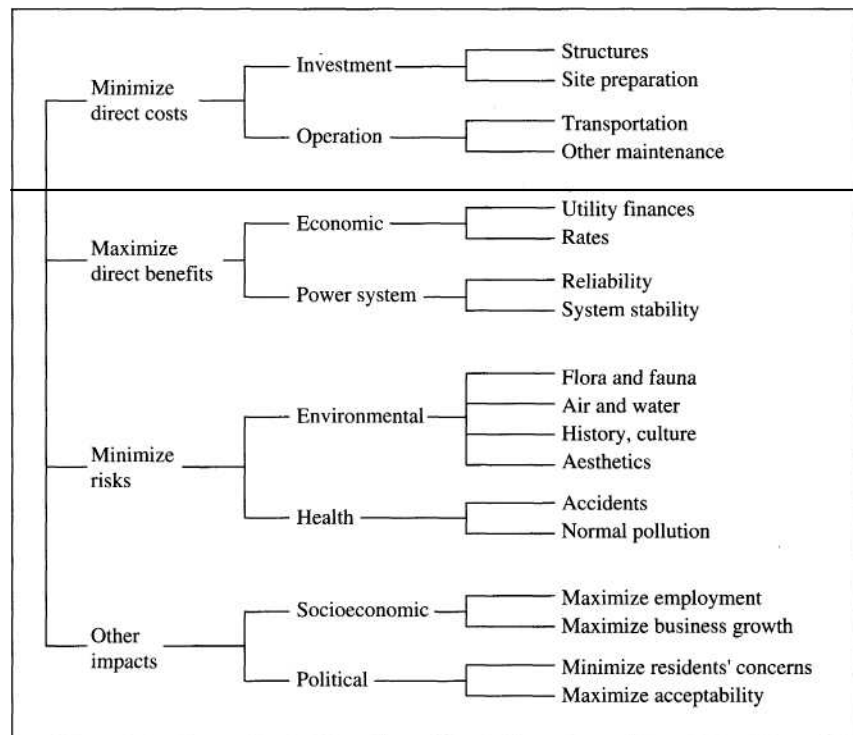
Objectives and Attributes

In Chapter 3 we discussed at length the notion of conflicting objectives in decision making. Understanding objectives is an important part of the structuring process. We stressed the importance of identifying fundamental objectives, the essential reasons that matter in any given decision context. Fundamental objectives are organized into a hierarchy in which the lower levels of the hierarchy explain what is meant by the higher (more general) levels. Figure 15.2 shows a fundamental-objectives hierarchy for evaluating alternative-energy technologies. We also discussed the notion of means objectives, which are not important in and of themselves but help, directly or indirectly, to accomplish the fundamental objectives. Distinguishing means and fundamental objectives is important because we would like to measure the available alternatives relative to the fundamental objectives, the things we really care about.

The discussion in Chapter 3 also introduced attribute scales. Attribute scales provide the means to measure accomplishment of the fundamental objectives. Some attribute scales are easily defined; if minimizing cost is an objective, then measuring cost in dollars is an appropriate attribute scale. Others are more difficult. In Figure 15.2, for example, there is no obvious way to measure risks related to aesthetic aspects of the environment. In cases like these, we discussed the use of constructed scales and related measurements as proxies.

For convenience, we use the term attribute to refer to the quantity measured on an attribute scale. For example, if an objective is to minimize cost, then the attribute scale might be defined in terms of dollars, and we would refer to dollar cost as the attribute for this objective. Use of the term "attribute" is common in the decision-analysis literature, and many authors use *multiattribute utility theory* (MAUT) to refer to topics covered in Chapters 15 and 16.

Figure 15.2
*Objectives hierarchy
 for evaluating alterna-
 tive-energy technolo-
 gies.* [Source: von
 Winterfeldt and
 Edwards (1986).]



There are a number of essential criteria for fundamental objectives and their attributes. Many of these we discussed in Chapters 3 and 4, and you are encouraged to review those sections. Here is an encapsulation:

- 1 The set of objectives, as represented by the fundamental-objectives hierarchy, should be complete; it should include all relevant aspects of a decision. The fact that important objectives are missing can be indicated by reluctance to accept the results of an analysis or simply the gnawing feeling that something is missing. If the results "just don't feel right," ask yourself what is wrong with the alternatives that the analysis suggests should be preferred. Careful thought and an honest answer should reveal the missing objectives.
- 2 At the same time, the set of objectives should be as small as possible. Too many objectives can be cumbersome and hard to grasp. Keep in mind that the objectives hierarchy is meant to be a useful representation of objectives that are important to the decision maker. Furthermore, each objective should differentiate the available alternatives. If all of the alternatives are equivalent with regard to a particular objective (as measured by its corresponding attribute), then that objective will not be of any help in making the decision.
- 3 The set of fundamental objectives should not be redundant. That is, the same objectives should not be repeated in the hierarchy, and the objectives should not be closely related.

- 4 As far as possible, the set of objectives should be *decomposable*. That is, the decision maker should be able to think about each objective easily without having to consider others. For example, in evaluating construction bids, the cost of the project and the amount of time required may be important attributes. In most cases, we can think about these attributes separately; regardless of the cost, it always would be preferable to complete the project sooner, and vice versa. Thus, the objectives would be decomposable into these two attributes, which can be considered independently. On the other hand, if you are deciding which course to take, you may want to choose the most interesting topic and at the same time minimize the amount of effort required. These attributes, however, are related in a way that does not permit decomposition; whether you want to put in a lot of effort may depend on how interested you are in the material. Hence, you may have to alter your set of objectives. For example, you might construct a scale that measures something like the extent to which the course inspires you to work harder and learn more.
- 5 Means and fundamental objectives should be distinguished. Even if fundamental objectives are difficult to measure and means objectives are used as proxies, it is important to remain clear on why the decision matters in the first place. Doing so can help avoid choosing an inappropriate alternative for the wrong reason.
- 6 Attribute scales must be operational. They should provide an easy way to measure performance of the alternatives or the outcomes on the fundamental objectives. Another way to put it is that the attribute scales should make it straightforward (if not easy) to fill in the cells in the consequence matrix.

Trading Off Conflicting Objectives: The Basics

The essential problem in multiobjective decision making is deciding how best to trade off increased value on one objective for lower value on another. Making these trade-offs is a subjective matter and requires the decision maker's judgment. In this section, we will look at a simple approach that captures the essence of trade-offs. We will begin with an example that involves only two objectives.

Choosing an Automobile: An Example

Suppose you are buying a car, and you are interested in both price and life span. You would like a long expected life span — that is, the length of time until you must replace the car — and a low price. (These assumptions are made for the purpose of this example; some people might enjoy purchasing a new car every three years, and for them a long life span may be meaningless.) Let us further suppose that you have narrowed your choices down to three alternatives: the Portalo (a relatively expensive sedan with a reputation for longevity), the Norushi (renowned for its reliability), and the Standard Motors car (a relatively inexpensive domestic automobile). You have

done some research and have evaluated all three cars on both attributes as shown in Table 15.1. Plotting these three alternatives on a graph with expected life span on the horizontal axis and price on the vertical axis yields Figure 15.3. The Portalo, Norushi, and Standard show up on the graph as three points arranged on an upward-sloping curve. That the three points are ordered in this way reflects the notion, "You get what you pay for." If you want a longer expected life span, you have to pay more money.


Occasionally alternatives may be ruled out immediately by means of a dominance argument. For example, consider a hypothetical car that costs \$15,000 and has an expected life span of seven years (Point A in Figure 15.3). Such a car would be a poor choice relative to the Norushi, which gives a longer expected life for less money. Thus, A would be dominated by the Norushi.

On the other hand, none of the cars under consideration is dominated. With this being the case, how can you choose? The question clearly is, "How much are you willing to pay to increase the life span of your car?" To answer this question, we will start with the Standard and assume that you will purchase it if the others are not better. Is it worthwhile to switch from the Standard to the Norushi? Note that the slope of the line connecting the Norushi and the Standard is \$666.67 per year. The switch would be worthwhile if you were willing to pay at least \$666.67 for each additional year of life span, or \$2000 to increase the expected life span by three years. Would you be willing to pay more than \$2000 to increase the expected life of your car by

Table 15.1
Automobile purchase alternatives.

	Portato	Norushi	Standard Motors
Price (\$ 1000s)	17	10	8
Life Span (Years)	12	9	6

Figure 15.3
Graph of three cars, comparing price and expected life span.



three years? This is a subjective assessment! If you would, then it is worthwhile to switch to the Norushi from the Standard. If not, do not switch.

For the sake of continuing the example, assume that you make the switch, and that you have decided that the Norushi is better for you than the Standard; you would pay at least \$2000 for the additional three years of life span. Now, should you switch to the Portalo? Notice that the slope of the line connecting the Norushi and the Portalo is \$2333.33 per year, or \$7000 for an additional three years. You were willing to pay at least \$666.67 for an extra year. Now, what about an extra \$7000 for another three years? Are the extra years of expected life span worth this much to you? If so, make the switch to the Portalo. If not, stick with the Norushi.

This simple procedure permits you to move systematically through the alternatives. The idea is to start at one corner of the graph (for example, the Standard) and then consider switching with each of the other alternatives, always moving in the same direction. Once a switch is made, there is never a need to reconsider the previous alternative. (After all, the one to which you changed must be better.) If there were many cars to choose from, the procedure would be to plot them all on a graph, eliminate the dominated alternatives, and then systematically move through the non-dominated set, beginning at the lower left-hand side of the graph and considering the alternatives that are to the upper right-hand side.

This procedure works well in the context of two conflicting objectives, and you can see how we are trading off the two. At each step we ask whether the additional benefit from switching (the increase in expected life span) is worth the cost (the increase in price). The same kind of procedure can be used with three or more attributes. The trade-offs become more complicated, however, and graphical interpretation is difficult.

The Additive Utility Function

The automobile example above is intended to give you an intuitive idea of how trade-offs work. This particular example is easy because it seems natural to many of us to think in terms of dollars, and we often can reduce nonmonetary attributes to dollars. But what if we wanted to trade off life span and reliability? We would like to have a systematic procedure that we can apply in any situation fairly easily. To do this, we must find satisfactory ways to answer two questions. The first question has to do with comparing the attribute levels of the available alternatives. We are comparing three different automobiles. How do they compare on the two attributes? Is the Portalo "twice as good" on life span as the Norushi? How does the Standard compare (quantitatively) to the Portalo on price? In the energy example in Figure 15.2, alternative technologies must be ranked on each of their attributes. For example, substantial differences may exist among these technologies on all of the detail attributes. To get anywhere with the construction of a quantitative model of preferences, we must assess numerical scores for each alternative that reflect the comparisons.

The second question asks how the attributes compare in terms of importance. In the automobile example, is life span twice as important as price, and exactly what does "twice as important" mean? In the alternative-energy example, how do the environmental and health risks compare in terms of importance within the "Minimize Risks" fundamental objective? As with the scores, numerical weights must be assessed for each attribute.

The model that we will adopt in this chapter is called the *additive utility function*. We assume that we have individual utility functions $U_1(x_1), \dots, U_m(x_m)$ for m different attributes x_1 through x_m . These are utility functions just like those that we discussed in Chapters 13 and 14. In particular, we assume that each utility function assigns values of 0 and 1 to the worst and best levels on that particular objective. The additive utility function is simply a weighted average of these different utility functions. For an outcome that has levels x_1, \dots, x_m on the m objectives, we would calculate the utility of this outcome as

$$\begin{aligned} U(x_1, \dots, x_m) &= k_1 U_1(x_1) + \dots + k_m U_m(x_m) \\ &= \sum_{i=1}^m k_i U_i(x_i) \end{aligned} \tag{15.1}$$

where the weights are k_1, \dots, k_m . All of the weights are positive, and they add up to 1. First, you can see this additive utility function also assigns values of 0 and 1 to the worst and best conceivable outcomes, respectively. To see this, look at what happens when we plug in the worst level (x_i^-) for each objective. The individual utility functions then assign 0 for each objective [$U(x_i^-) = 0$], and so the overall utility is also 0. If we plug in the best possible value for each objective (x_i^+), the individual utility functions are equal to 1 [$U(x_i^+) = 1$], and so the overall utility becomes

$$\begin{aligned} U(x_1^+, \dots, x_m^+) &= k_1 U_1(x_1^+) + \dots + \\ & k_m U_m(x_m^+) = k_1 + \dots + k_m = 1 \end{aligned}$$

You can see that this utility function is consistent with what we developed in Chapters 3 and 4. In Chapter 3 we discussed how to assess "scores" by using attribute scales; you can see that the individual utility functions U_i in Equation (15.1) define the attribute scales. In Chapters 3 and 4 we had the attribute scales ranging from 0 to 100, whereas here we have defined the utility functions to range from 0 to 1. The transformation is simple enough; think of the 0 and 100 in Chapters 3 and 4 as 0% and 100%. In the case of quantitative natural scales we showed in Chapter 4 how to standardize those to a scale from 0 to 100, and in the case of constructed scales we had to make a subjective assessment of the value of the intermediate levels on the scale from 0 to 100. As we proceed through this chapter, we will examine some alternative ways of coming up with the individual utility functions.

In Chapter 4 we turned to the assessment of weights. We used a simple technique of comparing the importance of the range of one attribute with another. Later in this chapter we will discuss this and several other approaches for assessing weights.

Choosing an Automobile: Proportional Scores

To understand the additive utility function better, let us work through a simple example with the automobile choice. We will begin with the determination of the individual utility values, following the same proportional-scoring technique discussed in Chapter 4 (except that we will have the utility functions range from 0 to 1 rather than 0 to 100). The first step is easy. The Standard is best on price and worst on life span, so assign it a 1 for price and a 0 for life span: $U_P(\text{Standard}) = 1$ and $U_L(\text{Standard}) = 0$, where the subscripts P and L represent price and life span, respectively. Do the opposite for the Portalo, which is worst on price and best on life span: $U_P(\text{Portalo}) = 0$ and $U_L(\text{Portalo}) = 1$. Now, how do we derive the corresponding scores for the Norushi? Because we have natural numerical measures for the objectives (dollars and years), we can simply scale the Norushi's price and life span. A general formula is handy. Calculate

$$U_i(x) = \frac{x - \text{Worst Value}}{\text{Best Value} - \text{Worst Value}}$$

$$= \frac{x - x_i^-}{x_i^+ - x_i^-}$$

For the Norushi's price,

$$U_P(\$10,000) = \frac{10,000 - 17,000}{8000 - 17,000}$$

$$= 0.78$$

Likewise, the Norushi's utility for life span is

$$U_L(9 \text{ years}) = \frac{9 - 6}{12 - 6}$$

The intuition behind these calculations is that 9 years is exactly halfway between 6 and 12 years [thus $U_L(\text{Norushi}) = 0.50$], whereas \$10,000 is 78% of the way from \$17,000 to \$8000. Moreover, you can see that the values of 0 and 1, which we assigned to the Standard and the Portalo, are consistent with the general formula; plugging in the worst value for x yields $U_i(\text{Worst}) = 0$, and plugging in the best value gives $U_i(\text{Best}) = 1$. The utilities for the three cars are summarized in Table 15.2. As long as the objectives have natural numerical attributes, it is a straightforward matter to scale those attributes so that the utility of the best is 1, the utility of the worst is 0, and the intermediate alternatives have scores that reflect the relative distance between the best and worst.

Table 15.2
Utilities for three cars
on two attributes.

	Portalo	Norushi	Standard Motors
Price (U_P)	0.00	0.78	1.00
Life Span (U_L)	1.00	0.50	0.00

Assessing Weights: Pricing Out the Objectives

Now we must assess the weights for price and life span. But before we decide on the weights once and for all, let us look at the implications of various weights. For the automobile example, we must assess k_P and k_L which represent the weights for price and life span, respectively.

Suppose you were to decide that price and expected life span should be weighted equally, or $k_P = k_L = 0.5$. In general, we are going to calculate

$$U(\text{Price, Life Span}) = k_P U_P(\text{Price}) + k_L U_L(\text{Life Span})$$

Thus, the weighted utilities would be

$$U(\text{Portalo}) = 0.5(0.00) + 0.5(1.00) = 0.50$$

$$U(\text{Norushi}) = 0.5(0.78) + 0.5(0.50) = 0.64$$

$$U(\text{Standard}) = 0.5(1.00) + 0.5(0.00) = 0.50$$

The Standard and the Portalo come out with exactly the same overall utility because of the way that price and life span are traded off against each other. Because the difference between 1 and 0 amounts to \$9000 in price versus six years in life span, the equal weight in this case says that one additional year of life span is worth \$1500. The Norushi comes out on top because you pay less than \$1500 per year for the three additional years in expected life span as compared to the Standard.

Suppose that you have little money to spend on a car. Then you might think that price should be twice as important as life span. To model this, let $k_P = 0.67$ and $k_L = 0.33$. Now the overall utilities for the cars are Portalo, 0.33; Norushi, 0.69; and Standard, 0.67. In this case the weights imply that an increase in life span of one year is only worth an increase in price of \$750. (You can verify this by calculating the utility for a car that costs \$8750 and is expected to last seven years; such a car will have the same weighted score as the Standard.) Again the Norushi comes out as being preferred to the Standard, because its three-year increase in life span (relative to the Standard) is accompanied by only a \$2000 increase in price, whereas the weights indicate that the additional three years would be worth \$2250.

You may not be happy with either scheme. Perhaps you have thought carefully about the relative importance of expected life span and price, and you have decided that you would be willing to pay up to \$600 for an extra year of expected life span. You have thus *priced out* the value of an additional year of expected life span. How can you translate this price into the appropriate weights? Take the Standard as your base case (although any of the three automobiles could be used for this). Essentially, you are saying that you would be indifferent between paying \$8000 for six years of expected life span and \$8600 for seven years of expected life span. Using Equation (15.2), we can find that such a hypothetical car (Car B) would score $1/6 = 0.167$ on expected life span (which is one-sixth of the way from the worst to the best case) and 0.933 on price (\$8600 is 0.933 of the way from the worst to the best case). Because you would be indifferent between the Standard and the hypothetical Car B, the weights must satisfy

$U(\text{Standard}) = U(\text{Car B}) = k_P(1.00) + k_L(0) = k_P(0.933) + k_L(0.167)$ Simplify this equation to find that

$$k_P(1.00 - 0.933) = k_L(0.167)$$

$$k_P = k_L \frac{0.167}{0.067}$$

or that

$$k_P = 2.50k_L$$

Including the condition that the weights must sum to 1, we have

$$k_P = 2.50(1 - k_P)$$

or

$$k_P = 0.714 \quad \text{and} \quad k_L = 0.286$$

Note that these weights are consistent with what we did above. The weight $k_P = 0.667$ implied a price of \$750 per additional year of expected life span. With a still lower price (\$600), we obtained a higher weight for k_P .

The final objective, or course, is to compare the cars in terms of their overall utilities:

$$U(\text{Portalo}) = 0.714(0.00) + 0.286(1.00) = 0.286$$

$$U(\text{Norushi}) = 0.714(0.78) + 0.286(0.50) = 0.700$$


$$U(\text{Standard}) = 0.714(1.00) + 0.286(0.00) = 0.714$$

The Standard comes out only slightly better than the Norushi. This is consistent with the switching approach described earlier. The weights here came from the assessment that one year of life span was worth only \$600, not the \$666.67 or more required to switch from the Standard to the Norushi.

Indifference Curves

The assessment that you would trade \$600 for an additional year of life span can be used to construct *indifference curves*, which can be thought of as a set of alternatives (some perhaps being hypothetical) among which the decision maker is indifferent. For example, we already have established that you would be indifferent between the Standard and hypothetical Car B, which costs \$8600 and lasts seven years. Thus, in Figure 15.4 we have a line that passes through the points for the Standard and the point for the hypothetical Car B (Point B). All of the points along this line represent cars that would be equivalent to the Standard; all would have the same utility, 0.714. Other indifference curves also are shown with their corresponding utilities. Note that the indifference curves have higher utilities as one moves down and to the right, because you would rather pay less money and have a longer life span. You can see that the Norushi and the Portalo are not preferred to the Standard because they lie above the 0.714 indifference curve.

Figure 15.4
Indifference curves for
the automobile
decision.



The slope of the indifference curves in Figure 15.4 is related to the trade-off rate that was assessed. Specifically, the slope is \$600 per year, the price that was assessed for each year of expected life span. This also is sometimes called the *marginal rate of substitution*, or the rate at which one attribute can be used to replace another.

When using the additive utility function, it is straightforward to calculate how much a utility point in objective *i* is worth in terms of objective *j*. Let us say that you want to know how much one utility unit of attribute *i* is worth in terms of utility units of attribute *j*. Then the marginal rate of substitution between *i* and *j* is simply k_i/k_j . Thus, in the automobile case, the marginal rate of substitution in terms of the utility scales is $0.286/0.714 = 0.40$. In other words, the increase of one utility point on the life-span scale is worth 40% of the increase of one point on the price scale.

