Myerson_FM.qxd 09-09-2004 15:56 Page i

# www.duxbury.com

www.duxbury.com is the World Wide Web site for Thomson Learning and is your direct source to dozens of online resources.

At www.duxbury.com you can find out about supplements, demonstration software, and student resources. You can also send email to many of our authors and preview new publications and exciting new technologies.

**www.duxbury.com**
Changing the way the world learns®

Myerson_FM.qxd 09-09-2004 15:56 Page ii

## Duxbury Titles of Related Interest

Albright, **VBA for Modelers**
Albright & Winston, **Spreadsheet Modeling and Applications**
Albright, Winston, & Zappe, **Data Analysis and Decision Making**
Berger & Maurer, **Experimental Design**
Bowerman & O'Connell, **Linear Statistical Models**
Bowerman, O'Connell, & Koehler, **Forecasting, Time Series and Regression**
Clemen & Reilly, **Making Hard Decisions with Decision Tools**
Davis, **Business Research for Decision Making**
Devore & Farnum, **Applied Statistics for Engineers and Scientists**
Dielman, **Applied Regression Analysis**
Durrett, **Probability: Theory and Examples**
Fourer, Gay, and Kernighan, **AMPL: A Modeling Language for Mathematical Programming**
Hildebrand, Ott, & Gray, **Basic Statistical Ideas for Managers**
Hoerl & Snee, **Statistical Thinking: Improving Business Performance**
Keller & Warrack, **Statistics for Management and Economics**
Kirkwood, **Strategic Decision Making**
Kuehl, **Design of Experiments: Statistical Principles for Research Design and Analysis**
Lapin & Whisler, **Quantitative Decision Making**
Lattin, Carroll, & Green, **Analyzing Multivariate Data**
Lohr, **Sampling: Design and Analysis**
Lunneborg, **Data Analysis by Resampling**
Minh, **Applied Probability Models**
Neuwirth & Arganbright, **Mathematical Modeling with Microsoft®Excel**
SAS Institute, Inc., **JMP-IN: Statistical Discovery Software**
Savage, **Decision Making with Insight**
Schaeffer, Mendenhall, & Ott, **Elementary Survey Sampling**
Seila, Ceric, & Tadikamalla, **Applied Simulation Modeling**
Shapiro, **Modeling the Supply Chain**
Stewart, **Introduction to Applied Econometrics**
Weiers, **Introduction to Business Statistics**
Winston, **Introduction to Probability Models**
Winston, **Operations Research: Applications and Algorithms**
Winston & Albright, **Practical Management Science**
Winston & Venkataramanan, **Introduction to Mathematical Programming**

To order copies, contact your local bookstore or call 1-800-354-9706.
For more information, go to www.duxbury.com

Myerson_FM.qxd 09-09-2004 15:56 Page iii

DUXBURY

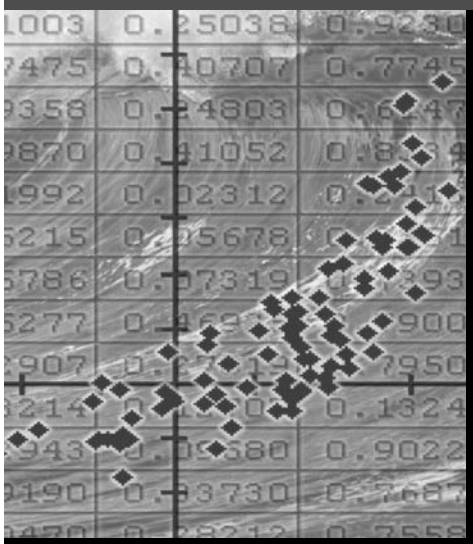

# Probability Models for Economic Decisions

**Roger B. Myerson**

University of Chicago

THOMSON
BROOKS/COLE

Australia ■ Canada ■ Mexico ■ Singapore ■ Spain
United Kingdom ■ United States

Myerson_FM.qxd 09-09-2004 15:56 Page iv

Publisher: *Curt Hinrichs*
Assistant Editor: *Ann Day*
Editorial Assistant: *Fiona Chong*
Technology Project Manager: *Burke Taft*
Marketing Manager: *Tom Ziolkowski*
Advertising Project Manager: *Nathaniel Bergson-Michelson*
Project Manager, Editorial Production: *Kelsey McGee*
Print/Media Buyer: *Doreen Suruki*
Permissions Editor: *Stephanie Lee*

© 2005 Thomson Brooks/Cole, a part of The Thomson Corporation. Thomson, the Star logo, and Brooks/Cole are trademarks used herein under license.

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, Web distribution, or information storage and retrieval systems, or in any other manner—without the written permission of the publisher.

Printed in the United States of America
1 2 3 4 5 6 7 08 07 06 05 04

For more information about our products, contact us at:
**Thomson Learning Academic Resource Center**
**1-800-423-0563**

For permission to use material from this text or product, submit a request online at
*http://www.thomsonrights.com.*
Any additional questions about permissions can be submitted by email to thomsonrights@Thomson.com.

Library of Congress Control Number: 2004109025
ISBN 0-534-42381-7

Production Service: *Matrix Productions/Aaron Downey*
Copy Editor: *Pam Rockwell*
Compositor: *International Typesetting and Composition*
Cover and Interior Designer: *Kaelin Chappell*
Cover and Chapter Opener Image: *Rick Doyle/CORBIS*
Cover Printing: *Coral Graphic Services*
Printing and Binding: *Malloy Incorporated*

**Brooks/Cole Thomson Learning**
**10 Davis Drive**
**Belmont, CA 94002**
**USA**

**Asia**
Thomson Learning
5 Shenton Way #01-01
UIC Building
Singapore 068808

**Australia/New Zealand**
Thomson Learning
102 Dodds Street
Southbank, Victoria 3006
Australia

**Canada**
Nelson
1120 Birchmount Road
Toronto, Ontario M1K 5G4
Canada

**Europe/Middle East/Africa**
Thomson Learning
High Holborn House
50/51 Bedford Row
London WC1R 4LR
United Kingdom

Myerson_FM.qxd 09-09-2004 15:56 Page v

To my parents, who have continued learning about the world and working to improve it, even in their 60th year of marriage.

Myerson_FM.qxd 09-09-2004 15:56 Page vi

Myerson_FM.qxd 09-09-2004 15:56 Page vii

# Brief Contents


1 **Simulation and Conditional Probability 1**

2 **Discrete Random Variables 47**

3 **Utility Theory with Constant Risk Tolerance 81**

4 **Continuous Random Variables 107**

5 **Correlation and Multivariate Normal Random Variables 151**

6 **Conditional Expectation 199**

7 **Optimization of Decision Variables 231**

8 **Risk Sharing and Finance 281**

9 **Dynamic Models of Growth and Arrivals 325**

Myerson_FM.qxd 09-09-2004 15:56 Page viii

Myerson_FM.qxd 09-09-2004 15:56 Page ix

# About the Author

**Roger B. Myerson** is the W. C. Norby Professor of Economics at the University of Chicago. He previously taught at the Kellogg School of Management at Northwestern University (1976–2001). His teaching interests include decision analysis, probability modeling, game theory, mathematical optimization theory, social choice and formal political theory, and economics of information. Dr. Myerson has received a number of professional awards, including Guggenheim Fellow (1983–1984), Sloan Foundation Fellow (1984–1986), Fellow of the Econometric Society (elected 1983), Fellow of the American Academy of Arts and Sciences (elected 1993), and an honorary doctorate from the University of Basel (2002). His research interests include game theory, economics of information, and analysis of voting systems.

Myerson_FM.qxd 09-09-2004 15:56 Page x

Myerson_FM.qxd 09-09-2004 15:56 Page xi

# Contents


**1 Simulation and Conditional Probability 1**

**1.0** Getting Started with Simtools in Excel 2
**1.1** How to Toss Coins in a Spreadsheet 3
**1.2** A Simulation Model of 20 Sales Calls 7
**1.3** Analysis Using Excel's Data-Table Command 15
**1.4** Conditional Independence 18
**1.5** A Continuous Random Skill Variable from a Triangular Distribution 19
**1.6** Probability Trees and Bayes' Rule 25
**1.7** Advanced Spreadsheet Techniques: Constructing a Table with Multiple Inputs 37
**1.8** Using Models 40
**1.9** Summary 42
Exercises 43

**2 Discrete Random Variables 47**

**2.1** Unknown Quantities in Decisions Under Uncertainty 49
**2.2** Charting a Probability Distribution 50
**2.3** Simulating Discrete Random Variables 53
**2.4** Expected Value and Standard Deviation 57

Myerson_FM.qxd 09-09-2004 15:56 Page xii

**2.5** Estimates from Sample Data 62
**2.6** Accuracy of Sample Estimates 66
**2.7** Decision Criteria 71
**2.8** Multiple Random Variables 75
**2.9** Summary 77
Exercises 78

## 3 Utility Theory with Constant Risk Tolerance 81

**3.1** Taking Account of Risk Aversion: Utility Analysis with Probabilities 82
**3.2** Utility Analysis from Simulation Data 92
**3.3** The More General Assumption of Linear Risk Tolerance 94
**3.4** Advanced Technical Note on Utility Theory 96
**3.5** Advanced Technical Note on Constant Risk Tolerance 101
**3.6** Summary 105
Exercises 105

## 4 Continuous Random Variables 107

**4.1** Normal Distributions 108
**4.2** EXP and LN 113
**4.3** Lognormal Distributions 116
**4.4** Generalized Lognormal Distributions 122
**4.5** Subjective Probability Assessment 124
**4.6** A Decision Problem with Discrete and Continuous Unknowns 131
**4.7** Certainty Equivalents of Normal Lotteries 135
**4.8** Other Probability Distributions 137
**4.9** Summary 145
Exercises 149

## 5 Correlation and Multivariate Normal Random Variables 151

**5.1** Joint Distributions of Discrete Random Variables 152
**5.2** Covariance and Correlation 156

Myerson_FM.qxd 09-09-2004 15:56 Page xiii

**5.3** Linear Functions of Several Random Variables 158
**5.4** Estimating Correlations from Data 162
**5.5** Making Multivariate Normal Random Variables with CORAND and NORMINV 166
**5.6** Portfolio Analysis with Multivariate Normal Asset Returns 171
**5.7** Excel Solver and Efficient Portfolio Design 176
**5.8** Subjective Assessment of Correlations 182
**5.9** Using CORAND with Non-Normal Random Variables 186
**5.10** More About Linear Functions of Random Variables 190
**5.11** Summary 193
Exercises 195

**6 Conditional Expectation 199**

**6.1** Dependence Among Random Variables 199
**6.2** Estimating Conditional Expectations and Standard Deviations 203
**6.3** The Expected-Posterior Law in a Discrete Example 206
**6.4** Backwards Analysis of Conditional Expectations in Tree Diagrams 211
**6.5** Conditional Expectation Relationships and Correlation 214
**6.6** Uncertainty About a Probability 216
**6.7** Linear Regression Models 220
**6.8** Regression Analysis and Least-Squared Errors 224
**6.9** Summary 227
Exercises 229

**7 Optimization of Decision Variables 231**

**7.1** General Techniques for Using Simulation in Decision Analysis 231
**7.2** Strategic Use of Information 241
**7.3** Decision Trees 246
**7.4** A Simple Bidding Problem 251
**7.5** The Winner's Curse 255
**7.6** Analyzing Competitive Behavior 263
**7.7** Summary 271
Exercises 273

Myerson_FM.qxd 09-09-2004 15:56 Page xiv

**8 Risk Sharing and Finance 281**

**8.1** Optimal Risk Sharing in a Partnership of Individuals with Constant Risk Tolerance 282
**8.2** Optimality of Linear Rules in the Larger Class of Nonlinear Sharing Rules 291
**8.3** Risk Sharing Subject to Moral-Hazard Incentive Constraints 295
**8.4** Piecewise-Linear Sharing Rules with Moral Hazard 301
**8.5** Corporate Decision Making and Asset Pricing in the Stock Market 305
**8.6** Fundamental Ideas of Arbitrage Pricing Theory 317
**8.7** Summary 322
Exercises 323