Unfortunately, knowing the marginal rate of substitution in utility terms is not so useful. If you are using the additive utility function and proportional scores as we have done here, then the marginal rate of substitution in terms of the original attributes is easily calculated. Let M_{ij} denote the marginal rate of substitution between attributes *i* and *j*. Then

$$M_{ij} = \frac{k_i/|x_i^+ - x_i^-|}{k_j/|x_j^+ - x_j^-|} \tag{15.3}$$

Thus, for the automobiles M_{LP} can be calculated as

$$\begin{aligned} M_{LP} &= \frac{0.286/|12 \text{ yr} - 6 \text{ yr}|}{0.714/|\$8000 - \$17,000|} \\ &= \$600 \text{ per year} \end{aligned}$$

In general, without proportional scores or with a nonadditive overall utility function, the marginal rate of substitution can vary depending on the values of the attributes

x_i and x_j , and the reason is that the indifference curves may be just that: curves instead of straight lines. Moreover, calculating M_{ij} can be difficult, requiring calculus to determine the slope of an indifference curve at a particular point. It is often possible to graph approximate indifference curves, however, and doing so can provide insight into one's assessed trade-offs. We will return to indifference curves in Chapter 16.

Assessing Individual Utility Functions


As you can see, the essence of using the additive utility function is to be able to assess the individual utility functions and the weights. In this section we look at some issues regarding the assessment of the individual utility functions.

Proportional Scores

In the example here and in Chapter 4, we have used the proportional scoring approach. In fact, we have just taken the original values and scaled them so that they now range from 0 (worst) to 1 (best). For the automobile example, Figure 15.5 graphs the price relative to a car's score on price. The graph shows a straight line, and we know about straight lines in this context: They imply risk neutrality and all of the unusual behavior associated with risk neutrality.

Let us think about what risk neutrality implies here with another simple example. Imagine that you face two career choices. You could decide to invest your life sav-

Figure 15.5
Proportional scores
for prices of three
cars.



ings in an entrepreneurial venture, or you could take a job as a government bureaucrat. After considering the situation carefully, you conclude that your objectives are purely monetary, and you can think in terms of income and savings. (Actually, more than just monetary outcomes should influence your career choice!) The bureaucratic job has a well-defined career path and considerable security. After 10 years, you know with virtual certainty that you will make \$30,000 per year and have \$60,000 in the bank toward retirement. On the other hand, becoming an entrepreneur is a risky proposition. The income and savings outcomes could range anywhere from zero dollars in the case of failure to some large amount in the event of success. To solve the problem, you decide to assess a continuous distribution and fit a discrete approximation. The decision tree in Figure 15.6 shows your probability assessments and alternatives. You can see that for the entrepreneurial alternative, the expected values are \$30,000 in income and \$60,000 in savings, the same as the certain values associated with the bureaucratic job.

Now we assign scores to the outcomes. For income, the best is \$60,000, the worst is 0, and \$30,000 is halfway between. These outcomes receive scores of 1, 0, and 1/2, respectively. Performing the same analysis for the savings dimension results in similar scores, and the decision tree with these scores appears in Figure 15.7.

We now calculate expected scores for the two alternatives; this requires calculating scores for each of the three possible outcomes. Let k_I and k_S represent the weights that we would assign to income and savings. The scores (U) for the possible outcomes are

Figure 15.6
A career decision.





Figure 15.7
A career decision with proportional scores.



$$U(\$60,000, \$120,000) = k_I(1) + k_S(1) = k_I + k_S$$

$$U(\$30,000, \$60,000) = k_I(1/2) + k_S(1/2) =$$

$$1/2 \{k_I + k_S\} \quad U(0,0) = k_I(0) + k_S(0) = 0$$

Now we can calculate the expected utility (EU) for the entrepreneurial venture. That expected score is obtained by averaging the scores over the three branches:


$$\begin{aligned} \text{EU}(\text{Entrepreneur}) &= 0.185(k_I + k_S) + 0.63[1/2(k_I + k_S)] + 0.185(0) = \\ &0.50(k_I + k_S) = 0.50 \end{aligned}$$

because $k_I + k_S = 1$. Of course, the expected score for the bureaucratic option also is 0.50, and so, on the basis of expected scores, the two alternatives are equivalent. (In fact, it is important to note that the two alternatives have the same score regardless of the specific values of k_I and k_S .) It is obvious, however, that being an entrepreneur is riskier. If each job really has the same expected scores in all relevant dimensions, would you be indifferent? (Probably not, but we will not guess about whether you would choose the riskier job or the more secure one!)

Ratios

Another way to assess utilities — and a particularly appropriate one for attributes that are not naturally quantitative — is to assess them on the basis of some ratio comparison. For example, let us return to the automobile example. Suppose that color is an important attribute in your automobile purchase decision. Clearly, this is not something that is readily measurable on a meaningful numerical scale. Using a ratio approach, you might conclude that to you blue is twice as good as red and that yellow is 2 1/2 times as good as red. We could accomplish the same by assigning some number of points between 0 and 100 to each possible alternative on the basis of performance on the attribute. In this way, for example, you might assign 30 points to red, 60 points to blue, and 75 points to yellow. This could be represented graphically as in Figure 15.8.

Figure 15.8
Graphically scoring
alternatives.



Now, however, we must scale these assessments so that they range from 0 to 1. We need to find constants a and b so that

$$0 = a + b(30)$$

$$1 = a + b(75)$$

Solving these two equations simultaneously gives

$$a = -\frac{2}{3}$$

$$b = \frac{1}{45}$$

Applying these scaling constants, we can calculate U_C , the utilities for the three colors:

$$U_C(\text{Red}) = -\frac{2}{3} + \frac{30}{45} = 0$$

$$U_C(\text{Blue}) = -\frac{2}{3} + \frac{60}{45} = \frac{2}{3}$$


$$U_C(\text{Yellow}) = -\frac{2}{3} + \frac{75}{45} = 1$$

Figure 15.9 shows the scaled scores, which now represent your relative preference for the different colors. They may be used to calculate weighted scores for different cars in a decision problem in which color is one attribute to consider. For example, with appropriate trade-off weights for price, color, and life span, the weighted score for a blue Portalo would be

$$U(\$17,000, 12 \text{ Years}, \text{Blue}) = k_P(0) + k_L(1) + k_C(0.667)$$

You can see how the ratio approach can be used to compare virtually any set of alternatives whether or not they are quantitatively measured.

Figure 15.9
Scaled scores for
colors.



Standard Utility-Function Assessment

In principle, assessing the individual utility functions need be no more complicated than the assessment procedures described in Chapters 13 and 14. For the automobile example, we would need to assess utility functions for price, life span, and color. If you face a decision that is complicated by both uncertainty and trade-offs, this is the best solution. The utility functions, having been assessed in terms of your preferences over uncertain situations, are models of your preferences in which your risk attitude is built in.

To see the effect of building in a risk attitude, let us return to the example above about the two jobs. Suppose we assess utility functions for income and savings and find that

$$\begin{aligned}U_I(\$60,000) &= 1 \\U_I(\$30,000) &= 0.75 \\U_I(0) &= 0\end{aligned}$$

and

$$\begin{aligned}U_S(\$120,000) &= 1 \\U_S(\$60,000) &= 0.68 \\U_S(0) &= 0\end{aligned}$$

With these utilities, the expected score for becoming an entrepreneur is

$$\begin{aligned}EU(\text{Entrepreneur}) &= 0.185(1k_I + 1k_S) + 0.63(0.75k_I + 0.68k_S) \\&+ 0.185(0k_I + 0k_S) = 0.658k_I + 0.613k_S\end{aligned}$$

Now compare $EU(\text{Entrepreneur})$ with $U(\text{Bureaucrat})$:

$$U(\text{Bureaucrat}) = 0.750k_I + 0.680k_S$$

Regardless of the specific values for the weights k_I and k_S , the expected score for being a bureaucrat always will be greater than the expected score for being an entrepreneur because the coefficients for k_I and k_S are larger for the bureaucrat than for the entrepreneur. This makes sense; the expected values (EMVs) for the two alternatives are equal for each attribute, but being an entrepreneur clearly is riskier. When we use a utility function for money that incorporates risk aversion, the riskier alternative ends up with a lower expected utility.

Assessing Weights

In the examples here and in Chapter 4, we have seen two different methods of assessing weights for the additive utility function. The approach in the automobile example is called *pricing out* because it involves determining the value of one objec-

tive in terms of another (usually dollars). The summer-job example in Chapter 4, though, used an approach that we will call *swing weighting*. We will discuss both of these methods here along with a third method based on a comparison of lotteries.

Pricing Out

The pricing-out method for assessing weights is no more complicated than what was described above in the automobile example. The essence of the procedure is to determine the marginal rate of substitution between one particular attribute (usually monetary) and any other attribute. Thus, we might say that a year of life span is worth \$600. Or, in a decision regarding whether to accept a settlement offer or pursue a lawsuit, one might assess that it would be worth \$350 in time and frustration to avoid an hour consulting with lawyers, giving depositions, or sitting in a courtroom. (And that would be on top of the cost of the lawyers' time!)

Assessing the marginal rate of substitution is straightforward in concept. The idea is to find your indifference point: the most you would pay for an incremental unit of benefit, or the least you would accept for an incremental unit of something undesirable. Although this idea seems straightforward, it can be a difficult assessment to make, especially for attributes with which we have little buying or selling experience.

As you can see, pricing out is especially appropriate for direct determination of the marginal rate of substitution from one attribute scale to another. Because the additive utility function implies a constant marginal rate of substitution, pricing out is consistent with the notion of proportional scores. However, when the individual utility functions are nonlinear or not readily interpretable in terms of an interval scale (as they might not be with a constructed scale), then pricing out makes less sense. One would have to ask, "How much utility on Attribute Scale 1 would I be willing to give up for a unit increase on Attribute Scale 2?" This is a more difficult assessment, and in situations like this, one of the following weight-assessment procedures might be easier.

Swing Weighting

The swing-weighting approach can be used in virtually any weight-assessment situation. It requires a thought experiment in which the decision maker compares individual attributes directly by imagining (typically) hypothetical outcomes. We demonstrate the procedure in this section with the automobile example.

To assess swing weights, the first step is to create a table like the one in Table 15.3 for the automobiles. The first row indicates the worst possible outcome, or the outcome that is at the worst level on each of the attributes. In the case of the automobiles, this would be a red car that lasts only six years and costs \$17,000. This "worst case" provides a benchmark. Each of the succeeding rows "swings" one of the attributes from worst to best. For example, the second row in Table 15.3 is for a red car that lasts 12 years and costs \$17,000. The last row is for a car that is worst on price and life span (\$17,000 and six years, respectively) but is your favorite color, yellow.

With the table constructed, the assessment can begin. The first step is to rank-order the outcomes. You can see in Table 15.3, a "4" has been placed in the "Rank"

Table 15.3
Swing-weight assessment table for automobile example.

Attribute Swung from Worst to Best	Consequence to Compare	Rank	Rate	Weight
(Benchmark)	6 years, \$17,000, red	4		
Life span	12 years, \$17,000, red	—	—	—
Price	6 years, \$8000, red	—	—	—
Color	6 years, \$17,000, yellow	—	—	—

column for the first row. There are four hypothetical cars to compare, and it is safe to assume that the benchmark car — the one that is worst on all the objectives — will rank fourth (worst) overall. The others must be compared to determine which ranks first (best), second, and third. Suppose that after some thought you conclude that the low-price car is the best, then the long life-span car, and finally the yellow car. Table 15.4 shows the partially completed table.

The next step is to fill in the "Rate" column in the table. You can see in Table 15.4 that two of the ratings are predetermined; the rating for the benchmark car is 0 and the rating for the top-ranked car is 100. The ratings for the other two cars must fall between 0 and 100. The comparison is relatively straightforward to make; how much less satisfaction do you get by swinging life span from 6 to 12 years as compared to swinging price from \$17,000 to \$8000? What about swinging color from red to yellow as compared to swinging price? You can even think about it in percentage terms; considering the increase in satisfaction that results from swinging price as 100%, what percentage of that increase do you get by swinging life span from worst to best?

Suppose that after careful thought, your conclusion is to assign 75 points to life span and 10 points to color. Essentially, this means that you think improving life span from worst to best is worth 75% of the value you get by improving the price from \$17,000 to \$8000. Likewise, changing the color from yellow to red is worth only 10% of the improvement in price. With these assessments, the table can be completed and weights calculated. Table 15.5 shows the completed table. The weights are the normalized ratings; recall that by convention we have the weights add up to 1. For example, k_p is calculated as $100/(100 + 75 + 10) = 0.541$. Likewise, we have $k_L = 75/(100 + 75 + 10) = 0.405$, and $k_C = 10/(100 + 75 + 10) = 0.054$.

Table 15.4
Swing-weight assessment table with ranks assessed.

Attribute Swung from Worst to Best	Consequence to Compare	Rank	Rate	Weight
(Benchmark)	6 years, \$17,000, red	4	0	
Life span	12 years, \$17,000, red	2		
Price	6 years, \$8000, red	1	100	
Color	6 years, \$17,000, yellow	3		

Table 15.5
Completed swing-weight assessment table.

Attribute Swung from Worst to Best	Consequence to Compare	Rank	Rate	Weight
(Benchmark)	6 years, \$17,000, red	4	0	
Life span	12 years, \$17,000, red	2	75	$0.405 = 75/185$
Price	6 years, \$8000, red	1	100	$0.541 = 100/185$
Color	6 years, \$17,000, yellow	3	10	$0.054 = 10/185$
	Total		185	1.000

With the weights determined, we can calculate the overall utility for different alternatives or outcomes. For example, we now can finish calculating the utility for a blue Portaló:

$$\begin{aligned} U(\$17,000, 12 \text{ Years, Blue}) &= k_P(0) + k_L(1) + k_C(0.667) \\ &= 0.541(0) + 0.405(1) + 0.054(0.667) = \\ &0.441 \end{aligned}$$

Why do swing weights work? The argument is straightforward. Here are the utilities for the hypothetical cars that you have considered:


$$\begin{aligned} U(\text{Worst Conceivable Outcome}) &= U(\$17,000, 6 \text{ Years, Red}) \\ &= k_P(0) + k_L(0) + k_C(0) = \\ &0 \end{aligned}$$

$$\begin{aligned} U(\$8000, 6 \text{ Years, Red}) &= k_P(1) + k_L(0) + k_C(0) = k_P \\ U(\$17,000, 12 \text{ Years, Red}) &= k_P(0) + k_L(1) + k_C(0) = \\ k_L U(\$17,000, 6 \text{ Years, Yellow}) &= k_P(0) + k_L(0) + k_C(1) \\ &= k_C \end{aligned}$$

From the first two equations, you can see that the increase in satisfaction from swinging price from worst to best is just k_P . Likewise, the improvement from swinging any attribute from worst to best is simply the value of the corresponding weight. When you compare the relative improvements in utility by swinging the attributes one at a time, you are assessing the ratios k_L/k_P and k_C/k_P . These assessments, along with the constraint that the weights add to 1, allow us to calculate the weights. Figure 15.10 graphically shows how swing weights work.

Swing weights have a built-in advantage in that they are sensitive to the range of values that an attribute takes on. For example, suppose you are comparing two personal computers, and price is an attribute. One computer costs \$3500 and the other \$3600. When you work through the swing-weight assessment procedure, you probably will conclude that the increase in utility from swinging the price is pretty small. This would result, appropriately, in a small weight for price. But if the difference in price is \$1000 rather than \$100, the increase in utility experienced by swinging from worst to best would be much larger, resulting in a larger weight for price.

Figure 15.10
Graphic representation of swing-weighting procedure.



If you have a hard time thinking about the "Worst Conceivable Outcome," you might try reversing this procedure. That is, use as a benchmark the "Best Conceivable Outcome," best on all attributes, and consider decreases in satisfaction from swinging attributes from high to low. Assess relative decreases in satisfaction, and use those assessments in exactly the same way that we used the relative utility increases.

Lottery Weights

It should come as no surprise that we also can use lottery-comparison techniques to assess weights. In fact, the technique we will use is a version of the probability-equivalent assessment technique introduced in Chapter 13. The general assessment setup is shown in Figure 15.11.

The assessment of the probability p that makes you indifferent between the lottery and the sure thing turns out to be the weight for the one odd attribute in the sure thing. We will see how this works in the case of the automobiles. Figure 15.12 shows the assessment decision for determining the weight associated with price.

Suppose that the indifference probability in Figure 15.12 turns out to be 0.55. Write down the equation that is implied by the indifference:

$$\begin{aligned}
 &k_P U_P(\$8000) + k_L U_L(6 \text{ Years}) + k_C U_C(\text{Red}) \\
 &= 0.55[k_P U_P(\$8000) + k_L U_L(12 \text{ Years}) + k_C U_C(\text{Yellow})] + \\
 &0.45[k_P U_P(\$17,000) + k_L U_L(6 \text{ Years}) + k_C U_C(\text{Red})]
 \end{aligned}$$

As before, the individual utilities range from 0 to 1. This means that

Figure 15.11
Assessing weights using a lottery technique. The task is to assess the probability p that makes you indifferent between the lottery (A) and the sure thing (B).


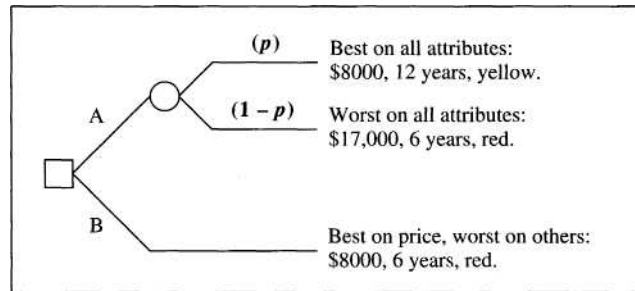


Figure 15.12
Assessing the weight for price.



$$\begin{aligned}
 U_p(\$8000) &= 1.00 & U_p(\$17,000) &= 0 \\
 U_L(12 \text{ Years}) &= 1.00 & U_L(6 \text{ Years}) &= 0 \\
 U_C(\text{Yellow}) &= 1.00 & U_C(\text{Red}) &= 0
 \end{aligned}$$

Substitute these values into the equation to obtain

$$k_p = 0.55(k_p + k_L + k_C)$$

Because $k_p + k_L + k_C = 1$ we have $k_p = 0.55$, which is simply the indifference probability that was assessed. Thus, we have a direct way to find the trade-off weight for the price attribute. Repeating this procedure one more time for the life-span attribute gives k_L , and then k_C follows because the weights must add to 1. Of course, a simple way to check for consistency is to repeat the procedure a third time for the color attribute. If assessed weights do not add to 1 and are not even close, then the additive model that we are using in this chapter is not appropriate. Chapter 16 discusses more complicated multiattribute utility models that may be able to accommodate such preferences.

Finally, you may have wondered why we always scale the individual utility functions from 0 to 1. The answer is contained in the equations above that underlie the weight-assessment techniques. In each case, we needed the values of 0 and 1 for best and worst cases in order to solve for the weights; the idea that scores are scaled from 0 to 1 is built into these assessment techniques. The result is that weights have very specific meanings. In particular, the swing-weight approach implies that we can interpret the weights in terms of improvements in utility that result from changing one

attribute from low to high, and those low and high values may be specific to the alternatives being considered. The lottery-assessment method suggests that the weights can be interpreted as an indifference probability in a comparison of lotteries. The specific low and high values among the alternatives are important anchors for interpreting these indifference probabilities.

Keeping Concepts Straight: Certainty Versus Uncertainty

Up to this point we have been doing something that is, strictly speaking, not correct. We have been mixing up decision making under conditions of certainty with decision making under conditions of uncertainty. For example, the decision regarding which automobile to buy, as we have framed it, is one that is made under certainty regarding the outcomes; the Norushi, Portalo, and Standard Motors cars have particular characteristics (price, color, and advertised life span) that you know for sure when you buy them. On the other hand, we could reframe the problem and consider life span to be uncertain. In this case, we would assess a probability distribution over possible life spans for each car. The decision model we would create would require us to assess a utility function that covers all of the possible-life span outcomes so that we could appropriately value the probability distributions or lotteries for life span that come with each car. As we also saw, the entrepreneur-bureaucrat example involves uncertainty, and we needed an appropriate utility function in order to evaluate those two options in a way that made sense.

Why does it matter whether we are talking about certainty or uncertainty? Some of the utility-assessment methods we have talked about are appropriate for decisions under uncertainty, and some are not. When making decisions under certainty, we can use what are called *value functions* (also known as *ordinal utility functions*). Value functions are only required to rank-order sure outcomes in a way that is consistent with the decision maker's preferences for those outcomes. There is no concern with lotteries or uncertain outcomes. In fact, different value functions that rank sure outcomes in the same way may rank a set of lotteries in different ways.

If you face a decision under uncertainty, you should use what is called a *cardinal utility function*. A cardinal utility function appropriately incorporates your risk attitude so that lotteries are rank-ordered in a way that is consistent with your risk attitude. All of the discussion in Chapters 13 and 14 concerned cardinal utility; we constructed preference models to incorporate risk attitudes.

The good news is that most of the assessment techniques we have examined are appropriate for conditions of both certainty and uncertainty. Of the methods for assessing individual utility functions, only the ratio approach is, strictly speaking, limited only to decisions under certainty; nothing about this assessment method encodes the decision maker's attitude about risk. The proportional-scores technique is a special case; it may be used under conditions of uncertainty by a decision maker who is risk-neutral for the specified attributes. Of course, all of the lottery-based utility-assessment methods from Chapters 13 and 14 are appropriate for decisions under

uncertainty. And all of the weight-assessment methods described in this chapter can be used regardless of the presence or absence of uncertainty; once a specific multiattribute preference model is established, the weight-assessment methods amount to various ways of establishing indifference among specific lotteries or consequences. The weights can then be derived on the basis of these indifference judgments.

How much does the distinction between ordinal and cardinal utility models really matter? Although the distinction can be important in theoretical models of economic behavior, practical decision-analysis applications rarely distinguish between the two. Human decision makers need help understanding objectives and trade-offs, and all of the assessment techniques described in this chapter are useful in this regard. As always, we adopt a modeling view. The objective of assessing a multiattribute utility function is to gain insight and understanding of the decision maker's trade-offs among multiple objectives. Any assessment method or modeling technique that helps to do this — which includes all that are mentioned in this chapter — should result in a reasonable preference model that can lead the decision maker to a clear choice among available alternatives.

In a way, the situation is similar to the discussion on decreasing and constant risk aversion in Chapter 13, where we argued that it is difficult enough to determine the extent of a given individual's risk aversion, let alone whether it is decreasing or not. When facing multiple objectives, any approach that will help us to understand trade-offs among objectives is welcome. Attention to preference ordering of lotteries rather than sure outcomes can be important in some situations but would come after obtaining a good basic understanding of the objectives and approximate trade-off weights.

A similar argument can be made for the additive utility function itself. Although strictly speaking it has some limitations and special requirements that will be explained in Chapter 16, it is an exceptionally useful and easy way to model preferences in many situations. Even if used only as an approximation, the additive utility function takes us a long way toward understanding our preferences and resolving a difficult decision.

An Example: Library Choices

With the discussion of assessing utility functions and weights behind us, we now turn to a realistic problem. This example will demonstrate the development of an additive utility function when there are many objectives. The issue is site selection for a new public library.

THE EUGENE PUBLIC LIBRARY

In 1986 a solution was sought for the overcrowded and inadequate conditions at the public library in Eugene, Oregon. The current building, with approximately 38,000

square feet of space, had been built in 1959 with the anticipation that it would serve satisfactorily for some 30 years. In the intervening years, Eugene's population had grown to the point that, on a per capita basis, the Eugene Public Library was one of the most heavily used public libraries in the western United States. All available space had been used. Even the basement, which had not been designed originally for patron use, had been converted to a periodicals reading room. Low-circulation materials had to be removed from the stacks and placed in storage to make room for patrons and books. Expansion was imperative; consultants estimated that 115,000 square feet of space were needed. The current building could be expanded, but because of the original design it could not be operated as efficiently as a brand new building. Other potential sites were available, but all had their own benefits and drawbacks. After much thought, the possibilities were reduced to four sites in or near downtown Eugene, one of which was the current site. Some were less expensive for one reason or another, others had better opportunities for future expansion, and so on. How could the city choose from among them?

In evaluating the four proposed library sites, the first task was to create a fundamental-objectives hierarchy. What aspects of site location were important? The committee in charge of the study created the hierarchy shown in Figure 15.13 with seven root-level fundamental objectives and a second level of detail objectives. Parking, for example, was broken down into patron parking, off-site parking, night staff parking, and bookmobile space. (Without knowledge of Eugene's financial system, certain attributes in the "Related Cost" category may not make sense.) An important objective that is conspicuous by its absence is minimizing the cost of building a library. The committee's strategy was to compare the sites on the seven fundamental criteria first, and then consider the price tags for each. We will see later how this comparison can be made.

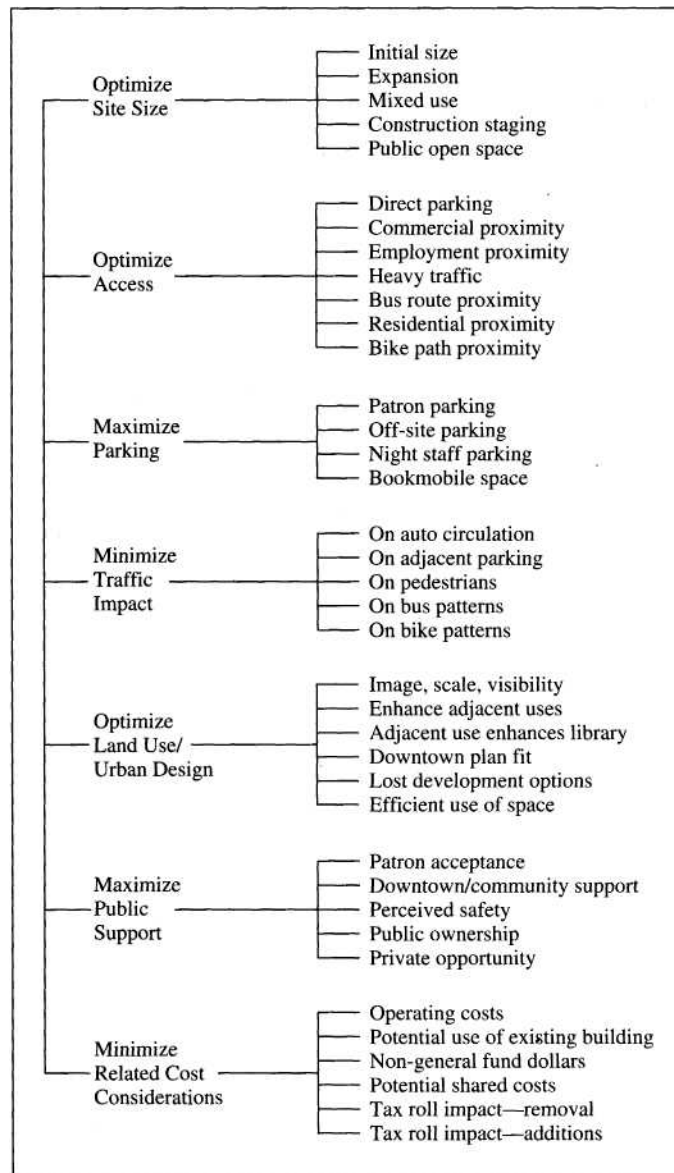
Comparing the four candidate sites (alternatives) required these four steps:

- 1 Evaluate the alternatives on each attribute.
- 2 Weight the attributes. This can be done with any of the three weight-assessment methods (although pricing out may be difficult because the overall cost is not included in this part of the analysis).
- 3 Calculate overall utility with the additive utility function.
- 4 Choose the alternative with the greatest overall utility.

Application of this procedure in the analysis of the four library sites produced the matrix shown in Table 15.6. The relative weights and the individual utilities are all shown in this table. (Actually, the committee's scores did not range from 0 to 1. Their scores have been thusly rescaled here, and the weights have been adjusted to be consistent with the new numbers. Six attributes on which all four sites were scored the same also have been eliminated.)

The overall utility is calculated and shown on a scale from 0 to 100, rather than from 0 to 1; this just makes the table easier to read. Under each fundamental objective is a subtotal for each site. For example, the subtotal for Site 1 under "Parking" is

Figure 15.13
Fundamental objectives for the Eugene Public Library site-evaluation study.



$$\begin{aligned} \text{Subtotal (Parking}_1) &= [(0.053 \times 0.20 \times 1.00) + (0.053 \times 0.60 \times 0.00) \\ &\quad + (0.053 \times 0.20 \times 1.00)] \times 100 \\ &= 2.12 \end{aligned}$$

(The factor of 100 at the end simply changes the scale so that it ranges from 0 to 100.) Once all individual utilities have been calculated for each fundamental objective, they are simply added to find the overall utility for each site. For example,

Table 15.6
Matrix of weights and utilities for four library sites.

Attributes	%	Utilities			
		Site 1	Site 2	Site 3	Site 4
Site Size (21.1%)					
Initial	38	1.00	0.00	1.00	1.00
Expansion (Horizontal)	13	0.00	0.00	0.00	1.00
Mixed Use	25	0.00	1.00	1.00	1.00
Construction Staging	12	1.00	0.00	0.00	1.00
Public Open Space	12	1.00	0.00	0.00	0.00
Subtotals		13.08	5.28	1329	18.57
Access (20.6%)					
Direct Parking	8	0.00	1.00	0.00	0.00
Commercial Proximity	23	0.00	1.00	0.67	1.00
Employment Proximity	15	0.50	1.00	0.00	1.00
Heavy Traffic	23	0.33	0.33	1.00	0.00
Bus Route Proximity	15	0.00	0.50	0.50	1.00
Residential Proximity	16	1.00	0.00	1.00	0.50
Subtotals		6.40	12.55	- 12.75	12.57
Parking (53%)					
Patron Parking	20	1.00	0.00	1.00	1.00
Off-Site Parking	60	0.00	1.00	0.33	0.33
Bookmobile Parking	20	1.00	0.00	1.00	1.00
Subtotals		2.12	3.18	3.17	3.17
Traffic Impacts (4.5%)					
Auto Circulation	47	0.00	0.75	1.00	0.00
Adjacent Parking	29	0.00	0.00	1.00	0.00
Bus Patterns	24	1.00	1.00	1.00	0.00
Subtotals		1.08	2.67	4.50	0.00
Land Use/Design (8.4%)					
Image/Scale/Visibility	13	0.00	1.00	0.00	0.00
Enhance Adjacent Uses	13	0.00	1.00	1.00	1.00
Adj. Uses Enhance Lib*	38	0.00	1.00	1.00	0.00
Downtown Plan Fit	13	1.00	0.00	1.00	1.00
Lost Devel. Options	23	1.00	0.00	0.00	0.00
Subtotals		3.02	5.38	5.38	2.18

Table 15.6
(Continued)

Attributes	%	Site 1	Site 2	Site 3	Site 4
Public Support (19.0%)					
Patron Acceptance	25	1.00	0.33	0.67	0.00
DT/Community Support	25	1.00	0.67	0.33	0.00
Perceived Safety	25	1.00	0.33	1.00	0.00
Public Ownership	17	0.00	1.00	1.00	0.00
Private Opportunity	8	1.00	0.00	1.00	1.00
<i>Subtotals</i>		15.77	9.55	14.25	1.52
Related Costs (21.1%)					
Operating Costs	20	0.00	1.00	1.00	1.00
Use of Existing Building	20	1.00	0.00	0.00	0.00
No General Fund \$	30	0.00	1.00	1.00	1.00
Tax Roll Impact, Removal	10	0.00	1.00	1.00	0.00
Tax Roll Impact, Added	20	0.00	1.00	1.00	1.00
<i>Subtotals</i>		4.22	16.88	16.88	14.77
<i>Weighted Score</i>		45.70	55.51	70.22	52.78

Source: Adapted from Robertson/Sherwood/Architects (1987) "Preliminary Draft Report: Eugene Public Library Selection Study. Executive Summary." Eugene, OR: Robertson/Sherwood.

$$\begin{aligned}
 U(\text{Site 1}) &= 13.08 + 6.40 + 2.12 + 1.08 + 3.02 + 15.77 + 4.22 \\
 &= 45.70
 \end{aligned}$$

You can see that the overall weight given to a specific attribute is the product of the specific weight at the lower level and the overall weight for the fundamental objective. For example, Site 1 has $U(\text{Construction Staging}) = 1.00$. This utility then is multiplied by 0.12, the weight for construction staging, and then multiplied by 0.211, the weight for the fundamental attribute of "Site Size." Thus, in this grand scheme, the utilities are weighted by a product of the weights at the two levels in the value tree. To express this more formally, let k_i represent the weight of the i th fundamental objective, and k_{ij} and U_{ij} the weight and utility, respectively, for the j th attribute under fundamental objective i . If there are m fundamental objectives and m_i detail attributes under fundamental objective i , then the overall utility for a site is

$$\begin{aligned}
 U(\text{Site}) &= k_1(k_{11}U_{11} + k_{12}U_{12} + k_{13}U_{13} + \dots + k_{1m_1}U_{1m_1}) \\
 &\quad + k_2(k_{21}U_{21} + k_{22}U_{22} + k_{23}U_{23} + \dots + k_{2m_2}U_{2m_2}) \\
 &\quad + \dots + k_m(k_{m1}U_{m1} + k_{m2}U_{m2} + k_{m3}U_{m3} + \dots + k_{mm_k}U_{mm_k}) \\
 &= k_1k_{11}U_{11} + k_1k_{12}U_{12} + \dots + k_mk_{mm_k}U_{mm_k} \\
 &= \sum_{i=1}^m k_i \left[\sum_{j=1}^{m_i} k_{ij}U_{ij} \right]
 \end{aligned}$$

From these expressions, we can see that the utilities on the individual attributes are being weighted by the product of the appropriate weights and then added. Thus, we still have an additive score that is a weighted combination of the individual utilities, just as we did in the simpler two- and three-attribute examples above. Moreover, it also should be clear that as the hierarchy grows to have more levels, the formula also grows, multiplying the individual utilities by all of the appropriate weights in the hierarchy.

The result of all of the calculations for the library example? Site 3 ranked the best with 70.22 points, Site 2 was second with 55.51 points, and Sites 4 and 1 (the current location) were ranked third and fourth overall with 52.78 and 45.70 points, respectively.

There is another interesting and intuitive way to interpret this kind of analysis. Imagine that 100 points are available to be awarded for each alternative, depending on how a given alternative ranks on each attribute. In the library case, 21.1 of the 100 points are awarded on the basis of "Site Size," 20.6 on the basis of "Access," 5.3 for "Parking," and so on. Within the "Site Size" category, the weights on the detail objectives determine how the 21.1 points for "Site Size" will be allocated; 38% of the 21.1 points (or 8.02 points) will be awarded on the basis of "Initial Size," 13% of the 21.1 points (2.74 points) will be awarded on the basis of "Expansion," and so on. We can see how this subdivision could continue through many layers in a hierarchy. Finally, when the ends of the branches are reached, we must determine utilities for the alternatives. If the utilities range from 0 to 1, then the utility indicates what proportion of the available points are awarded to the alternative for the particular detail attribute being considered. For example, Site 3 has a utility of 0.67 on "Commercial Proximity," so it receives 67% of the points available for this detail attribute. How many points are available? The weight of 23% tells us that 23% of the total points for "Access" are allocated to "Commercial Proximity," and the 20.6% weight for "Access" says that 20.6 points total are available for "Access." Thus, Site 3 earns $0.67 \times 20.6 = 13.802$ points for "Commercial Proximity." For each detail attribute, calculate the points awarded to Site 3. Now add those points; the total is Site 3's overall utility (70.22) on a scale from 0 to 100.


Recall that cost was an important attribute that was not included in the analysis of the library sites. The committee studying the problem decided to ignore construction costs until the sites were well understood in terms of their other attributes. But now that we have ranked the sites on the basis of the attributes, we must consider money. Table 15.7 shows the costs associated with each site, along with the overall utilities from Table 15.6, and Figure 15.14 shows the same information, graphically plotting cost against overall utility.

Table 15.7 and Figure 15.14 show clearly that Sites 1 and 4 are dominated. That is, if you like Site 4, then you should like Sites 2 or 3 even better, because each has a greater utility for less money. Likewise, Site 2 clearly dominates Site 1. Thus, Sites 1 and 4 can be eliminated from the analysis altogether on the basis of dominance. This leaves Sites 2 and 3. Is it worthwhile to pay the additional \$2.72 million to gain 14.71 additional points in terms of overall utility? Alternatively, is an increase of one point in the utility worth \$184,908?

Table 15.7
Cost and overall utility
for four library sites.

Site	Cost (\$Million)	Overall Utility
1	21.74	45.70
2	18.76	55.51
3	21.48	70.22
4	24.80	52.78

Figure 15.14
Library site costs plotted against overall utility.



Obviously, we are trying to price out the value of a single point on our 1-to-100 scale, just as we previously priced out the value of changes in attributes. But answering this question now is difficult because one point in the overall utility may have many components. One possible approach is to return to the detail attributes, look for specific attributes on which Site 3 is ranked higher than Site 2, and consider how much we would be willing to pay (in dollars) to bring Site 2 up to Site 3's level on this attribute. For example, Site 3 scored much higher than Site 2 for access during heavy traffic periods. The difference is that Site 2 is in a relatively congested area of downtown and on one of the city's main thoroughfares. In contrast, Site 3 is located at the edge of downtown, and access would be through relatively low-volume streets. How much would altering the traffic patterns be worth so that Site 2 would be just as good as Site 3 for this attribute? One million dollars? More?

Let us suppose that we have assessed that the difference between the sites in terms of access during heavy traffic is indeed worth \$1 million. It turns out that Site 3's advantage increases its weighted score by 3.17 points. (This is the difference between the sites in points awarded for this attribute.) Thus, the assessment would indicate that 3.17 points of overall utility are worth \$1 million, or \$315,457 per point of utility. This is greater than the \$184,908 per point required to switch from Site 2 to Site 3.

A pricing-out approach such as this can be used to assess the dollar value of one point of overall utility. Rather than making only one assessment, however, we should make several on different attributes where the two alternatives differ. For some of these assessments it may be possible to make estimates of true costs in terms of the market price of adjacent land, redesigning traffic patterns, constructing parking lots, and so on, and these cost estimates then might be helpful in assessing the dollar value of specific differences between sites. If the assessments in terms of prices per point of overall utility come out fairly close, then the average of these assessments can be used as a reasonable measure of the value of one point of overall utility. In the case of the library sites, if the final price per point is less than \$184,908, then Site 2 is preferred; otherwise Site 3 is preferred. If the price per point turns out to be very close to \$184,908, then the two sites are approximately equivalent.

The analysis probably would have been more complete had the committee elected to include minimizing cost as one of its fundamental objectives. In fact, all of the other attributes could have been priced out in terms of dollars, thus making the comparisons and trade-offs more intuitive. But it is not unusual for groups to ignore costs in an initial stage. The motivation often is political; it may be easier to gain constituent support by playing up the benefits of a project such as a new library early before talking about the "bottom line."

Using Software for Multiple-Objective Decisions

As you can imagine, multiple-objective decisions can become very complex. The library example involved only four alternatives and a relatively simple objectives hierarchy. Still, there were many assessments to make, and the calculations involved, though straightforward, are tedious. Given the ubiquitous nature of multiple-objective decisions, it should come as no surprise that many different multiple-objective decision programs are available. All of them provide some combination of different assessment methods and ways to analyze the alternatives. For example, in Chapter 3 we took a quick look at Logical Decisions. This program allows one to judge trade-offs graphically, with the swing-weight approach, and still other methods we have not mentioned. Alternatives can be compared with each other, and sensitivity analysis on the weights can be performed. Two other available programs are HIVIEW and HIPRE 3+. Both of these are similar to Logical Decisions in creating a modeling environment that is especially tuned to use of the additive utility function.

SUMMARY

This chapter has introduced the basics of making decisions that involve trade-offs. The basic problem is to create a model of the decision maker's preferences in terms of the various objectives. The first step in doing so, as discussed in detail in Chapter 3, is to understand which objectives are most important in making your decision; this requires

introspection as to the fundamental objectives. The goal of this step is to create a fundamental-objectives hierarchy and a set of operational attribute scales to measure the performance of each alternative or outcome on each of the fundamental objectives.

To evaluate the alternatives, we introduced the additive utility function, which calculates an overall utility for an alternative or outcome as a weighted sum of individual utility functions for each fundamental objective. Assessing the individual utility functions is best done through a consistent method of comparison that results in utilities ranging from 0 to 1. We discussed three methods: calculation of proportional scores, assessment of ratios, and the conventional lottery-based assessment discussed in Chapters 13 and 14. Once these utility functions are established, weights must also be assessed. Several assessment procedures are possible here, all providing mechanisms for determining the rate at which the attributes can be traded off against one another. Sometimes trade-off rates can be assessed in terms of dollars by pricing out the other attributes. The swing-weighting and lottery-based assessment techniques can be used even if it is difficult to think of the attributes in dollar terms, and these techniques lead to clear interpretations of the assessed weights. With weights and individual utility functions determined, overall utilities are calculated by applying the additive utility formula. We demonstrated the procedure with an example in which potential sites for a public library were evaluated.

EXERCISES

- 15.1 Why is it important to think carefully about decision situations that involve multiple conflicting objectives?
- 15.2 Explain the general decision-analysis approach to dealing with multiple objectives.
- 15.3 Explain what is meant by the term *indifference curve*.
- 15.4 Explain the idea of dominance in the context of multiobjective decision making.
- 15.5 Imagine that you are working on a committee to develop an employment-conditions policy for your firm. During the discussion of safety, one committee member states that employee safety is twice as important as benefits such as flexible work time, employer-sponsored day care, and so on.
 - a How might you respond to such a statement?
 - b Suppose the statement was that employee safety was twice as important as insurance benefits. How might you translate this statement into a form (perhaps based on dollar values) so that it would be possible to evaluate various policy alternatives in terms of safety and insurance in a way that would be consistent with the statement?
- 15.6 Explain why proportional scores represent a risk attitude that is risk-neutral.
- 15.7 An MBA student evaluating weather outcomes for an upcoming party has concluded that a sunny day would be twice as good as a cloudy day, and a cloudy day would be three times as good as a rainy day. Use these assessments to calculate utilities that range from 0 to 1 for sunny, rainy, and cloudy days.