**9 Dynamic Models of Growth and Arrivals 325**

**9.1** Net Present Value 326
**9.2** Forecasting Models 328
**9.3** Forecasting Example: Goeing Case 332
**9.4** Brownian-Motion Growth Models 339
**9.5** Log-Optimal Investment Strategies 343
**9.6** Exponential Arrival Models 349
**9.7** Queuing Models 356
**9.8** A Simple Inventory Model 362
**9.9** Project Length and Critical Tasks 366
**9.10** Summary 370
Exercises 372

**Appendix: Excel Add-Ins for Use with This Book** 377

**Bibliographic Note** 387

**Index** 389

Myerson_FM.qxd 09-09-2004 15:56 Page xv

# List of Figures

**1.1** Simulation of coin tossing in a spreadsheet 5
**1.2** Coin tossing with adjustable probabilities 6
**1.3** Simple model of independent sales calls 8
**1.4** Model of 20 sales calls with uncertainty about salesperson's skill 10
**1.5** Simulation table for model of 20 sales calls 12
**1.6** Simulation data and analysis 14
**1.7** Data table of results for different numbers of sales 16
**1.8** Frequencies of sales in simulation data (total and by skill level) 18
**1.9** Sales model with skill level from a Triangular distribution 22
**1.10** Simulation model of oil exploration example 26
**1.11** Probability tree for oil exploration example 28
**1.12** Alternative tree and model for oil exploration example 32
**1.13** Binomial probability computations for salesperson example 35
**1.14** Computing frequencies of different outcomes in oil exploration example 37
**2.1** Discrete probability distribution for the number of entrants (K) 51
**2.2** Inverse cumulative probability curve for number of entrants (K) 52
**2.3** Simulation with a discrete probability distribution (four equivalent formulas) 54
**2.4** Making a simple simulation model of competitors and profit 57
**2.5** Expected values and standard deviations of discrete random variables 59

Myerson_FM.qxd 09-09-2004 15:56 Page xvi

**2.6** Estimating an expected value and standard deviation from simulation data 63

**2.7** Estimate of an inverse cumulative distribution from simulation data 66

**2.8** A spreadsheet for studying the properties of a sample mean 67

**2.9** Analysis of Superior Semiconductor case 73

**2.10** Cumulative risk profile from simulation data 74

**2.11** A simulation model with three random variables affecting profit 76

**3.1** Utility function with constant risk tolerance for Superior Semiconductor (B) 86

**3.2** Analysis of Superior Semiconductor (B) with constant risk tolerance 89

**3.3** Estimating local risk tolerances, with constant or linear risk tolerance 96

**3.4** A spreadsheet illustrating the derivation of utility theory 98

**3.5** Derivation of exponential utility from constant risk tolerance 104

**4.1** The Normal distribution 110

**4.2** Compound interest and growth rate calculations with EXP and LN 114

**4.3** The multiplication and addition properties of EXP and LN 115

**4.4** Repeated investment strategies yielding payoffs that are approximately Normal or Lognormal 118

**4.5** Subjectively assessed quartiles and a Generalized-Lognormal distribution 130

**4.6** Simulation analysis of the Superior Semiconductor (C) Case 133

**4.7** Certainty equivalent of a Normal lottery 136

**4.8** Comparison of Lognormal, Gamma, and Normal distributions 142

**5.1** Joint probability distribution of two random variables 154

**5.2** Correlation of two discrete random variables and portfolio analysis 155

**5.3** Price data and annual growth ratios in the 1980s for six mutual funds 162

**5.4** Computing the correlation of Funds 1 and 2 163

**5.5** Making a model of the six funds' annual growth ratios 165

**5.6** Extending historical data by simulation 169

**5.7** Simulations with independent funds 170

Myerson_FM.qxd 09-09-2004 15:56 Page xvii

**5.8** Portfolio calculations for investments in the six funds 173
**5.9** A simple example of how Solver can fail 181
**5.10** Computing conditional medians for subjective correlation assessment 183
**5.11** Non-Normal random variables and normalized rank correlations 187
**5.12** Making Multivariate Normal random variables 191
**6.1** Conditional expected skills, given sales results, from a Beta prior 202
**6.2** Prior and posterior inverse cumulative distributions for salesperson's skill 206
**6.3** Conditional probability and conditional expectation in a discrete example 207
**6.4** Backwards analysis of conditional expectations in a tree diagram 212
**6.5** An example of dependence among random variables with correlation 0 215
**6.6** Modeling uncertainty about a probability 218
**6.7** Estimating a linear regression relationship among random variables 221
**6.8** Minimizing squared errors of linear estimates in regression analysis 225
**7.1** Simulation analysis of one production quantity in the Scotia case 234
**7.2** Simulation analysis of four alternative production quantities in the Scotia case 237
**7.3** A spreadsheet for finding the optimum among all possible production quantities for Scotia 239
**7.4** Strategic value of information about weather in the Scotia case 242
**7.5** Conditional expected profits given weather, for Scotia's optimal strategies 245
**7.6** Tree diagram for Scotia's decision without forecast information 246
**7.7** Tree diagram for Scotia's decision with forecast information 249
**7.8** Model of bidding to sell in the Bates case, part A (no winner's curse) 253
**7.9** Finding optimal bids for the Bates case, part A (no winner's curse) 255
**7.10** Bidding to buy a "lemon" 257

Myerson_FM.qxd 09-09-2004 15:56 Page xviii

**7.11** Bidding model for the Bates case (part B) with winner's curse 260

**7.12** Inverse cumulative distribution for the opposing bid-to-beat in Bates (B) 262

**7.13** Finding optimal bids for the Bates case (part B) with winner's curse 263

**7.14** Alpha's bidding for Valdez when others bid 80% of their estimated values 265

**7.15** Alphas' bidding for Valdez when others bid 60% of their estimated values 268

**7.16** Game model of bidding for Valdez 270

**8.1** Sharing a Normal gamble 285

**8.2** Optimal risk sharing among three partners with constant risk tolerance 290

**8.3** A spreadsheet to evaluate nonlinear sharing rules 292

**8.4** Optimal risk sharing in a discrete example 294

**8.5** Optimal risk sharing with moral hazard, in a discrete example 296

**8.6** Optimal linear incentive plan for an agent with moral hazard 299

**8.7** Optimal piecewise-linear incentive plan for an agent with moral hazard 302

**8.8** Comparing investments that differ in correlation with the market portfolio 307

**8.9** Computing an asset's value in a market with constant-risk-tolerant investors 312

**8.10** A simple example of arbitrage pricing theory 319

**8.11** Solver Sensitivity Report for the problem in Figure 8.10 320

**9.1** Net present value of a future income stream 327

**9.2** Forecasting model for two stock funds 329

**9.3** *Global Sports* case 331

**9.4** Goeing case, with Forecast Model 1 (extrapolating income growth only) 333

**9.5** Goeing case, Forecast Model 2 (extrapolating sales growth and income/sales) 335

**9.6** Goeing case, Forecast Model 3 (sales growth, autoregressive income/sales) 337

**9.7** Historical data on annual end-of-year prices for a stock 340

**9.8** Forecast model of monthly stock prices and two derivative assets 342

Myerson_FM.qxd 09-09-2004 15:56 Page xix

**9.9** Comparison of four investment strategies for a repeated investment process 345

**9.10** Computing the log-optimal strategy for the repeated investment process 347

**9.11** No-memory property of Exponential waiting times 350

**9.12** Poisson and Gamma random variables from Exponential waiting times 351

**9.13** Approximating a Poisson random variable by a Binomial random variable 354

**9.14** Arrivals of different types, with different Exponential waiting times 355

**9.15** A queuing model simulated as an iterative process 357

**9.16** A simple inventory model 363

**9.17** A project-length model 368

Myerson_FM.qxd 09-09-2004 15:56 Page xx

Myerson_FM.qxd 09-09-2004 15:56 Page xxi

# Preface

This book is an introduction to the use of probability models for analyzing risks and economic decisions. Some prior study of probability may be desirable but is not assumed here, as the basic ideas of probability are introduced in the book itself. Throughout, the focus is on showing the reader how to use probability in complex, realistic situations. All the analytical work in this book is done in Microsoft Excel® spreadsheets, because the spreadsheet medium enables students to handle much more complex problems than they could handle using traditional ways of representing mathematical formulas. As a result of this emphasis on spreadsheet modeling, many readers may also learn sophisticated spreadsheet skills from this book. But the main goal is to make the practical power of probability analysis accessible to students at the undergraduate or MBA level.

The ultimate goal of any quantitative analysis course should be to teach students the art of making quantitative models that can give practical insights into real decision problems. But students usually have difficulty with complex models that involve more than two variables, and so quantitative methods are often taught only in the context of simple applications that lack any realistic complexity. The typical result is that even good students who show a mastery of mathematical concepts on their final exams often cannot see how to apply these concepts in the real world.

This problem was particularly frustrating to me during the 1980s as I did theoretical research in information economics while I taught probability to MBA students at the Northwestern University's Kellogg School of Management. New advances in information economics were teaching economic theorists that analysis of uncertainty has vital importance for understanding the competitive behavior of economic agents, and probability theory is the foundation for all analysis of uncertainty. But it seemed difficult or impossible to communicate this practical importance of probability analysis to my MBA students, even though their career goal was to become the competitively successful economic agents that we theorists

Myerson_FM.qxd 09-09-2004 15:56 Page xxii

were studying. So in the late 1980s, I began searching for radically new ways to teach probability analysis so that students, after a 10-week MBA-level course, should be able to apply probability analysis to really interesting economic problems and cases. An important motivation in this search was my perception that the powerful new spreadsheet programs had great unexplored potential to transform the way that quantitative analysis is taught. This book is the fulfillment of this long process of experimentation and pedagogical development.

It has now become commonplace for probability and statistics books to do calculations in Excel spreadsheets. In most cases, however, these books have simply taken the old material that professors used to teach on blackboards and moved it into a spreadsheet. Only rarely have authors asked harder questions about how the new spreadsheet medium should change the content of an introductory quantitative analysis course. But if we really want to maximize the practical value of the skills that a student can take away from the course, then such questions need to be asked. Analytical methods that seemed too difficult or complex for an audience of applied business and economics students when we worked on paper and blackboards now may be straightforward for the same students to master in spreadsheets. On the other hand, there is less need for students to memorize basic computational formulas now that computer programs with built-in functions and help screens are so universally available.

The most important advantage of teaching spreadsheets is that they make multivariable models much easier for students to visually understand. Most students lose the big picture when they are asked to think about more than two variables that are represented by letters like $x$ and $\alpha$ on a blackboard, but the same students can visually understand a spreadsheet model with many variables that are represented as cells on a spreadsheet page. Working in spreadsheets brings down the barrier that prevents most students from becoming sophisticated quantitative modelers. So throughout this book, I have tried to show how the methods of probability analysis can be applied to examples of a realistic complexity.

The emphasis on Monte Carlo randomized simulation of probability models grew gradually in the development of this book. At first I emphasized other methods of computing probabilities and expected values, but I realized that other computational methods generally require special assumptions that limit the scope of applications that we can consider. Monte Carlo simulation is the most versatile general framework for modeling any kind of situation that involves uncertainty, and it has the advantage of letting us really see the uncertainty in our models, as our unknown quantities change at every recalculation. With Monte Carlo simulation, we have a general framework that can be used to illustrate and motivate any advanced concept in probability and decision analysis. We can even give new and effective illustrations of conventional statistical topics such as confidence intervals (see Section 2.6) and regression (see Sections 6.7 and 6.8).

Myerson_FM.qxd 09-09-2004 15:56 Page xxiii

As I developed this course, I simultaneously worked to develop my own add-in program (Simtools.xla) to extend Excel's capabilities for probability analysis. I did not use one of the commercially produced add-ins for statistical risk analysis (like Crystal Ball or @Risk) because I wanted to go farther in decision analysis and economic modeling with less alteration of Excel's structure. But I found that there was a synergy between the development of the course and the development of the software. Students would regularly ask me to add a function to simplify a difficult part of the course; and as the new software simplified some topics, it became possible to teach other important topics that had previously seemed too advanced. As a result of this process, there are a number of new functions for facilitating decision analysis that are currently unique to Simtools (see GENLINV, CORAND, and CE, for example) but may eventually be imitated in other simulation add-ins.