- 15.8** Explain in your own words why swing weights produce meaningful weights for evaluating alternatives.
- 15.9** A decision maker is assessing weights for two attributes using the swing-weight method. When he imagines swinging the attributes individually from worst to best, he concludes that his improvement in satisfaction from Attribute A is 70% of the improvement from swinging Attribute B. Calculate k_A and k_B .
- 15.10** Explain in your own words why the lottery method works to produce meaningful weights for evaluating alternatives.
- 15.11** A decision maker is assessing weights for three attributes using the lottery-assessment method. In considering the lotteries, she concludes that she is indifferent between:
- A** Win the best possible combination with probability 0.34
Win the worst possible combination with probability 0.66
- and
- B** Win a combination that is worst on Attributes 1 and 3 and best on 2
- She also has concluded that she is indifferent between
- C** Win the best possible combination with probability 0.25
Win the worst possible combination with probability 0.75
- and
- D** Win a combination that is worst on Attributes 2 and 3 and best on 1
- Find weights k_1 , k_2 , and k_3 .

QUESTIONS AND PROBLEMS

- 15.12** Suppose that you are searching for an apartment in which to live while you go to school. Apartments near campus generally cost more than equivalent apartments farther away. Five apartments are available. One is right next to campus, and another is one mile away. The remaining apartments are two, three, and four miles away.
- a** Suppose you have a tentative agreement to rent the apartment that is one mile from campus. How much more would you be willing to pay in monthly rent to obtain the one next to campus? (Answer this question on the basis of your own personal experience. Other than rent and distance from campus, the two apartments are equivalent.)
- b** Now suppose you have a tentative agreement to rent the apartment that is four miles away. How much more would you be willing to pay in monthly rent to move to the apartment that is only three miles from campus?
- c** What are the implications of your answers to parts a and b? Would it be appropriate to rank the apartments in terms of distance using the proportional-scoring technique?

- d** Sketch an indifference curve that reflects the way you would trade off rent versus proximity to campus. Is your indifference curve a straight line?
- 15.13** A friend of yours is in the market for a new computer. Four different machines are under consideration. The four computers are essentially the same, but they vary in price and reliability. The least expensive model is also the least reliable, the most expensive is the most reliable, and the other two are in between.
- a** Describe to your friend how you would approach the decision.
- b** Define reliability in a way that would be appropriate for the decision. Do you need to consider risk?
- c** How might your friend go about establishing a marginal rate of substitution between reliability and price?
- 15.14** Continuing Problem 15.13, the computers are described as follows:
- A** Price: \$998.95 Expected number of days in the shop per year: 4
- B** Price: \$1300.00 Expected number of days in the shop per year: 2
- C** Price: \$1350.00 Expected number of days in the shop per year: 2.5
- D** Price: \$1750.00 Expected number of days in the shop per year: 0.5
- The computer will be an important part of your friend's livelihood for the next two years. (After two years, the computer will have a negligible salvage value.) In fact, your friend can foresee that there will be specific losses if the computer is in the shop for repairs. The magnitude of the losses are uncertain but are estimated to be approximately \$180 per day that the computer is down.
- a** Can you give your friend any advice without doing any calculations?
- b** Use the information given to determine weights k_P and k_R , where R stands for reliability. What assumptions are you making?
- c** Calculate overall utilities for the computers. What do you conclude?
- d** Sketch three indifference curves that reflect your friend's trade-off rate between reliability and price.
- e** What considerations other than losses might be important in determining the trade-off rate between cost and reliability?
- 15.15** Throughout the chapter we have assessed individual utility functions that range from 0 to 1. What is the advantage of doing this?
- 15.16** You are an up-and-coming developer in downtown Seattle and are interested in constructing a building on a site that you own. You have collected four bids from prospective contractors. The bids include both a cost (millions of dollars) and a time to completion (months):

Contractor	Cost	Time
A	100	20
B	80	25
C	79	28
D	82	26

The problem now is to decide which contractor to choose. B has indicated that for another \$20 million he could do the job in 18 months, and you have said that you would be indifferent between that and the original proposal. In talking with C, you have indicated that you would just as soon pay her an extra \$4 million if she could get the job done in 26 months. Who gets the job? Explain your reasoning. (It may be convenient to plot the four alternatives on a graph.)

- 15.17** Once you decide that you are in the market for a personal computer, you have many different considerations. You should think about how you will use the computer, and so you need to know whether appropriate software is available and at what price. The nature of the available peripheral equipment (printers, disk drives, and so on) can be important. The "feel" of the computer, which is in some sense determined by the operating system and the user interface, can be critical. Are you an experienced user? Do you want to be able to program the machine, or will you (like most of us) rely on existing or over-the-counter software? If you intend to use the machine for a lot of number crunching, processor speed may be important. Reliability and service are other matters. For many students, an important question is whether the computer will be compatible with other systems in any job they might eventually have. Finally, of course, price and operating costs are important.

Create a fundamental-objectives hierarchy to compare your options. Take care in doing this; be sure that you establish the fundamental objectives and operational attributes that will allow you to make the necessary comparisons. (Note that the attributes suggested above are *not* exhaustive, and some may not apply to you!)

Use your model to evaluate at least three different computers (preferably from different manufacturers). You will have to specify precisely the packages that you compare. It also might be worthwhile to include appropriate software and peripheral equipment. (Exactly what you compare is up to you, but make the packages meaningful.) Be sure that your utilities are such that the best alternative gets a 1 and the worst a 0 for each attribute. Assess weights using pricing out, swing weighting, or lottery weights. Calculate overall utilities for your alternatives.

Try using the utility functions for money and computer reliability that you assessed in Problems 14.12 and 14.13. You may have to rescale the utility functions to obtain scores so that the best alternative in this problem scores 1, and the worst scores 0.

If possible, use a computer-based multiattribute decision program to do this problem.

- 15.18** When you choose a place to live, what objectives are you trying to accomplish? What makes some apartments better than others? Would you rather live close to campus or farther away and spend less money? What about the quality of the neighborhood? How about amenities such as a swimming pool?

Create a fundamental-objectives hierarchy that allows you to compare apartment options. Take care in doing this; be sure to establish the fundamental objectives and operational attributes that will allow you to make the necessary comparisons.

Once you are satisfied with your hierarchy, use it to compare available housing alternatives. Try ranking different apartments that are advertised in the classified section of the newspaper, for example. Be sure that your individual utilities follow the rules: Best takes a 1 and worst takes 0. (Try using the utility function for money that you assessed in Problem 14.12. You may have to rescale it so that your best alternative gets a 1 and worst

gets a 0.) Assess weights using pricing out, swing weighting, or lottery weights. Evaluate your alternatives with the additive utility function.

If possible, use a computer-based multiattribute decision program to do this problem.

- 15.19** What is important to you in choosing a job? Certainly salary is important, and for many people location matters a lot. Other considerations might involve promotion potential, the nature of the work, the organization itself, benefits, and so on.

Create a fundamental-objectives hierarchy that allows you to compare job offers. Be sure to establish the fundamental objectives and operational attributes that will allow you to make the necessary comparisons.

Once you are satisfied with your hierarchy, use it to compare your job offers. You also may want to think about your "ideal" job in order to compare your current offers with your ideal. You also might consider your imaginary worst possible job in all respects. (For the salary attribute, try using the utility function for money that you assessed in Problem 14.12, or some variation of it. You may have to rescale it so that your best alternative is 1 and worst 0.) Assess weights for the attributes using pricing out, swing weighting, or lottery weights, being careful to anchor your judgments in terms of both ideal and worst imaginable jobs. Evaluate your various job offers with the additive utility function.

If possible, use a computer-based multiattribute decision program to do this problem.

- 15.20** How can you compare your courses? When you consider those that you have taken, it should be clear that some were better than others and that the good ones were, perhaps, good for different reasons. What are the important dimensions that affect the quality of a course? Some are obvious, such as the enthusiasm of an instructor, topic, and amount and type of work involved. Other aspects may not be quite so obvious; for example, how you perceive one course may depend on other courses you have had.

In this problem, the objective is to create a "template" that will permit consistent evaluation of your courses. The procedure is essentially the same as it is for any multiattribute decision, except that you will be able to use the template to evaluate future courses. Thus, we do not have a set of alternatives available to use for the determination of scores and weights. (You want to think, however, about current and recent courses in making your assessments.)

First, create a fundamental-objectives hierarchy that allows you to compare courses. Be sure to establish the fundamental objectives and operational attributes that will allow for the necessary comparisons. Constructing a set of objectives for comparing courses is considerably more difficult than comparing computers, apartments, or jobs. You may find that many of the attributes you consider initially will overlap with others, leading to a confusing array of attributes that are interdependent. It may take considerable thought to reduce the degree of redundancy in your hierarchy and to arrive at one that is complete, decomposable, small enough to be manageable, and involves attributes that are easy to think about.

Once you are satisfied with your objectives and attributes, imagine the best and worst courses for each attribute. Create attribute scales (constructed scales where appropriate) for each objective. The idea is to be able to return to these scales with any new course and determine utilities for each objective with relative ease. (Try using the homework utility function that you developed in Problem 14.15. You may have to rescale it so that your best alternative gets a 1 and the worst a 0.)

Once you have created the attribute scales, you are ready to assess the weights. Try the swing-weighting or lottery approach for assessing the weights. (Pricing out may be difficult to do in this particular example. Can you place a dollar value on your attributes?)

Finally, with scales and weights established, you are ready to evaluate courses. Try comparing three or four of your most recent courses. (Try evaluating one that you took more than a year ago. Can you remember enough about the course to assess the individual utilities with some degree of confidence?)

If possible, implement your course-evaluation template using a computer-based multiattribute decision program. Alternatively, you might create a spreadsheet template that you could use to evaluate courses.

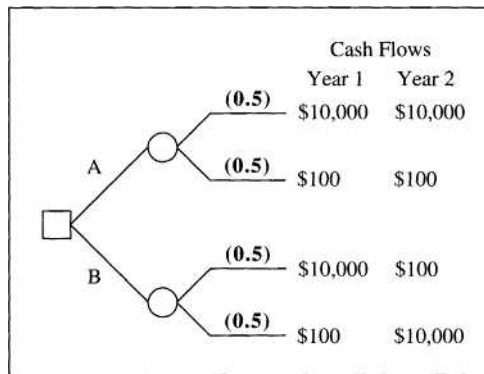
- 15.21** Refer to the discussion of the automobiles in the section on "Trading Off Conflicting Objectives: The Basics." We discussed switching first from the Standard to the Norushi, and then from the Norushi to the Portalo. Would it make sense to consider a direct switch from the Standard to the Portalo? Why or why not?
- 15.22** In Chapter 2 we discussed net present value (NPV) as a procedure for evaluating consequences that yield cash flows at different points in time. If x_i is the cash flow at year i and r is the discount rate, then the NPV is given by

$$\text{NPV} = \sum \frac{x_i}{(1+r)^i}$$

where the summation is over all future cash flows including the current x_0 .

- a** Explain how the NPV criterion is similar to the additive utility function that was discussed in this chapter. What are the attributes? What are the weights? Describe the way cash at time period i is traded off against cash at time period $i + 1$. (*Hint*: Review Chapter 2!)
- b** Suppose that you can invest in one of two different projects. Each costs \$20,000. The first project is riskless and will pay you \$10,000 each year for the next three years. The second one is risky. There is a 50% chance that it will pay \$15,000 each year for the next three years and a 50% chance that it will pay only \$5000 per year for the next three years. Your discount rate is 9%. Calculate the NPV for both the riskless and risky project. Compare them. What can you conclude about the use of NPV for deciding among risky projects?
- c** How might your NPV analysis in part b be modified to take risk into account? Could you use a utility function? How does the idea of a risk-adjusted discount rate fit into the picture? How could the interest rate be adjusted to account for risk? Would this be the same as using a utility function for money?
- 15.23** Following up Problem 15.22, even though we take riskiness into account, there still is difficulty with NPV as a decision criterion. Suppose that you are facing the two risky projects shown in Figure 15.15.
- Project A pays either \$10,000 for each of two years or \$100 for those two years. Project B pays \$10,000 either in the first year or the second year and \$100 in the other year. Assume that the cash flows are annual new profits.
- a** Which of these two risky investments would you prefer? Why?
- b** Calculate the expected NPV for both projects, using the same 9% interest rate from Problem 15.22. Based on expected NPV, in which project would you invest?

Figure 15.15
 Decision tree for
 problem 15.23.
 Which of these two
 risky investments
 would you prefer?



c After careful assessment, you have concluded that you are risk-averse and that your utility function can be adequately represented by $U(X_i) = \ln(X_i)$, where X_i represents cash flow during year i . Calculate the expected net present utility for each project. Net present utility is given by

$$NPU = \sum \frac{U(X_i)}{(1+r)^i}$$

d NPU in part c should incorporate your attitude toward risk. Are your NPU calculations in part c consistent with your preferences in part a? What is there about these two projects that is not captured by your utility function? Can you think of any other way to model your preferences?

15.24 Instead of calculating a "discounted" utility as we did in Problem 15.23, let us consider calculating $U(\text{NPV})$. That is, calculate NPV first, using an appropriate interest rate, and then calculate a utility value for the NPV. For your utility function, use the exponential utility function $U(\text{NPV}) = 1 - e^{-\text{NPV}/5000}$. Use this approach to calculate the expected utility of Projects A and B in Figure 15.15. Which would you choose? Are there any problems with using this procedure for evaluating projects?


15.25 A policy maker in the Occupational Safety and Health Administration is under pressure from industry to permit the use of certain chemicals in a newly developed industrial process. Two different versions of the process use two different chemicals, A and B. The risks associated with these chemicals are not known with certainty, but the available information indicates that they may affect two groups of people in the following ways:

- *Chemical A* There is a 50% chance that Group 1 will be adversely affected, while Group 2 is unaffected; and a 50% chance that Group 2 is adversely affected, while Group 1 is unaffected.
- *Chemical B* There is a 50% chance that both groups will be adversely affected, and a 50% chance that neither group will be affected.

Assume that "adversely affected" means the same in every case — an expected increase of one death in the affected group over the next two years. The decision maker's problem looks like the decision tree in Figure 15.16.

a Calculate the expected number of deaths for each chemical.

Figure 15.16
Deciding between
alternative chemicals
in problem 15.25.



- b** A decision maker who values consequences using an overall utility might calculate the utility for each consequence as

$$U(\text{Chemical}) = k_1 U_1(\text{Group 1 Deaths}) + k_2 U_2(\text{Group 2 Deaths})$$

For both U_1 and U_2 , the best and worst possible outcomes are 0 deaths and 1 death, respectively. Thus,

$$\begin{aligned} U_1(1 \text{ Death}) &= 0 & U_1(0 \text{ Deaths}) &= 1 \\ U_2(1 \text{ Death}) &= 0 & U_2(0 \text{ Deaths}) &= 1 \end{aligned}$$

Explain why k_1 and $1 - k_1$ may not be equal.

- c** Assume that $k_1 = 0.4$. Show that the decision maker who evaluates the two chemicals in terms of their expected overall utilities (as defined above) would be indifferent between them. Does the value of k_1 matter?
- d** Why might the decision maker *not* be indifferent between the two programs? (Most people think about the decision maker's risk attitude toward the number of deaths or lives saved. Besides this, think about the following: Suppose you are a member of Group 1, and the decision maker has chosen Chemical A. It turned out that Group 1 was affected. How would you feel? What would you do? What does this imply for the decision maker?)
- 15.26** Refer to the discussion of the three automobiles in the section "Trading Off Conflicting Objectives: The Basics." Suppose we had the following individual utility functions for price and life span:

Life span	Price
$U_L(6 \text{ years}) = 0.00$	$U_P(17,000) = 0.00$
$U_L(9 \text{ Years}) = 0.75$	$U_P(10,000) = 0.50$
$U_L(12 \text{ Years}) = 1.00$	$U_P(8000) = 1.00$

The additive utility model discussed in this chapter would give us the following:

$$U(\text{Price, Life Span}) = k_L U_L(\text{Life Span}) + (1 - k_L) U_P(\text{Price})$$

- a** With $k_l = 0.45$, calculate the utility for the three cars. Which would be chosen?
- b** Suppose that you are not completely comfortable with the assessment of $k_l = 0.45$. How *large* could k_l be before the decision changes, and what would be the new choice? How *small* could k_l be before the decision changes, and what would be the new choice? Would you say that the choice among these three cars is very sensitive to the assessment of k_l ?
- 15.27** Refer to Table 15.6. Your boss is a member of the library-site selection committee. She is not perfectly satisfied with the assessments shown in the table. Specifically, she wonders to what extent the assessed weights and individual utilities on the detail attributes could change without affecting the overall ranking of the four sites. Use an electronic spread sheet to answer this question, and write a memo that discusses your findings. (*Hint*: There is no specific way to attack a sensitivity analysis like this. One possibility is to establish reasonable ranges for the weights and then create a tornado diagram. Be sure that your weights add to 1 in each category!)
- 15.28** Refer to Problem 15.11. Suppose that the decision maker has made a third assessment, concluding that she is indifferent between
- E** Win the best possible combination with probability 0.18
Win the worst possible combination with probability 0.82
- and
- F** Win a combination that is worst on Attributes 1 and 2, and best on 3

What does this assessment imply for the analysis? Is it consistent with your answer for k_3 in Problem 15.11? What should you do now?

C A S E S T U D I E S

THE SATANIC VERSES

In early 1989, the Ayatollah Khomeini of Iran decreed that Salman Rushdie, the British author of *The Satanic Verses*, should be put to death. In many ways, Rushdie's novel satirized Islam and the Prophet Muhammed. Khomeini declared that Rushdie should die and that whoever killed him would go to Heaven.

Many bookstores in both Europe and the United States that carried *The Satanic Verses* found themselves in a bind. Some Muslims threatened violence unless the bookstores stopped selling the book. Some booksellers removed the book from their shelves and sold it only to customers who specifically asked for it. Others refused to sell it altogether on the grounds that it was too risky. Still others defied the threats.

One bookseller in Berkeley, California, continued selling the book on the grounds that he would not allow anyone to interfere with the principle of freedom of the press. His store was bombed, and damage was substantial. His reaction? He increased security.

Questions

- 1 Imagine that you are the owner of a bookstore faced with the decision of what to do about *The Satanic Verses*. In deciding on a course of action, there are several conflicting objectives. Develop an objectives hierarchy and operational attribute scales.
- 2 What alternatives do you have?
- 3 What risks do you face? Do the risks differ depending on the alternative you choose? Sketch a simple influence diagram or decision tree for your problem.

DILEMMAS IN MEDICINE

In *Alpha and Omega: Ethics at the Frontiers of Life and Death*, author Ernie Young specifies four fundamental principles that must be considered in making medical decisions: beneficence, nonmaleficence, justice, and autonomy. The following descriptions of these principles have been abstracted from the book (pp. 21-23):

- *Beneficence* implies that the physician's most important duty is to provide services that are beneficial to the patient. In many cases, this can mean taking measures that are intended to preserve the patient's life.
- *Nonmaleficence* is the duty not to cause harm to the patient. A medical aphorism of uncertain origin proclaims *primum non nocere* — above all, do no harm. Harm can mean different things in different situations and for different patients, and can include death, disability, separation from loved ones, or deprivation of pleasure or freedom. The difficulty is that many medical procedures entail at least some harm or the potential for harm. This always must be weighed against the potential benefits.
- *Justice* in this context refers to the fair use of resources. It is, after all, impossible to do absolutely everything that would be medically justifiable for all patients. Thus, decisions must be made regarding the allocation of scarce resources. The issue is how to make these decisions fairly or equitably. For example, how should we decide what patients have priority for receiving donated organs? Is it appropriate to admit a terminally ill patient to an intensive care unit?
- *Autonomy* requires allowing a patient to make his or her own decisions regarding medical treatment as far as is possible. The patient, operating as an independent, self-determining agent, should be able to obtain appropriate information and par-

ticipate fully in the decisions regarding the course of treatment and, ultimately, the patient's life.

In most medical situations, these principles do not conflict. That is, the physician can provide beneficial care for the patient without causing harm, the treatment can be provided equitably, and the patient can easily make his or her own decisions. In a few cases, however, the principles are in conflict and it is impossible to accomplish all of them at once. For example, consider the case of a terminally ill patient who insists that everything possible be done to extend his or her life. Doing so may violate both nonmaleficence and justice while at the same time providing limited benefit. But not providing the requested services violates autonomy. Thus, the physician would be in a very difficult dilemma.

Questions

- 1 Discuss the relationship between the medical ethics here and decision making in the face of conflicting objectives. Sketch an objectives hierarchy for a physician who must cope with difficult problems such as those described above. Can you explain or expand on the four fundamental objectives by developing lower-level objectives?
- 2 Neonatology is the study and treatment of newborn infants. Of particular concern is the treatment of low-birth-weight infants who are born prematurely. Often these babies are the victims of poor prenatal care and may be burdened with severe deformities. Millions of dollars are spent annually to save the lives of such infants. Discuss the ways in which the four principles conflict in this situation. Could you give any guidelines to a panel of doctors and hospital administrators grappling with such problems?
- 3 Terminally ill patients face the prospect of death within a relatively short period of time. In the case of cancer victims, their last months can be extremely painful. Increasing numbers of such patients consider taking their own lives, and much controversy has developed concerning euthanasia, or mercy killing. Imagine that a patient is terminally ill and mentions that he or she is considering suicide. If every thing reasonable has been done to arrest the disease without success, this may be a reasonable option. Furthermore, the principle of autonomy should be respected here, provided that the patient is mentally stable and sound and understands fully the implications. But how deeply should the physician or loved one be involved? There are varying degrees of involvement. First, the physician might simply provide counseling and emotional support. The next step would be encouraging the patient by removing obstacles. Third, the physician might provide information about how to end one's life effectively and without trauma. The next step would be to assist in the procurement of the means to commit suicide. Helping the patient to end his or her life represents still another step, and actually killing the patient — by lethal injection or removal of a life support system, for example — would represent full involvement.