I have used this book with MBA and undergraduate students at Northwestern University and the University of Chicago. Most of my students had some prior exposure to probability theory, but they often had difficulty recalling it when it was applied in our course. So this book includes an introductory discussion of all the basic probability concepts that we use. Chapter 1 introduces the basic ideas of probability and conditional probability. Chapter 2 covers discrete random variables, expected values, and standard deviations. In a short digression from probability theory, Chapter 3 introduces risk aversion and utility analysis with constant risk tolerance. Then Chapter 4 introduces continuous random variables, with an emphasis on Normal and Lognormal probability distributions. Chapter 5 covers joint distributions and correlation, and Chapter 6 discusses conditional expectation and regression models.

The later chapters of the book are organized around types of models. Simple decision models with one decision variable are introduced in Chapter 7, with an emphasis on decisions about production quantities (the "newsboy" problem) and decisions about prices (bidding in auctions). Chapter 8 focuses on risk sharing in partnerships and financial markets. Finally, Chapter 9 introduces a wide range of dynamic models for forecasting how quantities may change over time.

There is more than enough material here for a full semester course. In a recent 10-week course where many students needed a review of basic probability, I covered Sections 1.1–1.4, all of Chapter 2, Sections 3.1–3.2, Sections 4.1–4.7, Sections 5.1–5.8, Sections 6.1–6.2 and 6.6, Sections 7.1–7.2 and 7.4–7.6, and Sections 8.1 and 8.3. With students who have a stronger background in probability, I would try to spend less time in Chapters 2 and 4 and devote more time to Chapters 8 and 9. The heart of this book is in Chapter 5 (for relationships among random variables) and Chapter 7 (for decision variables), and these two chapters should get careful treatment in any course. I have also tried to make this book accessible for self-study or professional enrichment outside of the classroom. When

Myerson_FM.qxd 09-09-2004 15:56 Page xxiv

you open up the spreadsheet file for each chapter, you have everything that I would show you in my classroom.

Acknowledgments must begin with Howard Raiffa, who was my teacher and whose impact may be found throughout this book. Sid Deshmukh and Sam Savage have given me good colleagial suggestions and encouragement in this project over many years. It was Donald Jacobs who first told me to use more spreadsheets in my teaching, and Robert Weber first showed me how to use them. Last but not least is my debt to Rebecca Myerson, who carefully read every chapter of this book to find errors in both English and Excel. At the points where my MBA students had typically expressed some confusion, she kept asking questions until I found a way to rewrite it more clearly. I know that not every high school student would have been willing to proofread a long technical book by her father, and I am very grateful to Rebecca for helping to make this a better book.

*Chicago, Illinois*
*August 2004*

Myerson_FM.qxd 09-09-2004 15:56 Page xxv

# Acknowledgments

The author would like to express his gratitude to the reviewers of this book:

Luiz Brandao, University of Texas at Austin, McCombs School of Business
Sudhakar D. Deshmukh, Northwestern University, Kellogg
Sheldon Lin, University of Toronto
Erol Pekoz, Boston University
Sam Savage, Stanford University
Patrick Thompson, University of Florida

Myerson_FM.qxd 09-09-2004 15:56 Page xxvi

Myerson_FM.qxd 09-09-2004 15:56 Page xxvii

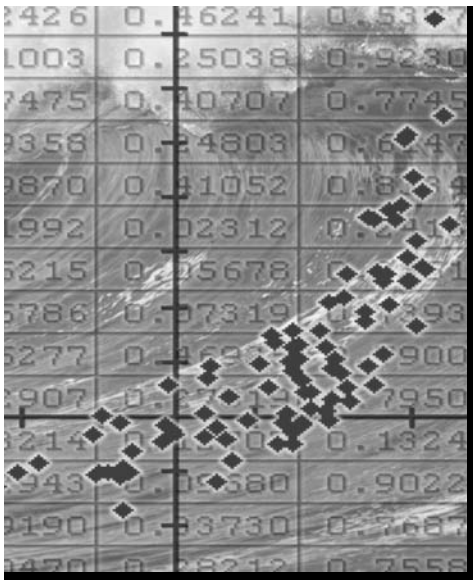

# Probability Models for Economic Decisions

Myerson_FM.qxd 09-09-2004 15:56 Page xxviii

Myerson_C01.qxd 08-09-2004 09:45 Page 1

# 1

# Simulation and Conditional Probability

The difficulties of making decisions under uncertainty are familiar to everyone. We all regularly have to make decisions where we lack important information that could significantly affect the outcomes. Decision analysis is the study of general techniques for quantitatively analyzing decisions under uncertainty. This book offers an introduction to decision analysis, with an emphasis on how to make sophisticated models of economic decisions that involve uncertainty.

The idea of quantitatively analyzing uncertainty may seem puzzling at first. How can we analyze what we do not know? The answer is that we can describe our information and uncertainty in the language of probability theory, which is the basic mathematics of uncertainty. Whenever there is something that we do not know, our uncertainty about it can (in principle) be described quantitatively by probabilities. This book includes an introduction to the basic ideas of probability theory, but the aim throughout is to show how probability theory can be applied to gain understanding and insights into practical decision problems.

In recent decades, economists have increasingly recognized the critical impact of uncertainty on decision making in the formulation of competitive strategies and financial plans. So there is good reason to believe that the study of probability theory should be an important technical topic in the training of business students. But traditional approaches to teaching probability have failed to make clear to most business students the connection between probability theory and the practical decisions under uncertainty that they will confront in their managerial careers. This disconnection occurs because the formulas that are traditionally taught in basic probability courses become hard to apply to problems that involve more than one or two unknown quantities. However, realistic decision problems typically involve many unknown quantities.

Recent advances in computer technology allow us to teach probability differently today in ways that can eliminate this disconnection between theory and application. Electronic spreadsheets, which were invented in

Myerson_C01.qxd 08-09-2004 09:45 Page 2

the late 1970s, offer an intuitive graphical display in which it is much easier to visualize and understand mathematical models, even when they involve large numbers of variables. Randomized (Monte Carlo) simulation methods offer a general way to compute probabilities without learning a lot of specialized formulas. Mathematicians used to object that such simulation methods were computationally slow and inefficient, but the increasing speed of personal computers has continually diminished such concerns. So the emphasis throughout this book is on formulating randomized simulation models in spreadsheets. With this approach, even in this first chapter, we can analyze a model that involves up to 21 random variables.

Chapter 1 is an introduction to working with probabilities in simulation models. Fundamental ideas of conditional probability are developed here in the context of simple examples about learning from observations. With these examples, we show how probabilities can be used to construct a simulation model and then how such a model can be used to estimate other probabilities.

## 1.0 Getting started with Simtools in Excel

All the work in this book is done in Microsoft Excel®. Within Excel, there are many different ways of telling the program to do any one task. Most commands can begin by selecting one of the options on the menu bar at the top of the screen (File, Edit, View, Insert, Format, Tools, Data, Window, Help) and then selecting among the secondary options that appear under the menu bar. Many common command sequences can also be entered by a shortcut keystroke (which is indicated in the pop-down menus) or by clicking on a button in a toolbar that you can display on the screen (try the View menu). But this book will describe commands as if you always use the full command sequence from the top menu. There are many fine books on Excel that you may consult for more background information.

Unfortunately, Excel by itself is not quite enough to do probabilistic decision analysis at the level of this book. To fill in its weak spots, Excel needs to be augmented by a decision-analysis add-in. In this book, we will use an add-in for Excel called Simtools.xla, which is available on the accompanying CD and on the Internet at

http://home.uchicago.edu/~rmyerson/addins.htm

Simtools.xla is comparable to a number of other commercially available add-ins for decision analysis, such as @Risk (www.palisade.com), Crystal Ball (www.decisioneering.com), and Insight.xla (www.analycorp. com).

When you copy or download Simtools.xla, you should save it in your macro library folder under the Microsoft Office folder on your computer's hard disk, which is often called C:\Program Files\Microsoft Office\Office\Library. Then, in Excel, you can install Simtools by using the Tools:Add-Ins

Myerson_C01.qxd 08-09-2004 09:45 Page 3

command sequence and then selecting the Simulation Tools option that will appear in the dialog box. Once installed, Simtools should appear as an option directly in Excel's Tools menu. (Also available with Simtools.xla is another add-in, Formlist.xla, for auditing formulas in spreadsheets. All spreadsheets in this book have lists of formulas that were produced with the help of this Formlist add-in, and I would recommend downloading and installing it along with Simtools.)

## 1.1 How to Toss Coins in a Spreadsheet

When we study probability theory, we are studying uncertainty. To study uncertainty with a spreadsheet, it is useful to create some uncertainty within the spreadsheet itself. Knowing this, the designers of Excel gave us one simple but versatile way to create such uncertainty: the RAND( ) function. This section describes how to use this function to create a spreadsheet that simulates tossing coins (a favorite first example of probability teachers).

With the cursor on cell A1 in the spreadsheet, type the formula

```
=RAND( )
```

and then press the [Enter] key. A number between 0 and 1 is displayed in cell A1. Then (by mouse or arrow keys) move the cursor to cell B1 and enter the formula

```
=IF(A1<0.5, "Heads", "Tails")
```

(The initial equal sign [=] alerts Excel to the fact that what follows is to be interpreted as a mathematical formula, not as a text label. The value of an IF function is its second parameter when the first parameter is a true statement, but when the first parameter is false, its value is the third parameter). Now Excel checks the numerical value of cell A1, and if A1 is less than 0.5, then Excel displays the text "Heads" in cell B1, but if A1 is greater than or equal to 0.5, then Excel displays the text "Tails" in cell B1.

If you have observed this operation carefully, you have noticed that the number in cell A1 changed when the formula was entered into cell B1. In fact, every time we enter anything into the spreadsheet, Excel recalculates everything in the spreadsheet and picks a new value for our RAND( ) function (assuming that Excel's calculation option is set to Automatic on your computer; if not, this setting can be changed under Excel's Tools:Options menu). You can also force such recalculation of the spreadsheet by pushing the Recalc button, which is the [F9] key in Excel. If you have set up this spreadsheet as just described, try pressing [F9] a few times, and watch how the number in cell A1 and the text in cell B1 change each time.

Now taking your hands away from the keyboard for a moment, ask the question: What will be the next value to appear in cell A1 after the next

Myerson_C01.qxd 08-09-2004 09:45 Page 4

time the [F9] key is pressed? The answer is that we do not know. The way that the Excel program determines the value of the RAND( ) function each time is, and should remain, a mystery to us. (The parentheses at the end of the function's name may be taken as a sign that the value depends on things that we cannot see.) My vague understanding is that the program does some very complicated calculations, which may depend in some way on the number of seconds past midnight on the computer's clock but which always generate a value between 0 and 1—nothing more specific than that. The only thing that we need to know about these RAND( ) calculations is that the value, rounded to any number of decimal digits, is equally likely to be any number between 0 and 1. That is, the first digit of the decimal expansion of this number is equally likely to be 0, 1, 2, 3, 4, 5, 6, 7, 8, or 9. Similarly, regardless of the first digit, the second decimal place is equally likely to be any of these digits from 0 to 9, and so on. Thus, the value of RAND( ) is just as likely to be between 0 and 0.1 as it is to be between 0.3 and 0.4. More generally, for any numbers $v$, $w$, $x$, and $y$ that are between 0 and 1, if $v - w = x - y$, then the value of the RAND( ) expression is as likely to be between $w$ and $v$ as it is to be between $y$ and $x$. This information can be summarized by saying that, from our point of view, RAND( ) is drawn from a *Uniform probability distribution over the interval from 0 to 1.*

The cell B1 displays "Heads" if the value of A1 is between 0 and 0.5, whereas it displays "Tails" if A1 is between 0.5 and 1. Because these two intervals have the same length ($0.5 - 0 = 1 - 0.5$), these two events are equally likely. That is, based on our current information, we should think that, after we next press [F9], the next value of cell B1 is equally likely to be Heads or Tails. So we have created a spreadsheet cell that behaves like a fair coin toss every time we press [F9]. We have our first simulation model.

After pressing [F9] to relieve your curiosity, you should press it a few more times to verify that although it is impossible to predict whether "Heads" or "Tails" will occur next in cell B1, they tend to happen about equally often when we recalculate many times. It would be easier to appreciate this fact if we could see many of these simulated coin tosses at once. This is easy to do by using the spreadsheet's Edit:Copy and Edit:Paste commands to make copies of our formulas in the cells A1 and B1. So let us make copies of this range A1:B1 in all of the first 20 rows of the spreadsheet. (Any range in a spreadsheet can be denoted by listing its top-left and bottom-right cells, separated by a colon.)

To copy in Excel, first select the range that you want to copy. This can be done by moving the cursor to cell A1 and then holding down the [Shift] key while moving the cursor to B1 with the right arrow key. (Pressing an arrow key while holding down the [Shift] key selects a rectangular range that has one corner at the cell where the cursor was when the [Shift] key was first depressed and has its opposite corner at the current cell. The selected range will be highlighted in the spreadsheet.) Then, having selected the range A1:B1 in the spreadsheet, open the Edit menu and choose Copy. Faint dots around the A1:B1 range indicate that this range has been copied to Excel's

Myerson_C01.qxd 08-09-2004 09:45 Page 5

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | 0.765196 | Tails | | | | | |
| 2 | 0.223048 | Heads | | | | | |
| 3 | 0.941351 | Tails | | | | | |
| 4 | 0.129491 | Heads | | | | | |
| 5 | 0.688211 | Tails | | | | | |
| 6 | 0.93142 | Tails | | | | | |
| 7 | 0.747859 | Tails | | | | | |
| 8 | 0.166433 | Heads | | | | | |
| 9 | 0.316867 | Heads | | | | | |
| 10 | 0.834753 | Tails | | | | | |
| 11 | 0.45684 | Heads | | | | | |
| 12 | 0.43939 | Heads | | | | | |
| 13 | 0.737932 | Tails | | | | | |
| 14 | 0.331283 | Heads | | | | | |
| 15 | 0.500878 | Tails | | | | | |
| 16 | 0.898716 | Tails | | | | | |
| 17 | 0.202086 | Heads | | FORMULAS FROM RANGE A1:D20 | | | |
| 18 | 0.656246 | Tails | | A1. =RAND() | | | |
| 19 | 0.518526 | Tails | | B1. =IF(A1<0.5,"Heads","Tails") | | | |
| 20 | 0.835384 | Tails | | A1:B1 copied to A1:B20 | | | |

FIGURE **1.1**
**Simulation of coin tossing in a spreadsheet**

Clipboard and is available to be pasted elsewhere. Next, select the range A1:A20 in the spreadsheet, and then open the Edit menu again and choose Paste. Now the spreadsheet should look something like Figure 1.1. (The descriptions that appear in the lower right corner of Figure 1.1 are just text that I typed into cells E17:E20 with the help of the add-in Formlist.xla, which is available along with Simtools.xla.)

In Figure 1.1, we have made 20 copies of the horizontal range A1:B1, putting the left-hand side of each copy in one of the cells in the vertical range A1:A20. So each of these 20 A cells contains the RAND( ) function, but the values that are displayed in cells A1:A20 are different. The value of each RAND( ) is calculated independently of all the other RAND( )s in the spreadsheet. The spreadsheet even calculates different RANDs within one cell independently, and so a cell containing the formula =RAND( )-RAND( ) could take any value between –1 and +1.

The word *independently* is being used here in a specific technical sense that is very important in probability theory. When we say that a collection of unknown quantities are *independent* of each other, we mean that learning the values of some of these quantities would not change our beliefs about the other unknown quantities in this collection. So when we say that the RAND( ) in cell A20 is independent of the other RAND( )s in cells A1:A19, we mean that knowing the values of cells A1:A19 tells us nothing at all about the value of cell A20. If you covered up cell A20 but studied the values of cells A1:A19 very carefully, you should still think that the value of cell A20 is drawn from a uniform distribution over the interval from 0 to 1 (and so, for example, is equally likely to be above or below 0.5), just as you would have thought before looking at any of these cell values.

Each of the 20 cells in B1:B20 contains a copy of the IF function originally entered into cell B1. If you run the cursor through the B cells, however, you

Myerson_C01.qxd 08-09-2004 09:45 Page 6

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | 0.129665 | Heads | | 0.25 | P(Heads) on each toss | | |
| 2 | 0.526313 | Tails | | | | | |
| 3 | 0.604349 | Tails | | | | | |
| 4 | 0.775623 | Tails | | | | | |
| 5 | 0.599284 | Tails | | | | | |
| 6 | 0.938539 | Tails | | | | | |
| 7 | 0.634633 | Tails | | | | | |
| 8 | 0.171716 | Heads | | | | | |
| 9 | 0.981668 | Tails | | | | | |
| 10 | 0.494835 | Tails | | | | | |
| 11 | 0.356306 | Tails | | | | | |
| 12 | 0.058359 | Heads | | | | | |
| 13 | 0.944084 | Tails | | | | | |
| 14 | 0.946007 | Tails | | | | | |
| 15 | 0.840926 | Tails | | | | | |
| 16 | 0.612527 | Tails | | | | | |
| 17 | 0.779595 | Tails | | FORMULAS FROM RANGE A1:D20 | | | |
| 18 | 0.219282 | Heads | | A1. =RAND() | | | |
| 19 | 0.4022 | Tails | | B1. =IF(A1<$D$1,"Heads","Tails") | | | |
| 20 | 0.267915 | Tails | | A1:B1 copied to A1:B20 | | | |

FIGURE **1.2**
**Coin tossing with adjustable probabilities**

should notice that the reference to cell A1 in cell B1 was adjusted when it was copied. For example, B20 contains the formula

=IF(A20<0.5, "Heads", "Tails")

Excel's Copy command treats references to other cells as *relative*, unless we preface them with dollar signs ($) to make them *absolute*. So each of the copied IF functions looks to the cell to the left for the number that it compares to 0.5 to determine whether "Heads" or "Tails" is displayed. Thus we have set up a spreadsheet in which cells B1:B20 simulate 20 independent tosses of fair coins.

Now let us change this spreadsheet so that it can do something that you could not do so easily with coins: simulate 20 independent tosses of an unfair coin that is not equally likely to be Heads or Tails. Into cell B1, enter the formula

=IF(A1<$D$1, "Heads", "Tails")

(You can use the Edit key [F2] to get into the old formula and revise it, simply changing the 0.5 to $D$1.) Next, copy cell B1 and paste it to B2:B20. The dollar signs in the formula tell Excel to treat the reference to D1 as an *absolute*, not to be adjusted when copied, and so the formula in cell B20 (for example) should now be =IF(A20<$D$1, "Heads", "Tails"). Now enter any number between 0 and 1 into cell D1. If you enter 0.25 into cell D1, for example, then your spreadsheet may look something like Figure 1.2.

In this spreadsheet, each cell in B1:B20 will display "Heads" only if the random number to the left is between 0 and 0.25; but it is three times more likely that the number will be above 0.25, and so we should expect substantially more "Tails" than "Heads." Introducing more language of probability theory, we may say that in each cell in the range B1:B20 in this

Myerson_C01.qxd 08-09-2004 09:45 Page 7

spreadsheet, the *probability* of getting "Heads" after the next recalculation is 0.25. If we entered any other number $p$ into cell D1, then this probability of getting a "Heads" in each B cell would change to this new probability $p$, independently of the other B cells.

More generally, when we say that the *probability* of some event $A$ is some number $q$ between 0 and 1, we mean that given our current information, we think that this event $A$ is just as likely to occur as the event that any single RAND( ) in a spreadsheet will take a value less than $q$ after the next recalculation. That is, we would be indifferent between a lottery ticket that would pay \$100 if the event $A$ occurs and a lottery ticket that would pay \$100 if the RAND( )'s next value is less than the number $q$. When this is true, we may write the equation $P(A) = q$.

## 1.2 A Simulation Model of 20 Sales Calls

Perhaps it seems somewhat frivolous to model 20 coin tosses. So let us consider instead a salesperson who will call on 20 customers this week. In each sales call, the salesperson may make a sale or not. To adapt the spreadsheet in Figure 1.2 to this situation, let us begin by re-editing the formula in cell B1 to simulate the first sales call. So instead of displaying "Heads" or "Tails", the cell should display either 1, to represent a sale, or 0, to represent no sale. The RAND( ) value that determines the value of this cell can actually be generated inside the cell's formula itself. So let us select cell B1 and change its formula to

```
=IF(RAND( ) < $D$1,1,0)
```

Next, copy this cell B1 and paste it to the range B1:B20, and you should see a column of random 0s and 1s in B1:B20, simulating the results of 20 sales calls.

We should have a label above this column of 0s and 1s to remind us that they represent sales. We also do not need the RAND column on the left any more. So first eliminate the current leftmost column by selecting cell A1 and using the command sequence Edit:Delete:EntireColumn. Next, insert a new top row by selecting (the new) A1 and using the command sequence Insert:Rows. Notice that the original B1 cell has now become cell A2, and its formula is now

```
=IF(RAND( ) < $C$2,1,0)
```

(Excel automatically adjusts all references appropriately when rows or columns are inserted or deleted in a spreadsheet. Dollar signs make a reference absolute only for the process of copying and pasting.)

Now enter the label 'Sales in the new empty cell A1. If we enter the value 0.5 into cell C2, then, with our new interpretation, this spreadsheet simulates

Myerson_C01.qxd 08-09-2004 09:45 Page 8

| | A | B | C | D | E |
|---|---|---|---|---|---|
| 1 | Sales? | | P(Sale in each call) | | |
| 2 | 1 | | 0.5 | | |
| 3 | 0 | | | | |
| 4 | 0 | | | | |
| 5 | 1 | | | | |
| 6 | 0 | | | | |
| 7 | 1 | | Total Sales in 20 calls | | |
| 8 | 1 | | 8 | | |
| 9 | 1 | | | | |
| 10 | 1 | | | | |
| 11 | 0 | | | | |
| 12 | 0 | | | | |
| 13 | 0 | | | | |
| 14 | 0 | | | | |
| 15 | 0 | | | | |
| 16 | 1 | | | | |
| 17 | 1 | | | | |
| 18 | 0 | | FORMULAS FROM RANGE A1:C21 | | |
| 19 | 0 | | A2. =IF(RAND()<$C$2,1,0) | | |
| 20 | 0 | | A2 copied to A2:A21 | | |
| 21 | 0 | | C8. =SUM(A2:A21) | | |

FIGURE **1.3**
**Simple model of independent sales calls**

a situation in which the salesperson has a probability of 1/2 of making a sale with each of the 20 customers. So we should enter the label

'P(Sale in each call)

into cell C1.

Indicating sales and no sales in the various calls by 1s and 0s makes it easy to count the total number of sales in the spreadsheet. We can simply enter the formula =SUM(A2:A21) into cell C8 (say) and enter the label 'Total Sales in 20 calls into the cell above. The result should look similar to Figure 1.3.

Figure 1.3 simulates the situation for a salesperson whose selling skill is such that the probability is 0.50 of making a sale in any call to a customer like these 20 customers. A more skilled salesperson might be more likely to make a sale in each call. To simulate 20 calls by a more (or less) skilled salesperson, we could simply replace the 0.50 in cell D2 with a higher (or lower) number that appropriately represents the probability that this salesperson will get a sale from any such call, given that person's actual level of skill in selling this product. In this sense, we may think of the probability of making a sale in any call to a customer like these as being a numerical measure of the salesperson's "skill" at this kind of marketing.

But once the number 0.50 is entered into cell C2 in Figure 1.3, the outcomes of the 20 simulated sales calls are determined independently, because each depends on a different RAND variable in the spreadsheet. In this spreadsheet, if we knew that the salesperson made sales to all of the first 19 customers, we would still think that the probability of making a sale to the 20th customer is 1/2 (as likely as a RAND( ) being less than 0.50). Such a strong independence assumption may seem very unrealistic. In real life, even if we knew that this company's salespeople generally make sales in about half of their calls, a string of 19 successful visits might cause us to

Myerson_C01.qxd 08-09-2004 09:45 Page 9

infer that this particular salesperson is very highly skilled, and so we might think that the likelihood of getting a sale on the 20th visit is very high. On the other hand, if we learned that this salesperson had a string of 19 unsuccessful calls, then we might infer lack of skill, and so we might think a sale on the 20th call unlikely. To take account of such dependencies, we need to revise our model to one in which the outcomes of the 20 sales calls are not completely independent.