Suppose that one of your loved ones were terminally ill and considering suicide. What issues would you want him or her to consider carefully? Draw an objectives hierarchy for the patient's decision.

Now suppose that the patient has asked you to assist in his or her suicide. What issues would you want to consider when deciding on your level of involvement? Sketch an objectives hierarchy for your own decision. Compare this hierarchy with the patient's.

Source: E. Young (1989) *Alpha and Omega: Ethics at the Frontiers of Life and Death*, Palo Alto, CA: Stanford Alumni Association.

A MATTER OF ETHICS

Paul Lambert was in a difficult situation. When he started his current job five years ago, he understood clearly that he would be working on sensitive defense contracts for the government. In fact, his firm was a subcontractor for some major defense contractors. What he did not realize at the time — indeed, he only discovered this gradually over the past two years — was that the firm was overcharging. And it was not just a matter of a few dollars. In some cases, the government was overcharged by as much as a factor of 10.

Three weeks ago, he inadvertently came across an internal accounting memo that documented one particularly flagrant violation. He quietly made a copy and locked it in his desk. At the time, he was amazed, then righteously indignant. He resolved to take the evidence immediately to the appropriate authorities. But the more he thought about it, the more confused he became. Finally, he called his brother-in-law, Jim Grillich. Jim worked for another defense-related firm and agreed to have lunch with Paul. After exchanging stories about their families and comments on recent sporting events, Paul laid his cards on the table.

"Looks as though you could really make some waves," Jim commented, after listening to Paul's story.

"I guess I could. But I just don't know. If I blow the whistle, I'd feel like I'd have to resign. And then it would be tough to find another job. Nancy and I don't have a lot of savings, you know." The thought of dipping into their savings made Paul shake his head. "I just don't know."

The two men were silent for a long time. Then Paul continued, "To make matters worse, I really believe that the work that the company is doing, especially in the research labs, is important. It may have a substantial impact on our society over the next 20 years. The CEO is behind the research 100 percent, and I gather from the few comments I've overheard that he's essentially funding the research by overcharging on the subcontracts. So if I call foul, the research program goes down the drain."

"I know what you mean." Jim went on to recount a similar dilemma that he faced a few years before.

"So what did you do?"

"The papers are still in my desk. I always wanted to talk to someone about it. I even thought about calling you up, but I never did. After a while, it seemed like it was pretty easy just to leave the papers in there, locked up, safe and sound."

Questions

- 1 What trade-offs is Paul trying to make? What appear to be his fundamental objectives?
- 2 Suppose that Paul's take-home pay is currently \$2400 per month. In talking to an employment company,, he is told that it will probably take two months to find a similar job if he leaves his current one, and he had better expect three months if he wants a better job. In looking at his savings account of \$10,500, he decides that he cannot justify leaving his job, even though this means keeping quiet about the overcharging incident. Can you say anything about an implicit trade-off rate between the fundamental objectives that you identified above?
- 3 Have you ever been in a situation in which it was difficult for you to decide whether to take an ethically appropriate action? Describe the situation. What made the decision difficult? What trade-offs did you have to make? What did you finally do?

FDA AND THE TESTING OF EXPERIMENTAL DRUGS

The Food and Drug Administration (FDA) of the federal government is one of the largest consumer-protection agencies in the world. One of the FDA's charges is to ensure that drugs sold to consumers do not pose health threats. As a result, the testing procedure that leads to a new drug's approval is rigorous and demanding. So much so, in fact, that some policy makers are calling for less stringent standards. Below are some of the dilemmas that FDA faces:

- If an experimental drug shows promise in the treatment of a dangerous disease such as AIDS, should the testing procedure be abbreviated in order to get the drug to market more quickly?
- The FDA already is a large and costly bureaucracy. By easing testing standards, substantial dollars could be saved. But would it be more likely that a dangerous drug would be approved? What are the costs of such a mistake?

A fundamental trade-off is involved here. What are we gaining in the way of assurance of safe drugs, and what are we giving up by keeping the drugs away from the general public for an additional year or two?

Questions

- 1 What are the consequences (both good and bad) of keeping a drug from consumers for some required period of rigorous testing?
- 2 What are the consequences (both good and bad) of allowing drugs to reach consumers with less stringent testing?
- 3 Imagine that you are the FDA's commissioner. A pharmaceutical company requests special permission to rush a new AIDS drug to market. On the basis of a first round

of tests, the company estimates that the new drug will save the lives of 200 AIDS victims in the first year. Your favorite pharmacologist expresses reservations, however, claiming that without running the complete series of tests, he fears that the drug may have as-yet-undetermined but serious side effects. What decision would you make? Why?

- 4 Suppose that the drug in Question 3 was for arthritis. It could be used by any individual who suffers from arthritis and, according to the preliminary tests, would be able to cure up to 80% of rheumatoid arthritis cases. But your pharmacologist expresses the same reservations as for the AIDS drug. Now what decision would you make? Why?

REFERENCES

The additive utility function has been described by many authors. The most comprehensive discussion, and the only one that covers swing weights, is that by von Winterfeldt and Edwards (1986). Keeney and Raiffa (1976) and Keeney (1980) also devote a lot of material to this preference model. Edwards and Barron (1994) discuss some heuristic approaches to assessing weights, including the use of only rank-order information about the objectives.

The basic idea of creating an additive utility function is fairly common and has been applied in a variety of settings. Moreover, this basic approach also has earned several different names. For example, a cost-benefit analysis typically prices out nonmonetary costs and benefits and then aggregates them. For an interesting critique of a cost-benefit analysis, see Bunn (1984, Chapter 5).

Other decision-aiding techniques also use the additive utility function implicitly or explicitly, including the Analytic Hierarchy Process (Saaty, 1980) and goal programming with nonpreemptive weights (see Winston 1987). Conjoint analysis, a statistical technique used in market research to determine preference patterns of consumers on the basis of survey data, often is used to create additive utility functions. In all of these, some kind of subjective judgment forms the basis for the weights, and yet the interpretation of the weights is not always clear. For all of these alternative models, extreme care must be exercised in making the judgments on which the additive utility function is based. There is no substitute for thinking hard about trade-off issues. This text's view is that the decision-analysis approach discussed in this chapter provides the best systematic framework for making those judgments.

Bunn, D. (1984) *Applied Decision Analysis*. New York: McGraw-Hill.

Edwards, W., and F. H. Barron (1994) "SMARTS and SMARTER: Improved Simple Methods for Multiattribute Utility Measurement." *Organizational Behavior and Human Decision Processes*, 60, 306-325.

Keeney, R. (1980) *Siting Energy Facilities*. New York: Academic Press.

Keeney, R., and H. Raiffa (1976) *Decisions with Multiple Objectives*. New York: Wiley.

Saaty, T. (1980) *The Analytic Hierarchy Process*. New York: McGraw-Hill.

von Winterfeldt, D., and W. Edwards (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press.

Winston, W. (1987) *Operations Research: Applications and Algorithms*. Boston: PWS-KENT.

EPILOGUE

What happened with the Eugene Public Library? After much public discussion, an alternative emerged that had not been anticipated. An out-of-state developer expressed a desire to build a multistory office building in downtown Eugene (at Site 2, in fact) and proposed that the library could occupy the lower two floors of the building. By entering into a partnership with the developer, the city could save a lot in construction costs. Many citizens voiced concerns about the prospect of the city's alliance with a private developer, others were concerned about the complicated financing arrangement, and still others disapproved of the proposed location for a variety of reasons. On the other hand, the supporters pointed out that this might be the only way that Eugene would ever get a new library. In March 1989, the proposal to accept the developer's offer was submitted to the voters. The result? They turned down the offer.

Eugene had still more chances to get a new library. In 1991, a study commissioned by the city concluded that a building vacated by Sears in downtown Eugene — only a few blocks from the current library — would be a suitable site for the new facility. In a March election of that year, nearly three-quarters of those who voted agreed with this assessment. A subsequent vote in May of 1994 to authorize issuance of bonds to pay for refurbishing the building and moving the library, however, was defeated by a handful of absentee ballots. The argument was made that the ballot measure was too complex; it proposed using the funds for other city facilities as well as the library. A measure that focused on funding only the library would surely fare better. In November of 1994, such a measure was placed on the ballot. And defeated by fewer than 100 absentee ballots.


CHAPTER 16. Conflicting Objectives II: Multiattribute Utility Models with Interactions

The additive utility function described in Chapter 15 is an easy-to-use technique. It is incomplete, however, because it ignores certain fundamental characteristics of choices among multiattribute alternatives. We discussed one problem with proportional scores — they assume risk neutrality. Problem 15.22 demonstrated this in the important context of the common NPV choice criterion. More subtle are situations in which attributes interact. For example, two attributes may be substitutes for one another to some extent. Imagine a CEO who oversees several divisions. The simultaneous success of every division may not be terribly important; as long as some divisions perform well, cash flows and profits will be adequate. On the other hand, attributes can be complementary. An example might be the success of various phases of a research-and-development project. The success of each individual project is valuable in its own right. But the success of all phases might make possible an altogether new technology or process, thus leading to substantial synergistic gains in many ways. In this case, high achievement on all attributes (success in the various R&D phases) is worth more than the sum of the value obtained from the individual successes.

Such interactions cannot be captured by the additive utility function. That model is essentially an additive combination of preferences for individual attributes. To capture the kinds of interactions that we are talking about here, as well as risk attitudes, we must think more generally. Let us think in terms of a *utility surface*, such as the one depicted in Figure 16.1 for two attributes. Although it is possible to think about many attributes at once, we will develop multiattribute utility theory concepts using only two attributes. The ideas are readily extended to more attributes, and at the end of the chapter we will say a bit about multiattribute utility functions for three or more attributes.

Much of this chapter is fairly abstract and technical. To do a good job with the material, the mathematics are necessary. After theoretical development, which is

Figure 16.1
A utility surface for two attributes.



sprinkled with illustrative examples, we will process a complete example that involves the assessment of a two-attribute utility function for managing a blood bank. We also discuss briefly how to deal with three or more attributes and demonstrate such a model in a large-scale electric-utility example.

Multiattribute Utility Functions: Direct Assessment

To assess a utility function like the one in Figure 16.1, we can use the same basic approach that we already have used. For example, consider the reference-gamble method. The appropriate reference gamble has the worst pair (x^-, y^-) and the best pair (x^+, y^+) as the two possible outcomes:

Win (x^+, y^+) with probability p
 (x^-, y^-) with probability $1 - p$

Now for any pair (x, y) , where $x^- \leq x \leq x^+$ and $y^- \leq y \leq y^+$, find the probability p to use in the reference gamble that will make you indifferent between (x, y) and the reference gamble. As before, you can use p as your utility $U(x, y)$ because $U(x^+, y^+) = 1$, and $U(x^-, y^-) = 0$. Figure 16.2 shows the decision tree that repre-

Figure 16.2
Directly assessing a multiattribute utility. The probability that makes you indifferent between the lottery and the sure thing is your utility value for (x, y) .





Figure 16.3
Sketching indifference curves. The point values are the assessed utility values for the corresponding (x, y) pair.



sents the assessment situation. This is simply the standard probability-equivalent utility assessment technique that we have seen before.

You can see that we will wind up with many utility numbers after making this assessment for a reasonable number of (x, y) pairs. There may be several pairs with the same utility, and you should be indifferent among such pairs. Thus, (x, y) pairs with the same utilities must fall on an indifference curve. One approach to understanding your multiattribute preferences is simply to plot the assessed points on a graph, as in Figure 16.3, and sketch rough indifference curves.

To find a good representation of preferences through direct assessment, however, there is a drawback: You must assess utilities for a substantial number of points. And even though it is straightforward to see how this approach might be extended to three or more attributes, the more being considered, the more points you must assess, and the more complicated graphical representations become. An easier way would be convenient.

Another approach that would ease the assessment burden would be to think about a multiattribute utility function that is made up of the individual utility functions. Mathematically, we might represent the most general case as

$$U(x,y) = f[U_X(x), U_Y(y)]$$

The $f[\bullet, \bullet]$ notation means that $U(x, y)$ is a function of the individual utility functions $U_X(x)$ and $U_Y(y)$. In Chapter 15, the form we used was

In this chapter we will consider

$$U(x, y) = c_1 + c_2 U_X(x) + c_3 U_Y(y) + c_4 U_X(x) U_Y(y)$$

The importance of any such formulation is that it greatly eases the assessment burden; as in Chapter 15, we only require the individual utility functions and enough information to put them together. Being able to break down the multiattribute utility

function this way sometimes is called *separability*; the overall utility function can be "separated" into chunks that represent different attributes.

Is such an arrangement possible? Yes, but it requires some interesting conditions for the combined utility function. These conditions concern how the preferences interact among the attributes, a point suggested in the beginning of this chapter. We will digress briefly to discuss these conditions.

Independence Conditions

Preferential Independence

One thing we need in order to have the kind of separability mentioned above is *mutual preferential independence*. An attribute Y is said to be preferentially independent of X if preferences for specific outcomes of Y do not depend on the level of attribute X . As an example, let Y be the time to completion of a project and X its cost. If we prefer a project time of 5 days to one of 10 days, assuming that the cost is 100 in each case, and if we also prefer a project time of 5 days to one of 10 days if the cost is 200 in both cases, then Y is preferentially independent of X ; it does not matter what the cost is — we still prefer the shorter completion time.

We need mutual preferential independence, so we also need the cost to be preferentially independent of the completion time. If we prefer lower cost no matter what the completion time, then X is preferentially independent of Y . Then we can say that the two attributes are mutually preferentially independent.

Preferential independence seems to be a pretty reasonable condition to assume, especially in cases like the one involving costs and time to completion. But it is easy to imagine situations in which preferential independence might not hold. For example, in Chapter 15 it was suggested that your preference for amount of homework effort might depend on course topic. Bunn (1984) relates a nice hypothetical example in which preferential independence might not hold. Consider a decision with outcomes that affect both the place where you live and the automobile that you drive. Let X be an outcome variable that could denote either Los Angeles or an African farm, and Y an outcome variable denoting either a Cadillac or a Land Rover. The value of X (whether you live in Los Angeles or on an African farm) may well affect your preference for a Cadillac or a Land Rover. Therefore, Y would not be preferentially independent of X . Consider the reverse: You may prefer Los Angeles to an African farm (or vice versa) regardless of the car you own. Thus, one attribute would be preferentially independent of the other, but the two are not mutually preferentially independent.

It probably is fair to say that mutual preferential independence holds for many people and many situations, or that at least it is a reasonable approximation. Mutual preferential independence is like the decomposability property for an objectives hierarchy. If a decision maker has done a good job of building a decomposable hierarchy, mutual preferential independence probably is a reasonable assumption. But it should never be taken for granted.

If you recall the discussion about value functions and utility functions from Chapter 15, you will realize that mutual preferential independence, being about sure

outcomes, is a condition that applies to value functions. In fact, mutual preferential independence is exactly the condition that is needed for the additive (ordinal) utility function to be appropriate when making a decision with no uncertainty. That is, when a decision maker's preferences display mutual preferential independence, then the additive utility function, assessed using any of the techniques described in Chapter 15, is appropriate for decisions under certainty. Once we move to decision under uncertainty, however, mutual preferential independence is not quite strong enough. Although it is a necessary condition for obtaining separability of a (cardinal) multiattribute utility function, it is not sufficient. Thus, we must look at stronger conditions.

Utility Independence

Utility independence is slightly stronger than preferential independence. An attribute Y is considered utility independent of attribute X if preferences for *uncertain choices* involving different levels of Y are independent of the value of X . Imagine assessing a certainty equivalent for a lottery involving only outcomes in Y . If our certainty equivalent amount for the Y lottery is the same no matter what the level of X , then Y is utility independent of X . If X also is utility independent of Y , then the two attributes are mutually utility independent.

Utility independence clearly is analogous to preferential independence, except that the assessments are made under conditions of uncertainty. For the project evaluation example above, suppose we assess that the certainty equivalent for an option giving, say, a 50% chance of $Y = 5$ and a 50% chance of $Y = 10$ does not depend on the level at which the cost X is fixed. As long as our preferences for lotteries in the completion-time attribute are the same (as, say, measured by their certainty equivalents) regardless of the fixed level of cost, then completion time is utility independent of cost.

Keeney and Raiffa (1976) discuss an example in which utility independence might not hold. Suppose that X and Y are the rates of serious crime in two precincts of a metropolitan police division. In determining the joint utility function for this region's police chief, the issue of utility independence for X and Y would be faced. With Y fixed at 5, a relatively low rate of crime, he may be quite risk-averse to rates of crime in region X . He may not want to appear as though he is neglecting a particular precinct. Thus, his certainty equivalent for an option giving a 50% chance of $X = 0$ and a 50% chance of $X = 30$ may be 22 when Y is fixed at 5. If Y were fixed at the higher rate of 15, however, his certainty equivalent may be less risk-aversely assessed at 17. Thus, one must not assume that utility independence will hold in all cases. Even so, almost all reported multiattribute applications assume utility independence and thus are able to use a decomposable utility function.

Determining Whether Independence Exists

How can you determine whether your preferences are preferentially independent? The simplest approach is to imagine a series of paired comparisons that involve one of the at-

tributes. With the other attribute fixed at its lowest level, decide which outcome in each pair is preferred. Once this is done, imagine changing the level of the fixed attribute. Would your comparisons be the same? Would the comparisons be the same regardless of the fixed level of the other attribute? If so, then preferential independence holds.

Determining whether independence holds is a rather delicate matter. The following sample dialogue is taken from Keeney and Raiffa (1976). The notation has been changed to correspond to our notation here.

ANALYST. I would now like to investigate how you feel about various Y values when we hold fixed a particular value of X . For example, on the first page of this questionnaire [this is shown to the assessor] there is a list of 25 paired comparisons between Y evaluations; each element of the pair describes levels on the Y attributes alone. On this first page it is assumed that, throughout, the X evaluations are all the same, that is, x_1 [the fixed value for X is shown to the assessor]. Is this clear?

ASSESSOR. Crystal clear, but you are asking me for a lot of work.

ANALYST. Well, I have a devious purpose in mind, and it will not take as much time as you think to find out what I want. Now on the second page of the questionnaire [this is shown to the assessor] the identical set of 25 paired comparisons are repeated, but now the fixed, common level on the X attribute is changed from x_1 to x_2 [value is shown to the assessor]. Are you with me?

ASSESSOR. All the way.

ANALYST. On page 3, we have the same 25 paired comparisons but now the common value of the X value is x_3 [shown to the assessor].

ASSESSOR. You said this would not take long.

ANALYST. Well now, here comes the punchline. Suppose that you painstakingly respond to all of the paired comparisons on page 1 where x_1 is fixed. Now when you go to the next page would your responses change to these same 25 paired comparisons?

ASSESSOR. Let's see. In the second page all paired comparisons are the same except x_1 is replaced with x_2 . What difference should that make?

ANALYST. Well, you tell me. If we consider this first comparison [pointed to on the questionnaire] does it make any difference if the X values are fixed at x_1 or x_2 ? There could be some interaction concerning how you view the paired comparison depending on the common value of X .

ASSESSOR. I suppose that might be the case in some other situation, but in the first comparison I prefer the left alternative to the right no matter what the X value is . . . as long as they are the same.

ANALYST. Okay. Would you now feel the same if you consider the second paired comparison?

ASSESSOR. Yes. And the third and so on. Am I being naive? Is there some trick here?

ANALYST. No, not at all. I am just checking to see if the X values have any influence on your responses to the paired comparisons. So I gather that you are telling me that your responses on page 1 carry over to page 2.

ASSESSOR. That's right.

ANALYST. And to page 3, where the X value is held fixed at x_3 [shown to assessor]?

ASSESSOR. Yes.

ANALYST. Well, on the basis of this information I now pronounce that for you attribute Y is preferentially independent of attribute X .

ASSESSOR. That's nice to know. **ANALYST.** That's all I wanted to find out.

ASSESSOR. Aren't you going to ask me to fill out page 1? **ANALYST.** No. That's too much work. There are less painful ways of getting that information.

Source: Keeney, R., and H. Raiffa (1976) *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: Wiley. Reprinted in 1993 by Cambridge University Press. Copyright Cambridge University Press.

The dialogue describes how to check for preferential independence. Checking for utility independence would be much the same, except that the paired comparisons would be comparisons between lotteries involving attribute Y rather than sure outcomes. As long as the comparisons remain the same regardless of the fixed value for X , then Y can be considered utility independent of X . Of course, to establish mutual preferential or utility independence, the roles of X and Y would have to be reversed to determine whether paired comparisons of outcomes or lotteries in X depended on fixed values for Y . If each attribute turns out to be independent of the other, then mutual utility or preferential (whichever is appropriate) independence holds.