This independence problem is important to consider whenever we make simulation models in spreadsheets. The RAND function makes it relatively easy for us to make many random variables and events that are all independent of each other. Making random variables that are not completely independent of each other is more difficult. In Excel, if we want the random values of two cells to *not* be independent of each other, then there must be at least one cell with a RAND function in it that directly or indirectly influences both of these cells. (You can trace all the cells that directly and indirectly influence any selected cell by repeatedly using the Tools:Auditing:TracePrecedents command sequence until it adds no more arrows. Then use Tools:Auditing:RemoveAllArrows.) So to avoid assuming that the sales simulated in A2:A21 are completely independent, we should think about some unknown quantities or factors that might influence all these sales events, and we should revise our spreadsheet to take account of our uncertainty about such factors.

Notice in the previous discussion that our concern about assuming independence of the 20 sales calls was really motivated by our uncertainty about the salesperson's skill level. This observation suggests that we should revise the model so that it includes some explicit representation of our uncertainty about the skill level. The way to represent our uncertainty is to make the skill level in cell C2 a random variable. When C2 is random, then the spreadsheet will indeed have a random factor that influences all 20 sales events. Making C2 random is appropriate because we do not actually know the salesperson's level of skill. As a general rule in probability modeling, whenever we have significant uncertainty about a quantity, it is appropriate to model this quantity as a random variable in the spreadsheet.

To keep things as simple as possible in this introductory example, let us suppose (for now) that the salesperson has just two possible skill levels: a high level of skills in which there is a 2/3 probability of making a sale in any call and a low level of skills in which there is a 1/3 probability of making a sale in any call. Under this assumption, we may simply say that the salesperson is "skilled" if there is a 2/3 probability of selling in any sales call; and we may say that the salesperson is "unskilled" if there is a probability of 1/3 of selling in any sales call. For simplicity, let us also suppose that before observing the outcomes of any sales calls, we think that this salesperson is equally likely to be either skilled or unskilled in this sense. To model this situation, we can modify the simple model in Figure 1.3 by entering the formula

```
=IF(RAND( ) < 0.5,2/3,1/3)
```

Myerson_C01.qxd 08-09-2004 09:45 Page 10

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | Sales? | | Salesperson's level of Skill | | | | |
| 2 | 1 | | 0.333333 | | | | |
| 3 | 0 | | (potential rate of sales in the long run) | | | | |
| 4 | 1 | | | | | | |
| 5 | 1 | | | | | | |
| 6 | 0 | | | | | | |
| 7 | 1 | | Total Sales in 20 calls | | | | |
| 8 | 0 | | 9 | | | | |
| 9 | 1 | | | | | | |
| 10 | 0 | | The results of these 20 sales-calls are | | | | |
| 11 | 0 | | conditionally independent given the Skill. | | | | |
| 12 | 1 | | | | | | |
| 13 | 0 | | FORMULAS FROM RANGE A1:C21 | | | | |
| 14 | 0 | | C2. =IF(RAND()<0.5,2/3,1/3) | | | | |
| 15 | 1 | | A2. =IF(RAND()<$C$2,1,0) | | | | |
| 16 | 0 | | A2 copied to A2:A21 | | | | |
| 17 | 0 | | C8. =SUM(A2:A21) | | | | |
| 18 | 1 | | | | | | |
| 19 | 0 | | | | | | |
| 20 | 1 | | | | | | |
| 21 | 0 | | | | | | |

FIGURE **1.4**
**Model of 20 sales calls with uncertainty about salesperson's skill**

into cell C2. Then we can enter the label 'Salesperson's level of Skill into cell C1, and the result should be similar to Figure 1.4.

If we repeatedly press the Recalc key [F9] for the spreadsheet in Figure 1.4, we can see the value of cell D2 changing between 0.666667 and 0.333333. When the value is 0.666667, the spreadsheet model is simulating a skilled salesperson. When the value is 0.333333, the spreadsheet model is simulating an unskilled salesperson. When the salesperson is skilled, more than half of the sales opportunities are usually successful; and when unskilled, usually more than half of the calls are failures. But if you recalculate this spreadsheet many times, you should occasionally see it simulating a salesperson who is skilled but who nevertheless fails in more than half of the 20 calls.

Let us now ask a question that might arise in the work of a supervisor of such salespeople. If the salesperson sold to exactly 9 of the 20 customers called on this week, then what should we think is the probability that the salesperson is actually skilled (but just had bad luck this week)? Remember we are assuming that we believed this person equally likely to be skilled or unskilled at the beginning of the week, but observing 9 sales in 20 gives us some information that should cause our beliefs to change.

To answer this question with our simulation model, we should recalculate the simulation many times. Then we can see how often skilled salespeople make only 9 sales and how often unskilled salespeople make 9 sales. The relative frequencies of these two events in many recalculated simulations will give us a way to estimate how much can be inferred from the evidence of only 9 sales.

There is a problem, however. Our model is two-dimensional (spanning many rows and several columns), and it is not so easy to make hundreds of

Myerson_C01.qxd 08-09-2004 09:45 Page 11

copies of it. In this simple example, we could put the whole model into 22 cells of a single row of the spreadsheet, and then we could make thousands of copies of the model in the rows of our spreadsheet, but recopying a random model thousands of times would make for slow calculations in the spreadsheet, and this technique would become unwieldy in even a slightly more complicated model.

But we do not really need a spreadsheet to hold thousands of copies of our whole model. In our simulation, we only care about two things: Is the salesperson skilled, and how many sales were made? So if we ask Excel to recalculate our model many times, then we only need it to make a table that records the answers to these two questions for each recalculated simulation. All the other information (about which customers among the 20 in A2:A21 actually bought) that is generated in the repeated recalculations of the model can be erased as the next recalculation is done. Excel has the capability to make just such a table, which is called a *data table* in the language of spreadsheets. Simtools gives us a special version of the data table, called a *simulation table*, which is particularly convenient for analyzing such models.

To make a simulation table with Simtools, the output from our model that we want to store in the table must be listed together in a single row. This model-output row must also have at least one blank cell to its left, and beneath the model-output row there must be many rows of blank cells where the simulation data will be written. In our case, the simulation output that we want is in cells C2 and C8, but we can easily repeat the information from these cells in an appropriate model-output row. So to keep track of whether the simulated salesperson's skill is high, let us enter the formula =IF(C2=2/3,1,0) into cell B35, and to keep track of the number of sales achieved by this simulated salesperson, let us enter the formula =C8 into cell C35. To remind ourselves of the interpretations of these cells, let us enter the labels 'Skill hi? and '#Sales into cells B34 and C34, respectively. In cells B34 and B35 here, we apply the convention that, whenever a Yes/No question is posed in a spreadsheet, the value 1 denotes "Yes" and the value 0 denotes "No."

Now we select the range in which the simulation table will be generated. The top row of the selected range must include the model-output range B35:C35 and one additional unused cell to the left (cell A35). So we begin by selecting cell A35 as the top-left cell of our simulation table. With the [Shift] key held down, we can then press the right arrow key twice to select the range A35:C35. Now, when we extend the selected range downward, any lower row that we include in the selection will be filled with simulation data. If we want to record the results of about 1000 simulations, then we should include in our simulation table about 1000 rows below A35:C35. So continuing to hold down the [Shift] key, we can press the [PgDn] key to expand our selection downward. Notice that the number of rows and columns in the range size is indicated in the formula bar at the top left of the screen (or in a pop-up box near the bottom of the selected range) while

Myerson_C01.qxd 08-09-2004 09:45 Page 12

| | A | B | C |
|---|---|---|---|
| 30 | FORMULAS | | |
| 31 | B35. =IF(C2=2/3,1,0) | | |
| 32 | C35. =C8 | | |
| 33 | | | |
| 34 | | Skill hi? | Sales |
| 35 | SimTable | 0 | 5 |
| 36 | 0 | 0 | 5 |
| 37 | 0.001 | 0 | 9 |
| 38 | 0.002 | 0 | 4 |
| 39 | 0.003 | 0 | 6 |
| 40 | 0.004 | 1 | 16 |
| 41 | 0.005 | 1 | 12 |
| 42 | 0.006 | 0 | 6 |
| 43 | 0.007 | 0 | 6 |
| 44 | 0.008 | 0 | 3 |
| 45 | 0.009 | 0 | 7 |
| 46 | 0.01 | 0 | 9 |

Data continues to row 1036

FIGURE **1.5**
**Simulation table for model of 20 sales calls**

we are using the [PgDn] or arrow keys with the [Shift] key held down. Let us expand the selection until the range A35:C1036 is selected, which gives us 1002 rows in our selected range: one row at the top for model output and 1001 rows underneath it in which to store data. (We will see that SimTable output looks a bit nicer when the number of data rows is 1 plus a multiple of 100.) Finally, with this range A35:C1036 selected, we enter the command sequence Tools:Simtools:SimulationTable. After a pause for some computations, the result should look similar to Figure 1.5.

The data from the 1001 simulations are stored in B36:C1036 as values that do not change when we recalculate the spreadsheet by pressing [F9]. Having these values fixed will be important, because it will allow us to sort the data and analyze it statistically without having it change every time we calculate another statistic. But above the data range, the model-output cells B35:C35 still contain the formulas that link them to our simulation model, and so these two cells will change when [F9] is pressed.

In the cells A36:A1036, on the left edge of the simulation table, Simtools has entered percent rank values that show, for each row of simulation data, what fraction of the other data rows are above this row. Because we selected a range with 1001 data rows, these percentile numbers increase by 1/1000 per row, increasing from 0 in the first data row (row 36) to 1 in the last data row (row 1036). (Using 1001 data rows gives us nice even 1/1000 increments here because each data row has 1000 "other data rows" above and below it. These percentile ranks will be used later for making cumulative distribution charts, after sorting the simulation data.)

Now, recall our basic question: What can we infer about skill level if the salesperson gets 9 sales in 20 calls? Paging down the data range, we may find some rows where a skilled salesperson got 9 sales, but most of the cases where 9 sales occurred are cases where the salesperson was unskilled. To get more precise information out of our huge data set, we need to be able to

Myerson_C01.qxd 08-09-2004 09:45 Page 13

count efficiently the number of times that 9 sales (or any other number of sales) occurred with either skill level.

To analyze the skill levels in the simulations where a salesperson made exactly 9 sales, let us first enter the number 9 into cell E33. (Cell E33 will be our comparison cell that we can change if we want to count the number of occurrences of some other number of sales.) Next, into cell E36, enter the formula

```
=IF(C36=$E$33,B36,"..")
```

Then copy cell E36 and paste it to the range E36:E1036. Now the cells in the E column display the value 1 in each data row where a skilled salesperson gets 9 sales, the value 0 in each data row where an unskilled salesperson gets 9 sales, and the label ".." in all other data rows. (Notice the importance of having absolute references to cell $E$33 in the preceding formulas—so that they do not change when they are copied.) With the cells in E36:E1036 acting as counters, we can count the number of times that 9 sales occurred in our simulation data by entering

```
=COUNT(E36:E1036)
```

into cell E28 (say). In this formula we are using the fact that Excel's COUNT function only counts cells that have numerical values and ignores all cells that have nonnumerical values, like ".." here. We can also use the same E cells to tell us how many times that 9 sales occurred for a skilled salesperson by entering the formula

```
=SUM(E36:E1036)
```

into cell E27. The SUM function also ignores the nonnumerical cells, and it does nothing with 0s, and so it is effectively counting the 1s in this range, just as the COUNT function is counting the 1s and 0s.

In the simulation data, the COUNT formula shows 9 sales occurred in 68 of our 1001 simulations, and the SUM formula shows that salespeople with high skills were responsible for 13 of these 68 simulations where 9 sales occurred. Now, suppose that we have a new salesperson whose skill level is not known to us, but we learn that this person also made 9 sales in 20 calls. If our uncertainty about the skill and sales totals for this person are like those in our simulation data, then the result should look like another draw from the population that gave us these 68 simulated salespeople who made 9 sales. Thus, given this sales information, we can estimate that the probability of this person's being skilled is approximately 13/68 = 0.191. This result could be computed by the formula =E27/E28 or by using Excel's AVERAGE function, which is equivalent to the SUM divided by the COUNT. (Like SUM and COUNT, the AVERAGE function ignores nonnumerical cells.) So the formula

```
=AVERAGE(E36:E1036)
```

has been entered into cell E30 in Figure 1.6 to display the value 13/68 = 0.191.

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 23 | FORMULAS FROM RANGE A26:E1036 | | | | | | |
| 24 | B35. =IF(C2=2/3,1,0) | | | | | | |
| 25 | C35. =C8 | | | | With Sales=9: | | |
| 26 | E36. =IF(C36=$E$33,B36,"..") | | | | Frequency in Simtable: | | |
| 27 | E36 copied to E36:E1036 | | | | 13 | Skill hi | |
| 28 | E27. =SUM(E36:E1036) | | | | 68 | Total | |
| 29 | E28. =COUNT(E36:E1036) | | | | P(Skill hi \| Sales=9) | | |
| 30 | E30. =AVERAGE(E36:E1036) | | | | 0.191176 | | |
| 31 | E25. ="With Sales="&E33&":" | | | | | | |
| 32 | | | | Given Sales= | | | |
| 33 | | | | | 9 | | |
| 34 | | Skill hi? | Sales | | | | |
| 35 | SimTable | 1 | 16 | | Skill hi? | | |
| 36 | 0 | 0 | 5 | | .. | | |
| 37 | 0.001 | 0 | 9 | | 0 | | |
| 38 | 0.002 | 0 | 4 | | .. | | |
| 39 | 0.003 | 0 | 6 | | .. | | |
| 40 | 0.004 | 1 | 16 | | .. | | |
| 41 | 0.005 | 1 | 12 | | .. | | |
| 42 | 0.006 | 0 | 6 | | .. | | |
| 43 | 0.007 | 0 | 6 | | .. | | |
| 44 | 0.008 | 0 | 3 | | .. | | |
| 45 | 0.009 | 0 | 7 | | .. | | |
| 46 | 0.01 | 0 | 9 | | 0 | | |

FIGURE **1.6**
**Simulation data and analysis**

This number 13/68 = 0.191 in cell E30 may be called our estimated *conditional probability* of a salesperson having high skill given 9 sales in 20 calls, based on our simulation data. In the notation of probability theory, we often use the symbol "|" to denote the word *given*, and mathematicians often use "≈" to denote the phrase "is approximately equal to." With this notation, the result of our simulation may be summarized by writing $P(\text{Skill high} \mid \text{Sales} = 9) = 13/68$.

In a probability model like this one, where we learn about some unknown quantity by observing other events that depend on it, the probabilities that describe our beliefs before the observation are called *prior* probabilities, and the conditional probabilities that describe our beliefs after the observation are called *posterior* probabilities. With this terminology, we could say that, in this example, our prior probability of the salesperson's skill being high is 1/2, but our posterior probability of the skill being high after we observe 9 sales in 20 calls should be only about 0.2.

**NOTE** (The labels in cells E25 and E29 in Figure 1.6 have been set up so that the number 9 that is displayed will change when the value of cell E33 changes. Each of these cells actually contains a formula that builds up the label using &, which is Excel's concatenation operator. Thus, cell E29 actually contains the text formula

= "With Sales=" & E33 & ":"

which returns the label "With Sales=9:" when cell E33 has the value 9. Similarly, the formula

= "P(Skill hi|Sales=" & E33 & ")"

Myerson_C01.qxd 08-09-2004 09:45 Page 15

returns the label P(Skill hi|Sales=9) in cell E29 in Figure 1.6, even though this formula is not actually listed in the figure. Throughout this book, the figures that show spreadsheets will list all the formulas that are used in numerical computations but may omit some formulas that are used for generating such live text labels.)

## 1.3 Analysis Using Excel's Data-Table Command

Changing the 9 in cell E33 to other numbers, we can see the outcome frequencies for other sales numbers. But it would be helpful to make a table in which we can see together the results for all sales numbers from 0 to 20. Such a table can be easily made with Excel's Data-Table command (sometimes called a *what-if table* in other spreadsheet programs).

Here we will learn to use one form of data table called a *column-input data table*. The structure of a column-input data table is similar to the simulation table that we made in the previous section. (This similarity is not coincidental, because Simtools actually uses Excel's column-input data tables to make its simulation tables.)

To make our data table, we must first make a row that contains the outputs from our model that we want listed in the table, and there must be a blank range under this row where the table will be made. In this case, we want to make a table of the recalculated values of the frequency numbers and conditional probabilities listed in cells E27, E28, and E30 of Figure 1.6, but there these are not in a row and there are not enough blank cells underneath these cells to make the data table there. So we enter the formula =E28 into cell I34 to echo there the number of salespeople who made the given number of sales in cell E33. Next we enter the formula =E27 into cell J34 to display there the number of skilled people who made the given number of sales. To compute the number of unskilled people who made the given number of sales, we enter the formula =I34 – J34 into cell K34. To compute the fraction of skilled salespeople among those who made the given number of sales, we enter the formula =J34/(J34+K34) into cell L34. This range I34:L34 will be the output range at the top of the data table. Underneath, the data table will list the recalculated values of these cells as the parameter in cell E33 is adjusted from 9 to other values between 0 and 20.

The other values that we want to substitute into cell E33 must be listed in the column to the left of the output range, in the rows below it. So we must enter the numbers 0 to 20 into the cells from H35 to H55. (To do so quickly, first enter the number 0 into cell H35, and then select the range H35:H55 and use the command sequence Edit:Fill:Series with the Series dialog box options: Columns, Linear, StepValue 1, and StopValue 20.)

Now we select the range H34:L55 and use the command sequence Data:Table. When the Row and Column Input dialog box comes up, we

Myerson_C01.qxd 08-09-2004 09:45 Page 16

| | H | I | J | K | L | M |
|---|---|---|---|---|---|---|
| 26 | | | | FORMULAS FROM RANGE H34:L55 | | |
| 27 | | | | I34. =E28 | | |
| 28 | | | | J34. =E27 | | |
| 29 | | | | K34. =I34-J34 | | |
| 30 | | | | L34. =J34/(J34+K34) | | |
| 31 | | | | I35:L55. {=TABLE(,D34)} | | |
| 32 | | Frequencies: | | | | |
| 33 | Sales: | Total | Skill Hi | Skill Lo | P(Skill=Hi\|Sales) | |
| 34 | (,E33) | 68 | 13 | 55 | 0.1911765 | |
| 35 | 0 | 0 | 0 | 0 | #DIV/0! | |
| 36 | 1 | 1 | 0 | 1 | 0 | |
| 37 | 2 | 9 | 0 | 9 | 0 | |
| 38 | 3 | 25 | 0 | 25 | 0 | |
| 39 | 4 | 49 | 0 | 49 | 0 | |
| 40 | 5 | 82 | 0 | 82 | 0 | |
| 41 | 6 | 82 | 0 | 82 | 0 | |
| 42 | 7 | 105 | 3 | 102 | 0.0285714 | |
| 43 | 8 | 67 | 5 | 62 | 0.0746269 | |
| 44 | 9 | 68 | 13 | 55 | 0.1911765 | |
| 45 | 10 | 61 | 32 | 29 | 0.5245902 | |
| 46 | 11 | 65 | 55 | 10 | 0.8461538 | |
| 47 | 12 | 75 | 70 | 5 | 0.9333333 | |
| 48 | 13 | 75 | 73 | 2 | 0.9733333 | |
| 49 | 14 | 82 | 82 | 0 | 1 | |
| 50 | 15 | 76 | 75 | 1 | 0.9868421 | |
| 51 | 16 | 50 | 50 | 0 | 1 | |
| 52 | 17 | 23 | 23 | 0 | 1 | |
| 53 | 18 | 5 | 5 | 0 | 1 | |
| 54 | 19 | 1 | 1 | 0 | 1 | |
| 55 | 20 | 0 | 0 | 0 | #DIV/0! | |

FIGURE **1.7**
**Data table of results for different numbers of sales**

leave the Row Input entry blank, but we tab to the Column Input box and enter cell E33 as the Column Input Cell.

Following this Data:Table command, Excel will compute the data entries into the range I35:L55 as follows. For each row in this range, Excel first takes the value of the cell in column H (the leftmost column in our selected range H34:L55) and enters it into cell E33. Then Excel recalculates the whole spreadsheet. The new values of the output row at the top of the data table, I34:L34, are then copied down into the corresponding (I to L) cells in this row. When all the cells in I35:K55 have been filled in this way, Excel restores the original contents of the input cell E33 (the value 9). Notice in Figure 1.7 that the data in row 44 of the data table (I44:K44) is identical to the output above in I34:K34, because row 44 is based on the input value of 9 (from H44), which is the actual current value of cell E33 (as shown previously in Figure 1.6).

If you check the formulas in the cells from I35 to L55, you will find that they all share the special formula {=TABLE(,E33)}. The braces mean that this is an array formula, which Excel has entered into the whole range I35:L55 at once. (Excel will not let you change any one cell in an array; you have to change all or none. To emphasize that these cells together form an array, Figure 1.7 has a border around the data range I35:L55, placed there by using a Format:Cells command.) The TABLE formula tells us that this range contains the data range of a data table that has no row-input cell but has E33 as

Myerson_C01.qxd 08-09-2004 09:45 Page 17

its column-input cell. But recall that the whole range that we selected before invoking the Data:Table command also included one row above this data range and one column to the left of this data range. The column on the left side of the data table contains the alternative input values that are substituted one at a time into the designated column-input cell. The row at the top of the data table contains the output values that were recalculated and entered into the table below as these alternative substitutions were done.

The *conditional probability* of any event $A$ given some other event $B$, denoted by the formula $P(A|B)$, is the probability that we would assign to this event $A$ if we learned that the event $B$ occurred. So the ratios in the L column of the data table give us estimates of the conditional probability of a salesperson's having high skill, given the total number of sales made in 20 calls. Notice how these conditional probabilities of having high skill increase from 0 to 1 as the given sales total increases. If the salesperson made only 7 sales in 20 calls, then we would estimate that the conditional probability of being skilled was only about 0.02. But if the salesperson made 12 sales out of 20 calls, then we would estimate that the conditional probability of being skilled was about 0.93, according to the simulation data in this spreadsheet.

Our simulation table gives us no data in Figure 1.7 for the extreme cases of 0 or 20 sales. But it is obvious from the entries immediately below cell L35 that finding 0 successes in the 20 trials should make us almost sure that the salesperson is unskilled. Similarly, the entries just above cell L55 indicate clearly that 20 successes out of 20 trials should make us almost sure that the salesperson is skilled.