Using Independence

If a decision maker's preferences show mutual utility independence, then a two-attribute utility function can be written as a composition of the individual utility functions. As usual, the least preferred outcome (x^-, y^-) is assigned the utility value 0, and the most preferred pair (x^+, y^+) is assigned the utility value 1.

Under mutual utility independent preferences, the two-attribute utility function can be written as

$$U(x, y) = k_X U_X(x) + k_Y U_Y(y) + (1 - k_X - k_Y) U_X(x) U_Y(y),$$

where

$U_X(x)$ = utility function on X scaled so that $U_X(x^-) = 0$ and $U_X(x^+) = 1$

$U_Y(y)$ = utility function on Y scaled so that $U_Y(y^-) = 0$ and $U_Y(y^+) = 1$

$$k_X = U(x^+, y^-)$$

$$k_Y = U(x^-, y^+)$$

The product term $U_X(x) U_Y(y)$ in this utility function is what permits the modeling of interactions among attributes. The utility functions U_X and U_Y are *conditional utility functions*, and each must be assessed with the other attribute fixed at a particular level. (For example, in assessing U_Y , imagine that X is fixed at a specific level.) To understand the last two conditions, all we must do is plug the individual utilities into the equation. For example,

$$U(x^+, y^-) = k_X U_X(x^+) + k_Y U_Y(y^-) + (1 - k_X - k_Y) U_X(x^+) U_Y(y^-) = k_X(1) + k_Y(0) + (1 - k_X - k_Y)(1)(0) = k_X$$

This multiattribute utility function, called a *multilinear* expression, is not as bad as it looks! Look at it from the point of view of the X attribute. Think about fixing Y at a value (say, y_a); you get a conditional utility function for X , given that Y is fixed at y_a :

$$U(x, y_a) = k_Y U_Y(y_a) + [k_X + (1 - k_X - k_Y) U_Y(y_a)] U_X(x)$$

Because Y is fixed at y_a , the terms $k_Y U_Y(y_a)$ and $[k_X + (1 - k_X - k_Y) U_Y(y_a)]$ are just constants. Thus, $U(x, y_a)$ is simply a scaled version of $U_X(x)$. Now change to another y (y_b). What happens to the utility function for X ? The expression now looks like

$$U(x, y_b) = k_Y U_Y(y_b) + [k_X + (1 - k_X - k_Y) U_Y(y_b)] U_X(x)$$

This is just another linear transformation of $U_X(x)$, and so $U(x, y_b)$ and $U(x, y_a)$ must be identical in terms of the way that lotteries involving the X attribute would be ranked. We have scaled the utility function $U_X(x)$ in two different ways, but the scaling does not change the ordering of preferences. Now, notice that we can do exactly the same thing with the Y attribute; for different fixed values of X (x_a and x_b), the conditional utility functions are simply linear transformations of each other:

$$\begin{aligned} U(x_a, y) &= k_X U_X(x_a) + [k_Y + (1 - k_X - k_Y) U_X(x_a)] U_Y(y) \\ U(x_b, y) &= k_X U_X(x_b) + [k_Y + (1 - k_X - k_Y) U_X(x_b)] U_Y(y) \end{aligned}$$

No matter what the level of one attribute, preferences over lotteries in the second attribute (Y) stay the same. This was the definition of utility independence in the last section. We have mutual utility independence because the conditional utility function for one attribute stays essentially the same no matter which attribute is held fixed.

Additive Independence

Look again at the multiattribute utility function

$$U(x, y) = k_X U_X(x) + k_Y U_Y(y) + (1 - k_X - k_Y) U_X(x) U_Y(y)$$

If $k_X + k_Y = 1$, then the utility function turns out to be simply additive:

$$U(x, y) = k_X U_X(x) + (1 - k_X) U_Y(y)$$

If this is the case, we only have to assess the two individual utility functions $U_X(x)$ and $U_Y(y)$ and the weighting constant k_X . This would be convenient: It would save having to assess k_Y . Of course, this is just the additive utility function from Chapter 15. How is this kind of multiattribute utility function related to the independence conditions? To be able to model preferences accurately with this additive utility function, we need *additive independence*, an even stronger condition than utility independence.

The statement of additive independence is the following: Suppose X and Y are mutually utility independent, and you are indifferent between Lotteries A and B:

- A** (x^-, y^-) with probability 0.5
 (x^+, y^+) with probability 0.5
- B** (x^-, y^+) with probability 0.5
 (x^+, y^-) with probability 0.5

If this is the case, then the utility function can be written as the weighted combination of the two utility functions, $U(x, y) = k_X U_X(x) + (1 - k_X) U_Y(y)$. You can see by writing out the expected utilities of the lotteries that they are equivalent:

$$\begin{aligned} EU(A) &= 0.5[k_X U_X(x^-) + (1 - k_X) U_Y(y^-)] + 0.5[k_X U_X(x^+) + (1 - k_X) U_Y(y^+)] = \\ &0.5[k_X U_X(x^-) + (1 - k_X) U_Y(y^-) + k_X U_X(x^+) + (1 - k_X) U_Y(y^+)] \end{aligned}$$

$$\begin{aligned} EU(B) &= 0.5[k_X U_X(x^-) + (1 - k_X) U_Y(y^+)] + 0.5[k_X U_X(x^+) + (1 - k_X) U_Y(y^-)] = \\ &0.5[k_X U_X(x^-) + (1 - k_X) U_Y(y^+) + k_X U_X(x^+) + (1 - k_X) U_Y(y^-)] = EU(A) \end{aligned}$$

The intuition behind additive independence is that, in assessing uncertain outcomes over both attributes, we only have to look at one attribute at a time, and it does not matter what the other attribute's values are in the uncertain outcomes. This sounds a lot like utility independence. The difference is that, in the case of additive independence, changes in *lotteries* in one attribute do not affect preferences for lotteries in the other attribute; for utility independence, on the other hand, changes in *sure levels* of one attribute do not affect preferences for lotteries in the other attribute. Here is another way to say it: When we are considering a choice among risky prospects involving multiple attributes, if additive independence holds, then we can compare the alternatives one attribute at a time. In comparing Lotteries A and B above, we are indifferent because (1) for attribute X , each lottery gives us a 50% chance at x^- and a 50% chance at x^+ ; and (2) for Y , each lottery gives us a 50% chance at y^- and a 50% chance at y^+ . Looking at the attributes one at a time, the two lotteries are the same.

The additive utility function from Chapter 15 requires additive independence of preferences across attributes in order to be an accurate model of a decision maker's preferences in decisions under uncertainty. Think back to some of the examples or

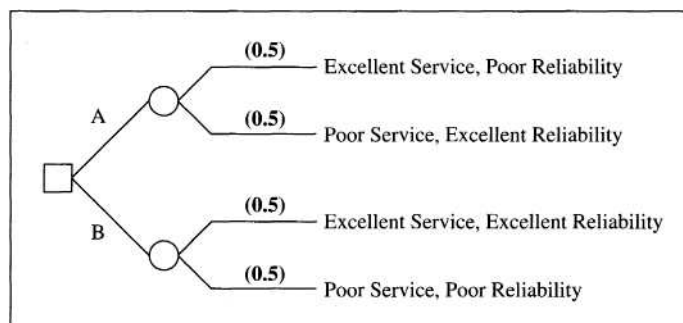
problems. Do you think this idea of additive independence makes sense in purchasing a car? Think about reliability and quality of service, two attributes that might be important in this decision. When you purchase a new car, you do not know whether the reliability will be high or low, and you may not know the quality of the service. To some extent, however, the two attributes are substitutes for each other. Suppose you faced the hypothetical decision shown in Figure 16.4. Would you prefer Lottery A or B? Most of us probably would take A. If you have a clear preference for one or the other, then additive independence cannot hold.

von Winterfeldt and Edwards (1986) discuss reports from behavioral decision theory that indicate additive independence usually does not hold. If this is the case, what is the justification for the use of the additive utility function? Many multiattribute decisions that we make involve little or no uncertainty, and evidence has shown that the additive model is reasonable for most situations under conditions of certainty. And in extremely complicated situations with many attributes, the additive model may be a useful rough-cut approximation. It may turn out that considering the interactions among attributes is not critical to the decision at hand.

Finally, it is possible to use simple approximation techniques to include interactions within the additive utility framework. For example, suppose that we have a decision problem with many attributes that are, for the most part, additively independent of one another. The additive representation of the utility function does not allow for any interaction among the attributes. If this is appropriate for almost all of the possible outcomes, then we may use the additive representation while including a specific "bonus" or "penalty" (depending on which is appropriate) for those outcomes with noticeable interaction effects.

Raiffa (1982) has an interesting example. Suppose a city is negotiating a new contract with its police force. Two attributes are (1) increase in vacation for officers who have less than five years of service and (2) increase in vacation for officers who have more than five years of service. The city loses points in an additive value function for increases in vacation for either group. If either group is held to no increase, then the city loses no points for that particular group. But the city would be happy if all its officers could be held to no increase in vacation time; thus, no precedent is set for the other group. To capture this interaction, the city gets a "bonus" of some points in its overall utility if there is no increase in vacation for either group.

Figure 16.4
An assessment lottery for a car purchase. If you have a clear preference for A or B, then additive independence cannot hold.



Substitutes and Complements

In the multiattribute utility function, the interaction between the attributes is captured by the term $(1 - k_X - k_Y) U_X(x) U_Y(y)$. How can we interpret this? Keeney and Raiffa (1976) give an interesting interpretation of the coefficient $(1 - k_X - k_Y)$. The sign of $(1 - k_X - k_Y)$ can be interpreted in terms of whether x and y are complements or substitutes for each other. Suppose $(1 - k_X - k_Y)$ is positive. Now examine the multiattribute utility function:

$$U(x, y) = k_X U_X(x) + k_Y U_Y(y) + (1 - k_X - k_Y) U_X(x) U_Y(y)$$

Preferred values of X and Y will give high values to the conditional utility functions, and the positive $(1 - k_X - k_Y)$ will drive up the overall utility for the pair even higher. Thus, if $(1 - k_X - k_Y)$ is positive, the two attributes complement each other. On the other hand, if $(1 - k_X - k_Y)$ is negative, high values on each scale will result in a high product term, which must be subtracted in the multiattribute preference value. In this sense, preferred values of each attribute work against each other. But if one attribute is high and the other low, the subtraction effect is not as strong. Thus, if $(1 - k_X - k_Y)$ is negative, the two attributes are substitutes.

Keeney and Raiffa (1976) offer two examples. In one, imagine a corporation with two divisions that operate in different markets altogether, and let profits in each division represent two attributes of concern to the president. To a great extent, success by the two divisions could be viewed as substitutes. That is, if profit from one division was down while the other was up, the firm would get along fine. Financial success by one division would most likely ensure the overall success of the firm.

For an example of the complementary case, Keeney and Raiffa consider the problem a general would face in a battle being fought on two fronts. If we let the consequences on the two fronts represent two distinct attributes, then these two attributes may be complementary. That is, defeat on one front may be almost as bad as defeat on both fronts, and a completely successful outcome may be guaranteed only by victory on both.

Assessing a Two-Attribute Utility Function

Now that we have seen the basics of two-attribute utility functions, we are ready to assess one. The procedure is relatively straightforward. First, we determine whether mutual utility independence holds. Provided that it does, we then assess the individual utility functions. Finally, the scaling constants are determined in order to put the individual utility functions together.

THE BLOOD BANK

In a hospital blood bank it is important to have a policy for deciding how much of each type of blood should be kept on hand. For any particular year, there is a "shortage rate," the percentage of units of blood demanded but not filled from stock because of shortages. Whenever there is a shortage, a special order must be placed to locate the required blood elsewhere or to locate donors. An operation may be postponed, but only rarely will a blood shortage result in a death. Naturally, keeping a lot of blood stocked means that a shortage is less likely. But there is also a rate at which blood is "outdated," or kept on the shelf the maximum amount of time, after which it must be discarded. Although having a lot of blood on hand means a low shortage rate, it probably also would mean a high outdated rate. Of course, the eventual outcome is unknown because it is impossible to predict exactly how much blood will be demanded. Should the hospital try to keep as much blood on hand as possible so as to avoid shortages? Or should the hospital try to keep a fairly low inventory in order to minimize the amount of outdated blood discarded? How should the hospital blood bank balance these two objectives? [*Source: Keeney and Raiffa (1976).*]

The outcome at the blood bank depends on uncertain demand over the year as well as the specific inventory policy (stock level) chosen. Thus, we can think of each possible inventory policy as a lottery over uncertain outcomes having two attributes, shortage and outdated. Shortage is measured as the annual percentage of units demanded but not in stock, while outdated is the percentage of units that are discarded due to aging. A high stock level probably will lead to less shortage but more outdated, and a low stock level will lead to more shortage and less outdated. To choose an appropriate stock level, we need to assess both the probability distribution over shortage and outdated outcomes for each possible stock level and the decision maker's utility function over these outcomes. Because each outcome has two attributes, we need a two-attribute utility function. Here we focus on the assessment of the utility function through the following steps:

- 1 The first step was to explain the problem to the nurse in charge of ordering blood. Maintaining an appropriate stock level was her responsibility, so it made sense to base an analysis of the problem on her personal preferences. She understood the importance of the problem and was motivated to think hard about her assessments. Without such understanding and motivation on her part, the entire project probably would have failed.
- 2 It was established that the annual outdated and shortage rates might range from 10% (worst case) to 0% (best case).
- 3 Did mutual utility independence hold? The nurse assessed a certainty equivalent for uncertain shortage rates (Attribute X), given a fixed outdated rate (Attribute Y). The certainty equivalent did not change for different outdated rates. Thus,

shortage was found to be utility independent of outdating. Similar procedures showed the reverse to be true as well. Thus, shortage and outdating were mutually utility independent, implying the multilinear form for the utility function.

- 4 The next step was to assess the conditional utility functions $U_X(x)$ and $U_Y(y)$. In each case, the utility function was assessed conditional on the other attribute being held constant at 0. To assess $U_X(x)$, it was first established that preferences decreased as x increased. Using the lotteries that had been assessed earlier in the utility independence step, an exponential utility function was determined. Setting $U_X(0) = 1$ (best case) and $U_X(10) = 0$ (worst case), the utility function was

$$U_X(x) = 1 + 0.375(1 - e^{x/7.692})$$

Likewise, the second utility function was determined using the previously assessed certainty equivalents, and again an exponential form was used. The utility function was

$$U_Y(y) = 1 + 2.033(1 - e^{y/25})$$

This utility function also has $U_Y(0) = 1$ and $U_Y(10) = 0$.

- 5 Assessing the weights k_X and k_Y is the key to finding the two-attribute utility function. The trick is to use as much information as possible to set up equations based on indifferent outcomes and lotteries, and then to solve the equations for the weights. Because we have two unknowns, k_X and k_Y , we will be solving two equations in two unknowns. To set up two equations, we will need two utility assessments.

Recall that the multilinear form can be written as

$$U(x, y) = k_X U_X(x) + k_Y U_Y(y) + (1 - k_X - k_Y) U_X(x) U_Y(y)$$

We also know that

$$U(10, 0) = k_Y$$

$$U(0, 10) = k_X$$

These follow from steps 3 and 4 above, substituting $x^- = 10, x^+ = 0, y^- = 10$, and $y^+ = 0$.

The nurse determined that she was indifferent between the two outcomes ($x = 4.75, y = 0$) and ($x = 0, y = 10$). This first assessment indicates that, for her, avoiding shortages is more important than avoiding outdating. We can substitute each one of these points into the expression for the utility function, establishing the first equation relating k_X and k_Y :

$$U(4.75, 0) = k_X U_X(4.75) + k_Y U_Y(0) + (1 - k_X - k_Y) U_X(4.75) U_Y(0) = k_X U_X(4.75) + k_Y(1) + (1 - k_X - k_Y) U_X(4.75) (1)$$

Because she was indifferent between (4.75, 0) and (0, 10), we have

$$\begin{aligned} U(4.75, 0) &= U(0, 10) && \text{(because she is indifferent)} \\ &= k_X && \text{[from } U(0, 10) = k_X, \text{ above]} \end{aligned}$$

Substituting, we obtain

$$\begin{aligned} k_X &= k_X U_X(4.75) + k_Y + (1 - k_X - k_Y) U_X(4.75) & (16.1) \\ &= k_Y + (1 - k_Y) U_X(4.75) = k_Y + (1 - \\ &k_Y)[1 + 0.375(1 - e^{4.75/7.692})] = k_Y + (1 - \\ &k_Y)0.68 = 0.68 + 0.32k_Y \end{aligned}$$

In the second assessment, the decision maker concluded that she was indifferent between the outcome (6, 6) and a 50-50 lottery between the outcomes (0, 0) and (10, 10). Using this assessment, we can find $U(6, 6)$:

$$\begin{aligned} U(6, 6) &= 0.5 U(0, 0) + 0.5 U(10, 10) = \\ &0.5(1) + 0.5(0) \\ &= 0.5 \end{aligned}$$

This is just a standard assessment of a certainty equivalent for a 50-50 gamble between the best and worst outcomes. Now substitute $U(6, 6) = 0.5$ into the two-attribute utility function to find a second equation in terms of k_X and k_Y :

$$\begin{aligned} 0.5 &= U(6, 6) \\ &= k_X U_X(6) + k_Y U_Y(6) + (1 - k_X - k_Y) U_X(6) U_Y(6) \end{aligned}$$

Substituting the values $X = 6$ and $Y = 6$ into the formulas for the individual utility functions gives

$$\begin{aligned} U_X(6) &= \\ 0.56 U_Y(6) & \\ &= 0.45 \end{aligned}$$

Now plug these into the equation for $U(6, 6)$ to get

$$0.5 = k_X(0.56) + k_Y(0.45) + (1 - k_X - k_Y)(0.56)(0.45)$$

which simplifies to

$$0.248 = 0.308 k_X + 0.198 k_Y \quad (16.2)$$

Now we have two linear equations in k_X and k_Y —Equations (16.1) and (16.2):

$$\begin{aligned} k_X &= 0.680 + 0.320 k_Y 0.248 \\ &= 0.308 k_X + 0.198 k_Y \end{aligned}$$

Solving these two equations simultaneously for k_X and k_Y , we find that $k_X = 0.72$ and $k_Y = 0.13$. Thus, the two-attribute utility function can be written as

$$U(x, y) = 0.72 U_X(x) + 0.13 U_Y(y) + 0.15 U_X(x) U_Y(y)$$

where $U_X(x)$ and $U_Y(y)$ are given by the exponential utility functions defined above. Now we can find the utility for any (x, y) pair (as long as the x 's and y 's are each between 0 and 10, the range of the assessments). Any policy for ordering blood can be evaluated in terms of its expected utility. Table 16.1 shows utilities for different

possible outcomes, and Figure 16.5 shows the indifference curves associated with the utility function. From Figure 16.5, we can verify the conditions and assessments that were used:


$$\begin{aligned}
 U(0, 0) &= 1 & U(10, 10) &= 0 \\
 U(10, 0) &= 0.13 = k_Y & U(0, 10) &= U(4.75, 0) = 0.72 = k_X \\
 U(6, 6) &= 0.50
 \end{aligned}$$

The final assessed utility function is readily interpreted. The large value for k_X relative to k_Y means that the nurse is much more concerned about the percentage of shortage than she is about the percentage of outdated. This makes sense; most of us would agree that the objective of the blood bank is primarily to save lives, and we probably would rather throw out old blood than not have enough on hand when it is

Table 16.1
Utility values for shortage and outdated in the blood bank.

x Values (Shortage)	y Values (Outdating)					
	0	2	4	6	8	10
0	1.00	0.95	0.90	0.85	0.79	0.72
2	0.90	0.86	0.81	0.76	0.70	0.64
4	0.78	0.74	0.69	0.64	0.59	0.54
6	0.62	0.58	0.54	0.50	0.45	0.40
8	0.40	0.37	0.34	0.31	0.27	0.23
10	0.13	0.11	0.08	0.06	0.03	0.00

Figure 16.5
Indifference curves for nurse's utility function for shortage and outdated. The numbers are $U(x, y)$ for the corresponding indifference curve.



needed. The fact that $k_X + k_Y < 1$ means that the two attributes are complements rather than substitutes. We can see this in Figure 16.5. For example, imagine an outcome of (8, 8), for which the utility is 0.27 (Point A). Now imagine improving the shortage percentage to zero. This would increase the utility value to $U(0, 8) = 0.79$, for an approximate net increase of 0.52. On the other hand, improving the outdated percentage the same amount results in $U(8, 0) = 0.40$, for a net increase of 0.13. If we increased both at the same time, we would have $U(0, 0) = 1.00$, an increase of 0.73. The increase in utility from increasing both at once is greater than the sum of the individual increases ($0.73 > 0.52 + 0.13 = 0.65$). This is the sense in which there is an interaction. This kind of phenomenon is impossible in the additive utility function.