**NOTE** (If you did not have the Simtools add-in for Excel, you could make a simulation table in the following way. Recall that we prepared to make our simulation table by putting our model output in cells B35 and C35, as shown in Figures 1.5 and 1.6. Then we selected a range that had these cells B35:C35 plus the next cell on the left, A35, as its top row, and that also included 1001 more rows below. With this range A35:C1036 selected, we used the command sequence Tools:Simtools:SimulationTable. If you did not have Simtools then, after selecting the range A35:C1036, you could instead use the command sequence Data:Table. When the dialog box asks you to specify input cells, leave the row-input box blank, and specify a column-input cell that has no effect on any of our calculations, such as cell A34 or A35. The result will be that Excel then recalculates the model 1001 times and stores the results in the rows of B36:C1036, just as before. The column-input substitutions affect nothing, but the recalculation of the RANDs gives us different results in each row. Unfortunately, Excel data tables are alive, in the sense that they will be recalculated every time we recalculate the spreadsheet. When we do statistical analysis, we will not want our simulation data to keep changing. To fix this, you could select the data range B36:C1036, copy it to the Clipboard by Edit:Copy, and then with B36:C1036 still selected use the

Myerson_C01.qxd 08-09-2004 09:45 Page 18

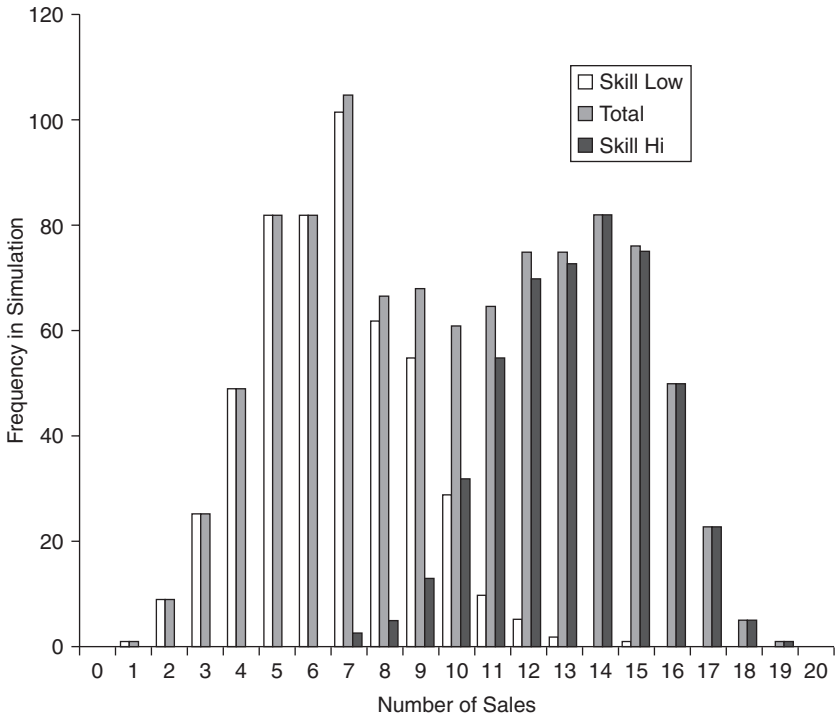


FIGURE **1.8**
**Frequencies of sales in simulation data (total and by skill level)**

command sequence Edit:PasteSpecial:Values. The result is that the TABLE formulas in the data range are replaced by the values that were displayed, and these numerical values now will not change when [F9] is pressed. The Simtools simulation-table procedure, as written in Excel's VBA macro language, actually tells Excel to do just such a data-table command followed by a Copy and PasteSpecial: Values command.)

The data in cells I35:K55 of Figure 1.7 can also be exhibited in a chart, as shown in Figure 1.8. (This chart can be made in our spreadsheet by first selecting the data range I35:K55 and then using the Insert:Chart command sequence, selecting the Column Chart type at the first dialog box. Then at a later dialog box, the three series can be renamed "Total," "Skill Hi," and "Skill Low," and the category-values range H35:H55 can be specified.)

## 1.4 Conditional Independence

Consider again the spreadsheet model in Figure 1.4. In this simulation model, the results of the 20 sales calls in the cells A2:A21 are not independent, because they all depend on the random skill level in cell C2. But notice these results also depend on the independent RAND( ) values in the

Myerson_C01.qxd 08-09-2004 09:45 Page 19

formulas of cells A2:A21, which represent the different customers' idiosyncratic feelings about our salesperson's product. In this case, we may say that the results of the 20 sales calls are conditionally independent of each other when the salesperson's skill level is given.

In general, when we say that some random variables, say **X** and **Y**, are *conditionally independent* given some other random variables, say **Z**, we mean that once you have learned the value of **Z**, getting further information about **X** would not affect your beliefs about **Y**, and getting further information about **Y** would not affect your beliefs about **X**. In a spreadsheet model, such conditional independence holds among random cells **X** and **Y** if the random cells **X** and **Y** are not both influenced by any random cells other than **Z**.

Conditional independence is an important and subtle idea. Because the results of the 20 sales calls are not independent in our model (Figure 1.4), learning the results of the first 19 calls could cause us to revise our beliefs about the probability of a successful sale resulting from the 20th call. But because the sales calls are conditionally independent given the skill level in this model, if we knew that the salesperson had a high skill level (because we could see that cell C2 contained the value 2/3), then we would think that the probability of making a sale in the 20th call was 2/3, even if there were no sales in any of the first 19 calls.

These concepts of conditional probability and conditional independence will be very important for describing what we do in our spreadsheet simulation models. With this terminology, the analysis of our model in Section 1.2 can be summarized as follows:

> The salesperson may either have a high skill level or a low skill level, each with probability 1/2. Each of 20 sales calls may result in either a sale or no sale. The results of 20 sales calls are conditionally independent of each other given the salesperson's level of skill. In each call, the conditional probability of a sale would be 2/3 given that the salesperson's skill is high, but the conditional probability of a sale would be 1/3 given that the salesperson's skill is low. Given this situation, we analyzed data from 1001 simulations of the model to estimate that the conditional probability that the salesperson's skill is high, given 9 sales in the 20 calls, would be approximately 0.2.

## 1.5 A Continuous Random Skill Variable from a Triangular Distribution

In this course we are learning how to make quantitative models of managerial situations that involve uncertainty. Even when our models seem complicated, they will always be simplifications that omit or distort much of the real situation that we are trying to study. Of course, it is not possible

Myerson_C01.qxd 08-09-2004 09:45 Page 20

to think about anything in the world without simplifying it. As powerful as our brains and computers may be, the complexity of the real world is always greater than their limited capacity. So perhaps we should not worry too much about simplification. We can look for useful insights from our analysis of a model like that in Figure 1.4, even while recognizing that it is a simplification of the real situation when new salespersons begin to build a reputation for skill (or lack thereof) in their first sales calls.

Of course, there is always a danger that we may have oversimplified and omitted from our model some important details of the real situation that would significantly change our results. So in applied analytical work, we should always be ready to look at one or more variations of our model, where each new variation is an attempt to add another fact of the real world to the model or to revise some of the simplifying assumptions that we have made. This process of analyzing different variations on our model to see how they may affect our conclusions is called *sensitivity analysis*.

Thus, although a model like the one in Figure 1.4 can give useful insights into a real situation of evaluating salespeople who have limited track records, nevertheless a real sales manager might be disturbed by some of the extreme simplifications that we used to make the model tractable. The first response of a consultant to this manager, would be to get a list of areas where the manager finds shortcomings in the model, and then to build one or more new models that include these suggestions for more realism in these areas. That is, although there can never be one perfect model, one should always be prepared to make a sequence of models where each new model addresses specific concerns that have been expressed about previous models. To do so, a versatile toolkit for making analytical models is needed. The ultimate goal of our study here is to develop such a toolkit.

So now suppose that a sales manager looks at our model in Figure 1.4 and reacts to our assumption of only two possible skill levels as an absurd oversimplification. Perhaps the manager is willing to accept our concept of a "skill level" that represents a person's potential long-run rate of success in sales calls but may tell us that such a "skill level" in this sense could be any number between 0 and 1, not just 1/3 or 2/3. We might then ask the manager to describe more fully beliefs about new salespeople before they make their first 20 sales calls. Suppose that, in response, the manager repeats that a salesperson's potential long-run success rate could be anywhere between 0 and 100% but adds that success rates close to 50% are probably the most common. So to take account of these beliefs, we should represent the skill level by a random variable that can take any value between 0 and 1 and that is more likely to be near 0.5 than any other number.

An unknown quantity that can take any possible value in an interval of numbers is called a *continuous* random variable. To generate such continuous random variables, we generally use one of several famous mathematical probability formulas. One of the simplest formulas that can be used to model beliefs like those described is the formula for a *Triangular* random variable. We can generate such Triangular random variables in our spreadsheets by a

Myerson_C01.qxd 08-09-2004 09:45 Page 21

Simtools function called TRIANINV. To learn about this statistical function, use the command sequence Insert:Function to launch the PasteFunction dialog box, and then search for TRIANINV among the Statistical functions. You will find that TRIANINV takes four parameters. After the first parameter, the latter three parameters are called "lowerbound," "mostlikely," and "upperbound." So the second, third, and fourth parameters should be 0, 0.5, and 1 when we want to make a random variable that could be any number from 0 to 1 but is most likely to be near 0.5. The first parameter, which is mysteriously called "probability," should be a RAND( ) function to make the result a random variable. Thus, the formula

=TRIANINV(RAND( ), 0, 0.5, 1)

returns a Triangular random variable that could be any number from 0 to 1 but is more likely to be near 0.5 than any other number in this interval.

What does a Triangular random variable really mean? One way to describe it is that its probability density has a simple triangular shape, positive over the interval from 0 to 1, with a peak at 0.5. But if that does not mean much to you, you should just enter this formula into a cell in a spreadsheet, press [F9] many times, and watch how this value jumps around. It can take any value between 0 and 1, but the first decimal place is more likely to be 4 or 5 than any other digit. So our uncertainty about the next value of this formula may be a good model for simulating the manager's beliefs about new salespeople.

Figure 1.9 shows a revised version of our old model from Figure 1.4, in which the salesperson's simulated skill level in cell C2 has been changed to a triangular random variable, generated by the formula

=TRIANINV(RAND( ),E3,E4,E5)

where the value of cell E3 is the lower bound 0, the value of E4 is the most likely value 0.5, and the value of cell E5 is the upper bound 1 for the random variable. (*Note:* If you want to make your own version of Figure 1.9, you should probably start on a clean spreadsheet page, because some cells will be different from Figure 1.6.) As before, we make 20 simulated sales calls that depend on the skill level in C2 by entering

=IF(RAND( )<$C$2,1,0)

into cell A2 and copying A2 to A2:A21. The total number of sales is computed by

=SUM(A2:A21)

in cell C8 of Figure 1.9.

Now we need to analyze the model and find how our beliefs about the salesperson should change after we observe the number of successful sales in 20 calls. In our simulation data, we will want to see the skill level (from cell C2) and the resulting number of successful sales (from cell C8) for each simulation. To echo the values of these random variables in one row that

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | Sales? | Salesperson's Skill level (potential long-run sales rate) | | | | | | |
| 2 | 0 | | 0.15444 | from a Triangular distribution with: | | | | |
| 3 | 0 | | | lower bound | 0 | | | |
| 4 | 0 | | | most likely | 0.5 | | | |
| 5 | 1 | | | upper bound | 1 | | | |
| 6 | 0 | | | | | | | |
| 7 | 0 | | Total Sales in 20 calls | | | | | |
| 8 | 0 | | 4 | | | | | |
| 9 | 1 | | | | | | | |
| 10 | 1 | | FORMULAS | | | | | |
| 11 | 0 | | C2. =TRIANINV(RAND(),E3,E4,E5) | | | | | |
| 12 | 1 | | A2. =IF(RAND()<$C$2,1,0) | | A2 copied to A2:A21 | | | |
| 13 | 0 | | C8. =SUM(A2:A21) | | | | | |
| 14 | 0 | | B35. =C2 | | C35. =C8 | | | |
| 15 | 0 | | E36. =IF(C36=$E$34,B36,"..") | | | | | |
| 16 | 0 | | E36 copied to E36:E1036 | | | | | |
| 17 | 0 | | B26. =PERCENTRANK(B36:B1036,$B$24) | | | | | |
| 18 | 0 | | B26 copied to E26 | | | | | |
| 19 | 0 | | B29. =PERCENTILE(B$36:B$1036,A29) | | | | | |
| 20 | 0 | | B29 copied to B29:B31 and E29:E31 | | | | | |
| 21 | 0 | | F37. =COUNT(E36:E1036) | | | | | |
| 22 | | | | | | | | |
| 23 | Compare to Skill cutoff: | | | | | | | |
| 24 | | 0.5 | | | | | | |
| 25 | | P(Skill<=B24) | | | P(Skill<=B24 \| Sales=E34) | | | |
| 26 | | 0.489 | | | 0.592 | | | |
| 27 | | | | | | | | |
| 28 | %Rank | Skill | | %Rank | Skill given Sales=E34 | | | |
| 29 | 0.90 | 0.778536 | | 0.90 | 0.611012 | | | |
| 30 | 0.50 | 0.506776 | | 0.50 | 0.477808 | | | |
| 31 | 0.10 | 0.221666 | | 0.10 | 0.365129 | | | |
| 32 | | | | | | | | |
| 33 | | | | Given Sales= | | | | |
| 34 | | Skill | Sales | | 9 | | | |
| 35 | SimTable | 0.15444 | 4 | | Skill \| Sales=E34 | | | |
| 36 | 0 | 0.594393 | 9 | | 0.594393 | Count(subsample) | | |
| 37 | 0.001 | 0.77063 | 16 | | .. | 73 | | |
| 38 | 0.002 | 0.515541 | 13 | | .. | | | |
| 39 | 0.003 | 0.513165 | 14 | | .. | | | |
| 40 | 0.004 | 0.187994 | 3 | | .. | | | |
| 41 | 0.005 | 0.438439 | 9 | | 0.438439 | | | |
| 42 | 0.006 | 0.288075 | 8 | | .. | | | |
| 43 | 0.007 | 0.22497 | 4 | | .. | | | |
| 44 | 0.008 | 0.664638 | 17 | | .. | | | |
| 45 | 0.009 | 0.309164 | 6 | | .. | | | |

FIGURE **1.9**
**Sales model with skill level from a Triangular distribution**

can be the top of our simulation table, enter the formula =C2 in cell B35 and the formula =C8 in cell C35. Then select the range A35:C1036 and enter the command sequence Tools:Simtools:SimulationTable. When Excel finishes calculating, the skill level and sales numbers from 1001 simulations are contained in the rows of the data range B36:C1036.