Three or More Attributes (Optional)

When the decision problem involves three or more objectives, modeling preferences is more difficult. Building a utility function that will permit interactions across many attributes can become complex. Under certain conditions, however, the *multiplicative* utility function can be used. Let X_i denote the i th attribute, and $U_i(x_i)$ and k_i the corresponding individual utility function and scaling constant. The multiplicative utility function for n different attributes is given by the equation

$$1 + k U(x_1, x_2, \dots, x_n) = \prod_{i=1}^n [k k_i U_i(x_i) + 1] \quad (16.3)$$

where k is a nonzero solution to the equation


$$1 + k = \prod_{i=1}^n (1 + k k_i)$$

The k_i 's have the same meaning they had in the two-attribute case:

$$U(x_1^-, \dots, x_{i-1}^-, x_i, x_{i+1}^-, \dots, x_n^-) = k_i$$

That is, k_i is the utility of an outcome having the best level on attribute X_i and worst on all others. Thus, we can assess the k_i 's directly through our standard reference-lottery approach, which is shown in Figure 16.6. We have the reference lottery with the best possible and worst possible outcomes versus the sure thing having the best level on attribute X_i and the worst level on all others. The probability p_i that makes us indifferent between the lottery and the sure thing in our utility for the sure thing and hence is equal to k_i . The decision maker can assess each of the k_i 's in this way and then find the scaling constant k that satisfies the condition given above. Finally, putting the individual utility functions together with the scaling constants gives the overall utility function.

Figure 16.6



Assessing scaling constants for the multiplicative utility function. The p_i that makes the decision maker indifferent between the lottery and the sure thing is the utility of the sure outcome and hence equals k_i .

The multiplicative utility function requires a fairly strong version of utility independence. Strictly speaking, each subset of the attributes must be utility-independent of the remaining attributes. This essentially means that we should be able to partition the attributes into two subsets in any way we want, and then consider lotteries in one subset, holding the attributes in the other subset fixed. As long as preferences for the lotteries do not depend on the level of the remaining attributes, the multiplicative utility function should provide a good model of the decision maker's preferences. If the multiplicative model is not appropriate, then a more general version of the multilinear model may be a possibility. For more details, consult Keeney and Raiffa (1976).

When Independence Fails

We have just dealt in depth with situations in which the assumption of mutual utility independence results in a reasonable model of preferences. This is not always true. Suppose we are interested in assessing a two-attribute utility function over attributes X and Y but have found that neither X nor Y is utility independent of the other. Then neither the multilinear nor additive forms for the utility function are appropriate. How can we obtain a reasonable $U(x, y)$ for decision-making purposes? Several possibilities exist. One is simply to perform a direct assessment as described at the beginning of this chapter. Pick the best and worst (x, y) pairs, assign them utility values of 1 and 0, and then use reference gambles to assess the utility values for other points. A second approach is to transform the attributes and proceed to analyze the problem with the new set. Of course, the new set of attributes still must capture the critical aspects of the problem, and they must be measurable. Take the example discussed above in which X and Y designate measures of the crime rates in two sections of a city. There may be a complicated preference structure for (x, y) pairs. For political reasons, the relative ordering of hypothetical lotteries for criminal activity in one section may be highly dependent on the level of crime in the other section. But sup-

pose we define $s = (x + y)/2$ and $t = |x - y|$. Then s may be interpreted as an average crime index for the city and t as an indicator of the balance of criminal activity between the two sections. Any (x, y) outcome implies a unique (s, t) outcome, so it is easy to transform from one set of attributes to the other. Furthermore, even though x and y may not be utility independent, it may be reasonable to model s and t as being utility independent, thus simplifying the assessment procedure.

One of the most difficult and subtle concepts in multiattribute utility theory is the notion of interaction among attributes. To use the additive utility function from Chapter 15, we must have no interaction at all. To use the multilinear utility function discussed in this chapter, we can have some interaction, but it must be of a limited form. (Any interactions must conform to the notion of utility independence.) The blood bank example demonstrated the nature of the interaction between two attributes that is possible with the multilinear utility function. In this section, we are concerned with situations in which the interactions are even more complicated than those that are possible in the multilinear case. Fortunately, evidence from behavioral research suggests that it is rarely, if ever, necessary to model extremely complex preference interactions.

Chapter 15 and 16 have presented many of the principles and decision-analysis techniques that are useful in making decisions in the face of conflicting objectives. Given the complexity of the techniques, it is important to keep in mind the modeling perspective that we have held all along. The objective of using the multiattribute decision-analysis techniques is to construct a model that is a reasonable representation of a decision maker's value structure. If minimal interactions exist among the attributes, then the additive utility function is appropriate. When attributes interact, then it may be necessary to consider multiattribute utility theory, just as we have in this chapter.

Multiattribute Utility in Action: BC Hydro

Can multiattribute utility modeling be of use? Consider the situation of the British Columbia Hydro and Power Authority (BC Hydro):

STRATEGIC DECISIONS AT BC HYDRO


In the late 1980s, BC Hydro found itself at an important juncture. Knowing that it would face many strategic decisions in the future, its directors wanted to prepare to make those decisions in the best possible way. It would have to make decisions regarding power-generation plants, placement of transmission lines, employee relations, how to communicate with public and special-interest groups, environmental impact of its activities and facilities, and policies for addressing large-scale problems such as

global warming. To ensure that many different decisions would be coordinated and would all serve the organization's interests, Mr. Ken Peterson was appointed as Director of Strategic Planning. In short, Peterson's job was to make BC Hydro the leader in strategic planning among North American utility companies. To help Peterson get started, BC Hydro had a general mission statement. Like most such statements, however, BC Hydro's was broad and lacked details. It was certainly not up to the task of coordinating the myriad of diverse decisions that BC Hydro would face.

Realizing that he would need a way to think systematically about BC Hydro's strategic alternatives, Peterson enlisted Ralph Keeney and Tim McDaniels to help identify the organization's objectives and construct a multiattribute utility function. The project is described in Keeney (1992) and Keeney and McDaniels (1992). The process began with interviews of key decision makers in the organization to identify fundamental objectives. Once a satisfactory set of fundamental objectives was established, the process continued through all of the steps necessary to define and assess

Figure 16.7

Two nonlinear utility functions for BC Hydro, (a) A risk-averse utility function for government dividend, (b) A risk-prone utility function for duration of outages to large customers.



Source: Keeney (1992, p. 364).

individual utility functions for the attributes and scaling constants for the multiattribute utility function. The final utility function involved 18 different attributes in six major groups (economics, environment, health and safety, equity, service quality, and public-service recognition). The hierarchy of objectives with corresponding attributes and ranges is shown in Table 16.2.

Keeney reports that the component utility functions were linear (i.e., proportional scores were used) for all but two of the attributes. The two that were not linear were annual government dividend and large customer outage duration. The individual utility functions for these two attributes are shown in Figure 16.7. Note that the dividend utility function shows risk aversion, but the outage utility function is risk-seeking.

The assessment proceeded by creating multiattribute utility functions for each of the groups and subgroups in Table 16.2. For example, two-attribute utility functions were created for service quality to small customers and large customers, each one covering the number of outages and outage duration for the corresponding customer type. These two utility functions were then combined with individual utility functions for new service and inquiries to create a four-attribute utility function for service. Finally, the service utility function was combined with five similar utility functions for each of the other major groups.

All of the multiattribute utility functions (including the overall function for the six major groups) were additive except for the economic utility function, which was multiplicative as in Equation (16.3). The reason for the nonadditive economic utility function is that the economic attributes were thought to be substitutes for each other to some extent. For the additive part of the model, it is possible to collapse the hierarchy and calculate the weight that any one attribute has in the overall utility function. This is done simply by multiplying the weight for the attribute times the weight for the group (or subgroup) in which it resides, and so on, until the top level of the objectives hierarchy is reached. (You will recall that we did this in the library example in Chapter 15.) The scaling weights, thusly calculated, are shown in Table 16.3 for the additive attributes. (This table is taken from Keeney's Table 12.5 (1992). Note that scaling constants are given for public and worker health-and-safety aggregates rather than for the individual mortality and morbidity attributes. The same is true for small and large customer service attributes.)

Table 16.3 also shows in parentheses the scaling constants for the six major attribute groups, thus indicating the relative importance of these top-level objectives. The most important of the six is the economic objective, accounting for almost 40% of the total weight. The table also shows the scaling constants for the three components in the economic utility function and the constant k that defines the interactions in this multiplicative submodel. The economic utility function is given by

$$1 - 0.76 U(x_1, x_2, x_3) = [1 - 0.76 (0.84) U_1(x_1)] [1 - 0.76 (0.18) U_2(x_2)][1 - 0.76 (0.30) U_3(x_3)] \quad (16.4)$$

where x_1 , x_2 , and x_3 represent levelized energy costs, annual dividend, and resource losses, respectively.

From Table 16.3, the overall utility of any given consequence can be calculated. First, find the individual for each attribute. Then, for the additive components,

Table 16.2		Worst Level	Best Level
Attributes and ranges for BC Hydro's strategic utility function.			
Economics			
Levelized cost of energy from new sources (1989 \$0.001/kWh)		55	35
Annualized dividend payable (1989 \$ million)		0	200
Economic cost or resource losses (1989 \$ million)		20	0
Environment			
Local Impacts			
Flora (hectares of mature forest lost)		10,000	0
Fauna (hectares of wildlife habitat lost)		10,000	0
Wilderness ecosystem (hectares of wilderness lost)		10,000	0
Recreation (hectares of recreational land lost)		10,000	0
Aesthetic (annual person-years of viewing high-voltage transmission lines in scenic terrain)		500,000	0
Global impact (megawatts generated from fossil fuels)		1000	0
Health and Safety			
Public			
Mortality (annual person-years of life lost)		100	0
Morbidity (annual person-years of "severe" disability)		1000	0
Employees			
Mortality (annual person-years of life lost)		100	0
Morbidity (annual person-years of lost work time)		1000	0
Equity			
Equitable pricing (constructed scale: see Keeney 1992, section 12.3)		0.5	0
Equitable compensation (annual average number of individuals who feel they are inequitably treated)		500	0
Service quality			
Small customers			
Outages (annual number per customer)		2	0
Outage duration (hours per outage)		24	0
Large customers			
Outages (annual number per customer)		2	0
Outage duration (hours per outage)		24	0
New service (installation time in workdays)		20	1
Inquiries (minutes until personal response)		1	0
Public-service recognition (constructed scale: see Keeney 1992, section 12.3)		0	4

Source: Keeney (1992, pp. 358-359).

Table 16.3
Scaling constants for
BC Hydro's strategic
utility function.

	Scaling Constant
Economics (0.395)	
<i>Constants for multiplicative function:</i>	
Levelized cost of energy	0.84
Annualized dividend	0.18
Resource cost	0.30
Scaling constant k	-0.76
<i>Constants for additive function:</i>	
Environment (0.250)	
Flora	0.023
Fauna	0.046
Wilderness ecosystem	0.093
Recreation	0.046
Aesthetic	0.023
Global impact	0.019
Health and Safety (0.089)	
Public	0.045
Worker	0.045
Equity (0.012)	
Equitable pricing	0.004
Equitable compensation	0.008
Service quality (0.250)	
Small customers	0.111
Large customers	0.125
New service	0.010
Inquiries	0.005
Public-service recognition (0.004)	0.004

multiply each utility value times its weight and add up the resulting products. For the economic attributes, calculate the utility from Equation (16.4) and multiply the result times 0.395, the weight given to the economic component. Finally, add this product to the sum of the weighted utilities from the previous step. Of course, no one would do this in such a painstaking way by hand. Fortunately, it is very straightforward to program such utility calculations either into a specialized utility program or even in an electronic spreadsheet.

Keeney (1992) shows how this utility function for BC Hydro can be used to generate important insights. For example, he shows the resource value in dollars of various attributes: Each hectare of wilderness lost is valued at \$2500, but a hectare

of flora lost is worth \$625. Each person-year of aesthetic deterioration is worth \$13. One statistical fatality, for either an employee or a member of the general public, is worth \$3 million. Many other economic trade-offs are listed in Keeney's Table 12.6 (1992). These are the amounts that BC Hydro should just be willing to sacrifice (but no more) in order to save a hectare, a life, or a person-year of aesthetic damage. In addition to these trade-offs, Keeney lists a number of decision opportunities that arose from the analysis, including studying the implications of resource losses on economic activity, developing a database for fatalities, studying public environmental values, and understanding the process of responding to telephone inquiries.

Did the definition of strategic objectives and a strategic utility function have an impact on BC Hydro? Ken Peterson is quoted as saying:

The structured set of objectives has influenced BC Hydro planning in many contexts. Two examples include our work to develop a decision framework for supply planning, and a case study of an investment to upgrade reliability, both of which have adopted a multiple objective structure. Less obvious has been an evolution in how key senior planners view planning issues. The notion of a utility function over a range of objectives (rather than a single objective, like costs) is evident in many planning contexts. The specific trade-offs in the elicitation process are less important than the understanding that trade-offs are unavoidable in electricity utility decisions and that explicit, well-structured, informed trade-offs can be highly useful (*Interfaces*, November-December 1992, p. 109).

The use of multiattribute utility models has been a source of insight for many decision makers in diverse decision situations. For example, Keeney and Raiffa (1976) report applications involving fire department operations, strategic decision making by a technical consulting firm, evaluation of computer systems, siting of nuclear power facilities, the analysis of sites for the Mexico City airport, and many others. Other applications are described in von Winterfeldt and Edwards (1986).

SUMMARY

We have continued the discussion of making decisions in the face of conflicting objectives. Much of the chapter has been a rather technical discussion and treatment of independence conditions: preferential independence, utility independence, and additive independence. The differences among these have to do with the presence or absence of uncertainty. Preferential independence says that preferences for sure outcomes in one attribute do not depend on the level of other attributes. Utility independence requires that preferences for gambles or lotteries in an attribute do not depend on the level of other attributes. Additive independence is still stronger: Preferences over lotteries in one attribute must not depend on lotteries in the other attributes. The blood bank example showed how to apply mutual utility independence to assess a two-dimensional utility function that includes an interaction term. We saw that the attributes of shortage and outdated were complementary; they work together to increase the decision maker's utility. We briefly saw what to do when

there are three or more attributes, and what to do if no independence properties hold. Finally, the BC Hydro application demonstrated multiattribute utility in a large-scale organizational setting.

EXERCISES

- 16.1** Explain what is meant when we speak of interaction between attributes. Why would the additive utility function from Chapter 15 be inappropriate if two attributes interact?
- 16.2** What are the advantages and disadvantages of directly assessing a multiattribute utility function?
- 16.3** Explain preferential independence in your own words. Can you cite an example from your own experience in which preferential independence holds? Can you cite an example in which it does not hold?
- 16.4** Explain in your own words the difference between preferential independence and utility independence.
- 16.5** Explain in your own words the difference between utility independence and additive independence. Why is it important to understand the concept of additive independence?
- 16.6** Suppose that a company would like to purchase a fairly complicated machine to use in its manufacturing operation. Several different machines are available, and their prices are more or less equivalent. But the machines vary considerably in their available technical support (Attribute X) and reliability (Attribute Y). Some machines have a high degree of reliability and relatively low support, while others are less reliable but have excellent field support. The decision maker determined what the best and worst scenarios were for both attributes, and an assessment then was made regarding the independence of the decision maker's preferences for these attributes. It was determined that utility independence held. Individual utility functions $U_X(x)$ and $U_Y(y)$ were assessed. Finally, the decision maker was found to be indifferent in comparing Lottery A with its alternative B:
- A** Best on both reliability and support with probability 0.67
Worst on both reliability and support with probability 0.33

B Best on reliability and worst on support

The decision maker also was indifferent between

C Best on both reliability and support with probability 0.48
Worst on both reliability and support with probability 0.52

D Worst on reliability and best on support

- a What are the values for k_X and k_Y ? Write out the decision maker's full two-attribute utility function for reliability and support.
- b Are the two attributes substitutes or complements? Explain, both intuitively and on the basis of k_X and k_Y .

QUESTIONS AND PROBLEMS

- 16.7** A hospital administrator is making a decision regarding the hospital's policy of treating individuals who have no insurance coverage. The policy involves examination of a prospective patient's financial resources. The issue is what level of net worth should be required in order for the patient to be provided treatment. Clearly, there are two competing objectives in this decision. One is to maximize the hospital's revenue, and the other is to provide as much care as possible to the uninsured poor. Attributes to measure achievement toward these two objectives are (1) prospective revenue (R), and (2) percentage of uninsured poor who are treated (P).

The two attributes were examined for independence, and the administrator concluded that they were mutually utility independent. Utility functions $U_R(r)$ and $U_P(p)$ for the two attributes then were assessed. Then two more assessments were made. Lottery A and its certain alternative B were judged to be equivalent by the administrator:

- A** Best on revenue, worst on treating poor with probability 0.65
Worst on revenue, best on treating poor with probability 0.35
- B** Levels of revenue and treatment of poor that give $U_R(r) = 0.5$
and $U_P(p) = 0.5$

In the second assessment, Lottery C and its certain alternative D were judged to be equivalent:

- C** Best on both revenue and treating poor with probability 0.46
Worst on both revenue and treating poor with probability 0.54
- D** Worst on revenue and best on treatment of poor

- a Find values for k_R and k_P . Should the administrator consider these two attributes to be substitutes or complements? Why or why not?
 - b Comment on using an additive utility model in this situation. Would such a model seriously compromise the analysis of the decision?
- 16.8** Suppose you face an investment decision in which you must think about cash flows in two different years. Regard these two cash flows as two different attributes, and let X represent the cash flow in Year 1, and Y the cash flow in Year 2. The maximum cash flow you could receive in any year is \$20,000, and the minimum is \$5000. You have assessed your individual utility functions for X and Y , and have fitted exponential utility functions to them:

$$U_X(x) = 1.05 - 2.86 e^{-x/5000}$$

$$U_Y(y) = 1.29 - 2.12 e^{-y/10,000}$$

Furthermore, you have decided that utility independence holds, and so these individual utility functions for each cash flow are appropriate regardless of the amount of the other cash flow. You also have made the following assessments:

- You would be indifferent between a sure outcome of \$7500 each year for two years and a risky investment with a 50% chance at \$20,000 each year, and a 50% chance at \$5000 each year.
- You would be indifferent between getting (1) \$18,000 the first year and \$5000 the second, and (2) getting \$5000 the first year and \$20,000 the second.
- a Use these assessments to find the scaling constants k_X and k_Y . What does the value of $(1 - k_X - k_Y)$ imply about the cash flows of the different periods?
- b Use this utility function to choose between Alternatives A and B in Problem 15.23 (Figure 15.15).
- c Draw indifference curves for $U(x, y) = 0.25, 0.50,$ and 0.75 .

16.9 Refer to Problem 15.25. A decision maker who prefers Chemical B might be said to be sensitive to equity between the two groups. The eventual outcome with Chemical A is not equitable; one group is better off than the other. On the other hand, with Chemical B, both groups are treated the same. Let X and Y denote the expected increase in the number of deaths in Groups 1 and 2, respectively, and denote a decision maker's utility function as

$$U(x, y) = k_X U_X(x) + k_Y U_Y(y) + (1 - k_X - k_Y) U_X(x) U_Y(y)$$

What can you say about the value of $(1 - k_X - k_Y)$? Are X and Y complements or substitutes?

16.10 In Problems 13.12 and 14.13 you assessed individual utility functions for money (X) and computer reliability (Y). In this problem, we will use these assessed utility functions to put together a two-attribute utility function for use in a computer-purchase decision. (If you have not already worked Problems 13.12 or 14.13, do so before continuing. Even if you already have assessed these utility functions, review them and confirm that they are good models of your preferences. When you assessed them, did you consciously think about keeping all other important attributes at the same level?)

- a Are your preferences mutually utility independent? Your utility for money probably does not depend on your computer's level of reliability. But would you be less nervous about computer reliability if you had more money in the bank? Imagine that you have \$5000 in the bank. Now assess a certainty equivalent (in terms of computer downtime) for a 50-50 gamble between your computer being down for 20 days next year or not breaking down at all. Now imagine that you have \$30,000 in the bank, and reassess your certainty equivalent. Is it the same? Are your preferences for money and computer reliability mutually utility independent?
- b Regardless of your answer to part a, let us assume that your preferences for money and computer reliability are mutually utility independent. Now make the following assessments:

- i. Assess a certainty equivalent (in terms of both attributes: computer downtime and money) for a 50-50 gamble between the worst outcome (\$1000 in the bank and 50 days of computer downtime) and the best outcome (\$20,000 in the bank and no downtime).
 - ii. Imagine the outcome that is \$1000 and no downtime (worst level in money, and best level in computer reliability). Assess a dollar amount x so that the outcome x dollars and 50 days of downtime is equivalent to \$1000 and no downtime.
- c Using your individual assessed utility functions from Problems 13.12 and 14.13 and assessments i and ii from part b, calculate k_X and k_Y . Write out your two-attribute utility function.

(Note: This problem can be done without having fit a mathematical expression to your individual utility function. It can be done with a utility function expressed as a table or as a graph.)

- 16.11** Someday you probably will face a choice among job offers. Aside from the nature of the job itself, two attributes that are important for many people are salary and location. Some people prefer large cities, others prefer small towns. Some people do not have strong preferences about the size of the town in which they live; this would show up as a low weight for the population-size attribute in a multiattribute utility function.
- Assess a two-attribute utility function for salary (X) and population size (Y):
- a Determine whether your preferences for salary and town size are mutually utility independent.
 - b If your preferences display mutual utility independence, assess the two individual utility functions and the weights k_X and k_Y . Draw indifference curves for your assessed utility function. If your preferences do not display mutual utility independence, then you need to think about alternative approaches. The simplest is to assess several utility points as described at the beginning of this chapter and "eyeball" the indifference curves.
 - c What other attributes are important in a job decision? Would the two-attribute utility function you just assessed be useful as a first approximation if many of the other attributes were close in comparing two jobs?
- 16.12** Refer to Problem 15.17. For some of the attributes in your computer decision, you face uncertainty. The additive utility function assessed in Problem 15.17 essentially assumes that additive independence among attributes is reasonable.
- a Check some of your attributes with formal assessments for additive independence. Follow the example in the text in setting up the two lotteries. (One lottery is a 50-50 gamble between best and worst on both attributes. The other is a 50-50 gamble between (i) best on X , worst on Y ; and (ii) worst on X , best on Y)
 - b How could you extend this kind of assessment for additive independence to more than two attributes?
- 16.13** Show that when $n = 2$, the multiplicative utility function is equivalent to the two-attribute multilinear utility function. (*Hint:* This is an algebra problem. You must show that one utility function is a scaled version of the other. Start by solving for the scaling constant k in the multiplicative function when $n = 2$. Then substitute your expression for k into the multiplicative utility function and simplify.)

- 16.14** In many cases a group of people must make a decision that involves multiple objectives. In fact, difficult decisions usually are dealt with by committees composed of individuals who represent different interests. For example, imagine a lumber mill owner and an environmentalist on a committee trying to decide on national forest management policy. It might make sense for the committee to try to assess a multiattribute "group utility function." But assessment of the weights would be a problem because different individuals probably would want to weight the attributes differently. Can you give any advice to a committee working on a problem that might help it to arrive at a decision? How should the discussions be structured? Can sensitivity analysis help in such a situation; if so, how?
- 16.15** In making a land-use policy decision, a planner had to consider three objectives: the development of the best economic mixture of industrial and residential uses, the preservation of sensitive environmental areas, and satisfying the largest industrial firm in the community. Attributes to measure achievement along these three objectives were developed. Let k_{econ} represent the scaling constant for the economy, k_{env} the scaling constant for the environmental concerns, and k_{firm} the scaling constant for satisfying the firm. After careful thought, the planner assessed $k_{\text{econ}} = 0.36$, $k_{\text{env}} = 0.25$, and $k_{\text{firm}} = 0.14$. Use Equation 16.3 for the multiplicative utility model to verify that the scaling constant k is approximately 1.303.

C A S E S T U D I E S

A MINING-INVESTMENT DECISION

A major U.S. mining firm faced a difficult capital-investment decision. The firm had the opportunity to bid on two separate parcels of land that had valuable ore deposits. The project involved planning, exploration, and eventually production of minerals. The firm had to decide how much to bid, whether to bid alone or with a partner, and how to develop the site if the bid were successful. Overall, the company would have to commit approximately \$500 million to the project if it obtained the land.

Figure 16.8 shows a skeleton version of the decision-tree model for this decision. Note that one of the immediate alternatives is not to bid at all, but to stay with and develop the firm's own property. Some of the key uncertainties are whether the bid is successful, the success of a competing venture, capital-investment requirements, operating costs, and product price.

Figure 16.9 shows cumulative distribution functions for net present value (NPV) from four possible strategies. Strategy 25 — develop own property with partner — stochastically dominates all of the other strategies considered, and hence appears to be a serious candidate for the chosen alternative. The decision makers realized, however, that while they did want most to maximize the project's NPV, they also had another objective, the maximization of product output (PO). Because of this, a

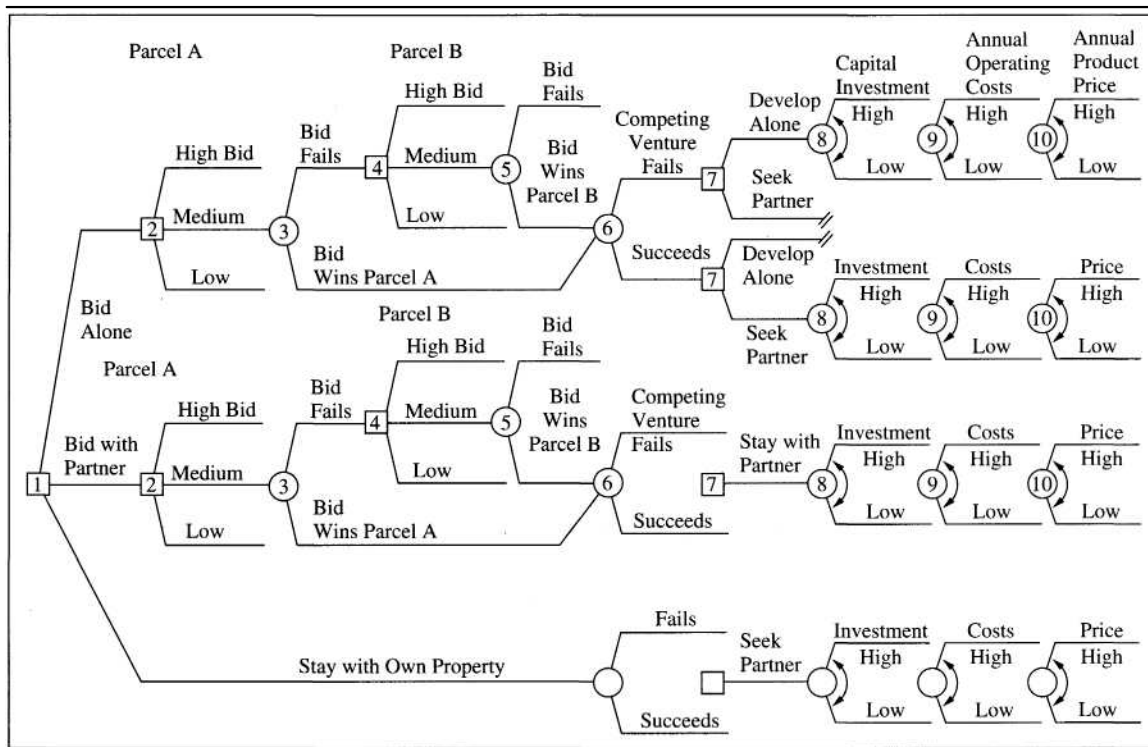


Figure 16.8 Skeleton decision tree for mining-investment decision.

Figure 16.9 Cumulative risk profiles for four strategies in mining-investment decision.




Table 16.4
Expected utilities (EU)
and certainty equivalents (CE) for strategies in mining-investment decision.

	Bid High Alone	Bid High with Partner	Bid Low with Partner	Develop Own Property with Partner	Do Nothing
EU	0374	0.577	0.564	0.567	0308
CE(NPV, in \$Million, PO = 50)	51	52	45	47	

two-attribute utility function for NPV and PO was constructed. The individual utility functions were assessed as exponential utility functions:

$$U_{PO}(po) = 1 - e^{-po/33.33}$$

$$U_{NPV}(npv) = 1 - e^{-(npv+100)/200}$$

The scaling constants also were assessed, yielding $k_{NPV} = 0.79$ and $k_{PO} = 0.16$. Using this two-attribute model, expected utilities and certainty equivalents—a certain (NPV, PO) pair, with PO set to a specific level — were calculated. Table 16.4 shows the results for some of the strategies.

Questions

- 1 Which of the alternatives should be chosen? Why? Discuss the apparent conflict between the stochastic dominance results (Figure 16.9) and the results from the utility model.
- 2 The values of the scaling constants suggest that NPV and PO are viewed by the firm as complements. Can you explain intuitively why these two attributes would be complements rather than substitutes?

Source: A. C. Hax and K. M. Wiig (1977) "The Use of Decision Analysis in a Capital Investment Problem." In D. Bell, R. L. Keeney, and H. Raiffa (eds.) *Conflicting Objectives in Decisions*, pp. 277-297. New York: Wiley. Figures 16.8 and 16.9 reprinted by permission.

REFERENCES

Chapter 16 represents only the tip of the iceberg when it comes to general multiattribute utility modeling. If you are interested in reading more, the standard reference for multiattribute utility theory is Keeney and Raiffa (1976). Keeney (1980) covers most of the same material and is a little easier to read. Bunn (1984) provides a somewhat less technical summary of many of the techniques, and von Winterfeldt and Edwards (1986) have extensive discussion with an emphasis on behavioral issues. Further references and examples of applications can be obtained from all of these texts.

Bunn, D. (1984) *Applied Decision Analysis*. New York: McGraw-Hill.

Keeney, R. (1980) *Siting Energy Facilities*. New York: Academic Press.

Keeney, R. (1992) *Value-Focused Thinking: A Path to Creative Decisionmaking*. Cambridge, MA: Harvard University Press.

Keeney, R., and T. McDaniels (1992) "Value-Focused Thinking about Strategic Decisions at BC Hydro." *Interfaces*, 22, 94-109.

Keeney, R., and H. Raiffa (1976) *Decisions with Multiple Objectives*. New York: Wiley.

Raiffa, H. (1982) *The Art and Science of Negotiation*. Cambridge, MA: Harvard University Press.

von Winterfeldt, D., and W. Edwards (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press.

EPILOGUE

The Mining-Investment Decision. Hax and Wiig (1977) reported that the decision was to adopt Strategy 2 — bid high with a partner. Just looking at NPV left the decision maker feeling uneasy. But with the product volume included in the two-attribute analysis, the decision maker was satisfied that the important objectives had been considered, so it was easy to adopt the strategy with the highest EU.

CHAPTER 17

Conclusion and Further Reading

We all have to face hard decisions from time to time. Sometimes we must make difficult personal decisions such as how to care for an elderly loved one or which one of several job offers to accept. Many policy decisions for corporations or governmental agencies are hard to make. In fact, as time goes by, it becomes increasingly clear that as a society we must grapple with some particularly thorny problems, such as competitiveness in the world marketplace, the risks associated with new technologies, and trade-offs between short-term economic benefits and long-term environmental stability.

The argument all along has been that decision analysis can help with such hard decisions. The cycle of structuring the decision, modeling uncertainty and preferences, analyzing and then performing sensitivity analysis can lead a decision maker systematically through the issues that make the decision complicated and toward a requisite decision model, one that captures all of the essential elements of the problem. The objective is to arrive at a decision model that explicates the complex parts in a way that the decision maker can choose from the alternatives with insight and understanding.

At the same time that decision analysis provides a framework for tackling difficult decisions, it also furnishes the decision maker with a complete tool kit for constructing the necessary models of uncertainty and preferences. We have spent much time in considering probability and how to use it to model the uncertainty that a decision maker faces. Subjective assessment, theoretical models, and the use of data and simulation are all tools in the decision analyst's kit. We also discussed the tools available for modeling preferences. We considered in some depth the fundamental trade-off between risk and return. Finally, the last two chapters focused on the modeling of preferences when the decision maker must try to satisfy conflicting objectives.

Throughout the book the view of decision making has been optimistic. We all are subject to human foibles, but a person interested in making better decisions can use the principles and tools that we have discussed in order to do a better job. In day-to-day decisions, it may be simply a matter of thinking in terms of an informally decomposed problem: What is the nature of the problem? What are the objectives? Are there trade-offs to make? What uncertainties are there? Is it a risky situation? More complicated situations may warrant considerable effort and careful use of the modeling tools.

Finally, in the process of reading the text and working through the problems, you may have learned something about yourself. You may have learned how you personally feel about uncertainty in your life, how you deal with risky situations, or what kinds of trade-offs are important to you. If you have learned something about the tools of decision analysis, gained an understanding of what it means to build a model of a decision problem, and learned a little about your own decision-making personality, then your work has been worthwhile. If you feel that you are more prepared to face some of the complicated decisions that we all must face as we move into the twenty-first century, then the goal of this text has been achieved.

A Decision-Analysis Reading List

Where should you go from here? The references at the end of each chapter can lead you to more information on specific topics. You undoubtedly noticed that many of the references reappeared several times. Several textbooks cover decision analysis at a variety of levels and from different perspectives. Here are some favorites:

- Max Bazerman and Margaret Neale (1992) *Negotiating Rationally*. New York: Free Press. An introduction to behavioral aspects of negotiation, written at a slightly lower level than the companion volume by Neale and Bazerman. Taking a decision-analysis perspective, the authors also emphasize prescriptive advice for negotiators.
- Derek Bunn (1984) *Applied Decision Analysis*. New York: McGraw-Hill. An excellent book written at about the same level as this one, although more theoretical and somewhat more terse. Excellent problems.
- Robyn Dawes (1988) *Rational Choice in an Uncertain World*. San Diego, CA: Harcourt Brace. An easy-to-read introduction to the behavioral issues in decision analysis.
- Simon French (1986) *Decision Theory: An Introduction to the Mathematics of Rationality*. London: Wiley. If you liked the chapter on utility axioms, you would love this book. The text covers a lot of new material, including consensus, group decisions, and non-Bayesian approaches. The problems are good, tending to be technical and theoretical rather than applied.
- Robin M. Hogarth (1987) *Judgment and Choice*, 2nd ed. New York: Wiley. An excellent introduction to behavioral decision theory, this is decision analysis

from a psychological perspective. This book covers a broad range of topics and is very easy to read.

- Ronald A. Howard and James Matheson (eds.) (1983) *The Principles and Applications of Decision Analysis* (2 volumes). Palo Alto, CA: Strategic Decisions Group. Since the early 1960s, Ron Howard has been practicing decision analysis and teaching the principles to students in the Engineering-Economic Systems Department at Stanford University. This two-volume set contains many papers presenting the principles and techniques that make up the "Stanford School" of decision analysis.
- Daniel Kahneman, Paul Slovic, and Amos Tversky (1982) *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press. A reader filled with classic articles on behavioral decision theory.
- Ralph Keeney (1992) *Value-Focused Thinking: A Path to Creative Decision Making*. Cambridge, MA: Harvard University Press. As you know by now, this book provides all the details on understanding one's objectives and using them as a basis for decision-analysis models and improved decision making.
- Ralph Keeney and Howard Raiffa (1976) *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: Wiley. Reprinted in 1993 by Cambridge University Press. This is the standard reference for multiattribute utility theory, although it also is good for decision analysis in general. Many applications are described. Much of the material is highly technical, although the mathematics are not difficult. Gems of insight and explanation are scattered through the technical material. Unfortunately, no problems are included.
- Dennis V. Lindley (1985) *Making Decisions*, 2nd ed. New York: Wiley. A classic by a founder in the field. Professor Lindley explains difficult concepts well. The problems tend to be somewhat abstract.
- M. Granger Morgan and Max Henrion (1990) *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge: Cambridge University Press. The authors provide an in-depth treatment of the use of uncertainty for risk analysis. Especially good on the elicitation and use of expert judgment.
- Margaret Neale and Max Bazerman (1991) *Cognition and Rationality in Negotiation*. New York: Free Press. An excellent summary of behavioral issues in negotiation and group decision making. Many of the results are closely related to parallel results in individual decision making.
- Scott Pious (1993) *The Psychology of Judgment and Decision Making*. New York: McGraw-Hill. The best up-to-date summary of behavioral decision theory.
- Howard Raiffa (1968) *Decision Analysis*. Reading, MA: Addison-Wesley. Professor Raiffa also is a founder of decision analysis, and like Lindley, he explains the material well. The problems tend to be abstract. The text covers the basics (and then some) and still is worthwhile after almost 30 years.
- Howard Raiffa (1982) *The Art and Science of Negotiation*. Cambridge, MA: Belknap Press. Raiffa discusses negotiations from a decision-theoretic point of

view. A well-written (and easy-to-read!) book that provides deep insights and top-flight analytical tools.

- Detlof von Winterfeldt and Ward Edwards (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press. An up-to-date and in-depth treatment of decision analysis from a behavioral perspective. Professor Edwards has been involved in decision analysis since its beginnings, and his history (Chapter 14) is an eye-opener. The authors have a strong slant toward applications and behavioral research that provides evidence on the applicability of decision-analysis tools. Like Keeney and Raiffa, there are no problems.
- Robert L. Winkler (1972) *Introduction to Bayesian Inference and Decision*. New York: Holt. An excellent introduction to decision theory. Professor Winkler is especially interested in Bayesian models of information, and the book is slanted more toward inference and statistics than toward applied decision analysis.

Software Sources

@RISK	Palisade Corporation 31 Decker Rd. Newfield, NY 14867 (800) 432-7475 (607) 277-8000	Best Fit	Palisade Corporation 31 Decker Rd. Newfield, NY 14867 (800) 432-4757 (607) 277-8000
Crystal Ball	Decisioneering, Inc. 2530 S. Parker Road, Suite 220 Aurora, CO 80014 (800) 289-2550	DATA	TreeAge Software, Inc. P.O. Box 990207 Boston, MA 02199-0207 (617) 536-2128
DPL	Applied Decision Analysis 2710 Sand Hill Rd. Menlo Park, CA 94025 (415)926-9251	Student editions of DPL available from: Duxbury Press 10 Davis Dr. Belmont, CA 94002 (800) 423-0563	
HIVIEW	KRYSALIS LTD. 28 Derwent Dr. Maidenhead Berks. SL6 6LB England 44 (0) 628-36861	HIPRE 3+	Santa Monica Software 30033 Harvester Rd. Malibu, CA 90265 (310)395-7635
Logical			
Decisions	Logical Decisions 1014 Wood Lily Dr. Golden, CO 80401 (303) 526-7536 (800) 355-6442		

Appendixes

A	Binomial Distribution: Individual Probabilities
B	Binomial Distribution: Cumulative Probabilities
C	Poisson Distribution: Individual Probabilities
D	Poisson Distribution: Cumulative Probabilities
E	Normal Distribution: Cumulative Probabilities
F	Beta Distribution: Cumulative Probabilities

Answers to Selected Numerical Exercises

The following are answers to selected numerical *exercises* only. Solutions and answers to *all* exercises, questions, and problems can be found in the instructor's manual.

Chapter 2: **2.9.** \$158.78. **2.10.** \$2003.90, 19.2%. **2.11a.** \$23.71, -\$28.44. **2.11b.** 14.5%. **2.12a.** -\$25.60. **2.12b.** \$25.60. **2.12c.** -\$0.08.

Chapter 4: **4.4.** 3.0, 3.2. **4.6.** 8.1, 3.2. **4.8.** 5.08, 4.5.

Chapter 7: **7.3.** 0.12, 0.29, 0.41, 0.65, 0.35, 0.29, 0.18, 0.17. **7.4.** 0.94. **7.6.** j, c, c, c, j, c, c, c, c, j. **7.7.** $\sigma_B = \$3.72$ million. $\sigma_c = 0$. **7.8.** 0.58, 0.34, 0.75, 0.42, 0.58, 0.66, 0.34, 0.25, 0.75. **7.9.** 0.90, 0.61, 0.61, 0.39, 0.61, 0.10, 0.90, 0.10, 0.90. **7.10.** 2.73, 0.60. **7.14.** 0.70, 0.98, 0.21, 0.79, 0.970, 0.030, 0.603, 0.397. **7.15.** 0.56. **7.16a.** 2.65, 0.728, 0.853. **7.16b.** 26.40, 2255.04, 47.49. **7.16c.** 0.632, 0.233, 0.482. **7.19a.** \$7000, 12250. **7.19b.** \$42,895, 41871. **7.20.** \$1594, 430,575. **7.22.** 0.24, 0.795, 0.001. **7.24.** 0.779, 0.041, 0.656, 0.134.

Chapter 8: **8.13.** (1) 39 years; (2) 4187 miles; (3) 13 countries; (4) 39 books; (5) 2160 miles; (6) 390,000 pounds; (7) 1756; (8) 645 days; (9) 5959 miles; (10) 36,198 feet.

Chapter 9: **9.1.** 0.915. **9.2.** 0.0918, 0.0918. **9.3.** 0.027, 0.706, 0.001. **9.4.** 0.33. **9.5a.** 0.037, 0.971, 0.191, 0.157; 0.205, 0.514, 0.317, 0.108. **9.5b.** 0.180, 0.594, 0.099, 0.068; 0.336, 0.130, 0.190, 0.310. **9.5c.** 0.993, 0.471, 0.0; 0.368, 0.002, 0.368. **9.5d.** 0.6915, 0.2604, 0.0062, 0.0; 0.9525, 0.1056, 0.5363, 0.6915. **9.5e.** 0.39, 0.01, 0.36; 0.31, 0.30, 0. **9.6.** -1.645, 0; 0.675, 1.28.

9.7. $\mu = 200$, $\sigma = 111.11$. **9.8.** 0.2854. **9.9.** 5.4. **9.10.** 1.743.

Chapter 10: **10.4.** 275,400. **10.6.** 9300, 9850.

Chapter 12: **12.2.** 1.2. **12.4.** 3.04, 1.20, 3.22.

Chapter 13: **13.6.** \$236. **13.7a.** 0.56, 0.48, 0, -1.81. **13.7b.** 0.052. **13.7c.** \$64.52, \$270.98. **13.7d.** \$106.18. **13.7e.** \$2363.

Chapter 15: **15.7.** 1, 0.4, 0. **15.9.** 0.41, 0.59. **15.11.** 0.25, 0.34, 0.41.

Chapter 16: **16.6.** 0.48, 0.67.