One question that we may ask is: Given any number $x$, how likely is the salesperson's skill level to be less than $x$ in our simulation model? To get an answer to this question from our simulation data, we can use the Excel function PERCENTRANK. Looking up this function with the Insert:Function command, you can find that PERCENTRANK takes two parameters, where

Myerson_C01.qxd 08-09-2004 09:45 Page 23

the first parameter is an array or range of cells that contain numerical data, and the second parameter is a number $x$. Then PERCENTRANK(array, $x$) returns the fraction of the numerical data in the array that is less than the given value $x$. (*Note:* Excel's PERCENTRANK reports this fraction rounded to the nearest 1/1000 unless you specify another significance level by an optional third parameter.)

In Figure 1.9, cell B26 contains the formula

=PERCENTRANK(B36:B1036, $B$24)

Cell B24 contains the value 0.5. So the value PERCENTRANK(B36:B1036, 0.5) = 0.489 in cell B30 tells us that 48.9% of the simulated skill levels in cells B36:B1036 are less than 0.5. If we changed the value of cell B24 in this spreadsheet to 0.2, then the value of cell B30 would change to about 0.08, telling us that only about 8% of the simulated skill levels in this data set are below 0.2.

It can actually be more useful to ask a second question that inverts the previous one: Given any number $k$ between 0 and 1, for what number $x$ would the salesperson's skill level have a probability $k$ of being less than $x$? When $k$ is 1/2, the answer to this question is called the *median* of the probability distribution. That is, the median skill level is number $x$ such that the salesperson's skill level is equally likely to be above or below $x$. When $k$ is 1/10, the answer to this question may be close to the lowest skill level that we should worry about, because the salesperson's probability of being below this skill level is only 1/10. On the other hand, when $k$ is 9/10, the answer may be close to the highest skill level that we could reasonably hope to find, because the probability of being more skilled is only 1/10.

To compute an answer to this second question from our simulation data, we can use the Excel function PERCENTILE. According to the Insert:Function dialog box, PERCENTILE takes two parameters, where the first parameter is an array or range of cells in the spreadsheet and the second parameter is a number $k$ that is between 0 and 1. Then PERCENTILE(array, $k$) returns the value $x$ such that the given fraction $k$ of the array's numerical values are less than $x$. (*Note:* Because there may be no data value whose fraction below is exactly the fraction $k$, the PERCENTILE function actually reports a weighted average of the data values whose fractions below are closest to $k$.)

In Figure 1.9, cell B29 contains the formula

=PERCENTILE(B$36:B$1036, A29)

and cell A29 has the value 0.90. Thus, the value 0.778 in cell B29 tells us that 90% of the skill values in E36:E1036 are less than 0.778, so we may say that 0.778 is the *0.90 percentile* for skill in this probability distribution. Cell B29 has been copied down to cells B30 and B31 in Figure 1.9 (where the absolute ($) before 36 and 1036 in its formula maintains the array reference to B$36:B$1036, but the second parameter adjusts to take the percent rank from the cell on the left). So with the value 0.50 in cell A30, the value PERCENTILE(B$36:B$1036, A30) = 0.506 in cell B30 tells us that half of the

Myerson_C01.qxd 08-09-2004 09:45 Page 24

simulated skill levels in B$36:B$1036 are less than 0.506. With the value 0.10 in cell A31, the value PERCENTILE(B$36:B$1036, A31) = 0.221 in cell B31 tells us that 10% of the simulated skill levels in B$36:B$1036 are less than 0.221.

These statistics in cells B27 and B30:B32 tell us something more about this Triangular distribution that, in practice, could be used to verify whether this Triangular distribution is an appropriate description of the manager's actual uncertainty about the new salespeople's skill levels. That is, we might ask the manager whether, before observing the results of any sales calls, a new salesperson's potential long-run sales rate is about equally likely to be above or below 50%, has about a 1/10 chance of being above 78%, and has about a 1/10 chance of being below 22%. If the manager felt that these probabilities were in reasonable agreement with prior beliefs about the skill level of the new salesperson, then our Triangular distribution with parameters (0, 0.5, 1) could be considered a reasonably good fit to those prior beliefs. If these numbers seemed seriously wrong, then we could try adjusting the parameters in E3:E5 and regenerating the simulation data until we found a distribution that generated percentile values in B30:B32 that better fit the manager's actual prior beliefs about the new salesperson.

We can now ask the harder question of how our beliefs about a salesperson should change if 9 sales (say) are made in the first 20 calls. This sales value 9 has been entered into cell E34 in Figure 1.9, so we can filter out the skill levels for individuals who made exactly 9 sales in their 20 calls by entering into cell E36 the formula

```
=IF(C36=$E$34,B36, "..")
```

and then copying cell E36 to the range E36:E1036. Then the range E36:E1036 contains the skill levels of all those who made 9 sales and the nonnumerical text ("..") for all the simulations that had other outcomes. For the spreadsheet shown in Figure 1.9, the simulation table actually contained data from 73 simulations in which 9 successful sales occurred, as indicated by the value of cell F37, which contains the formula

```
=COUNT(E36:E1036)
```

To say more about the distribution of skill levels in these 73 simulations, we copy the PERCENTRANK formula from cell B26 to cell E26, and we copy the PERCENTILE formulas from cells B29:B31 to cells E29:E31, with the numbers (0.90, 0.50, 0.10) also copied from cells A29:A31 to cells D29:D31. Because we did not put any absolute $ before the Bs in the formulas in cells B26 and B29, these copied PERCENTRANK and PERCENTILE formulas in the E column take their data from the subsample in cells E36:E1036, which includes only the skills of people who made 9 (E34) sales. (Like the other statistical functions that we have seen, PERCENTRANK and PERCENTILE ignore the nonnumerical cells in our filtered data range.) With 0.5 in cell B24, the value 0.592 in cell E26 tells us

PERCENTRANK(E36:E1036, 0.5) = 0.592

Myerson_C01.qxd 08-09-2004 09:45 Page 25

and so 59.2% of the simulations with 9 sales had skill levels below 0.5. That is, our simulation results suggest that

$$P(\text{Skill} < 0.5|\text{Sales} = 9) \approx 0.592$$

With the number 0.10 in cell D31, the value 0.365 in cell E31 tells us that

$$\text{PERCENTILE(E\$36:E\$1036, 0.10)} = 0.365$$

and so 10% of the simulations with 9 sales had skill levels below 0.365. That is, our simulation data suggest that

$$0.10 \approx P(\text{Skill} < 0.365|\text{Sales} = 9)$$

Similarly, the values of cells E30 and E29 in Figure 1.9 tell us that, among the simulations with 9 sales, 50% had skill levels below 0.477 and 90% had skill levels below 0.611. (Notice, however, that these estimates are based on a selected subsample of only 73 simulations.) To describe these results, we may say that 0.365 is the *conditional 0.10 percentile* for skill given 9 sales, 0.477 is the *conditional 0.50 percentile* or the *conditional median* for skill given 9 sales, and 0.611 is the *conditional 0.90 percentile* for skill given 9 sales.

By changing the value of cell E34 to something other than 9, we can see how our beliefs about the salesperson's skill level should change if there were different outcomes in the first 20 sales calls. For example, with the simulation data in Figure 1.9, entering 14 into cell E34 would show that, among salespeople who had 14 sales, 1/10 had skill levels below 0.54 (the new value of cell E31), half had skill levels below 0.66 (the new value of cell E30), and 9/10 had skill levels below 0.79 (the new value of cell E29). So these results indicate that 0.54 is the conditional 0.10 percentile for skill given 14 sales, 0.66 is the conditional median for skill given 14 sales, and 0.79 is the conditional 0.90 percentile for skill given 14 sales. (With different simulation data, you may get slightly different numbers.) So according to this simulation data, although 9 sales in the first 20 calls would lead us to believe that the potential long-run sales rate was very likely to be between 0.36 and 0.61, 14 sales would make us believe that the potential long-run sales rate was very likely to be between 0.55 and 0.79. These results may give us a good indication of how much can be learned about a new salesperson's skill from the results of the first 20 sales calls.

## 1.6 Probability Trees and Bayes' Rule

Students often find conditional probability to be hard to use because it involves the use of different counterfactuals. When you are asked to estimate $P(B|A)$, the conditional probability of some event $B$ given some other event $A$, you are being asked how likely you think $B$ is if you learn that $A$

Myerson_C01.qxd 08-09-2004 09:45 Page 26

| | A | B | C | D | E | F |
|---|---|---|---|---|---|---|
| 1 | | SIMULATION MODEL | | | | |
| 2 | | FavStr? | Oil@A? | Oil@B? | | |
| 3 | | 0 | 1 | 1 | | |
| 4 | | | | | | |
| 5 | FORMULAS | | | | | |
| 6 | B3. =IF(RAND()<0.6,1,0) | | | | | |
| 7 | C3. =IF(RAND()<IF($B$3=1,0.3,0.1),1,0) | | | | | |
| 8 | D3. =IF(RAND()<IF($B$3=1,0.3,0.1),1,0) | | | | | |
| 9 | | | | | | |
| 10 | | RESULTS FROM 1000 SIMULATIONS | | | | |
| 11 | | FavStr? | Oil@A? | Oil@B? | Frequency | |
| 12 | | Yes | Yes | Yes | 56 | |
| 13 | | Yes | Yes | No | 129 | |
| 14 | | Yes | No | Yes | 141 | |
| 15 | | Yes | No | No | 284 | |
| 16 | | No | Yes | Yes | 3 | |
| 17 | | No | Yes | No | 37 | |
| 18 | | No | No | Yes | 32 | |
| 19 | | No | No | No | 318 | |

FIGURE **1.10**
**Simulation model of oil exploration example**

is true. So even though you do not know that *A* is true, you have to be ready to say what you would think if it were true. Keeping track of what you can and cannot assume to be true gets confusing. One useful way to keep track of the various conditional probabilities that are being used in the construction of a probability model is to use a probability tree diagram. In this section, we introduce the use of such tree diagrams in the context of a simple example about exploration for new oil fields.

The government sometimes sells rights to drill for oil and gas in tracts of the offshore continental shelf. In this example, suppose that we are working for an oil exploration firm that has been invited to bid on a couple of tracts in an area of the continental shelf that has not been well explored geologically. If favorable geological strata exist underneath the surface in this area, then about 30% of the tracts in this area should contain oil. On the other hand, if favorable geological strata do not exist under this area, then only about 10% of the tracts should contain oil. Our firm is considering the possibility of bidding for the oil development rights in two specific tracts, Tract A and Tract B. At this time, given our current information, our geologists figure that the probability of favorable geological strata existing deep under this general area is 0.6.

A spreadsheet simulation model for this example is shown in Figure 1.10. Cell B3 contains the formula

=IF(RAND( )<0.6,1,0)

Thus, based on our current information, the probability of cell B3 taking the value 1 is equal to 0.6, which is the probability of favorable strata existing in this area of continental shelf. So our interpretation is that cell B3 simulates the existence of favorable strata when the value of B3 is 1, and it simulates the nonexistence of favorable strata when its value is 0. As before, we are using the convention that, in answer to a question, the value 1 means Yes and 0 means No.

Myerson_C01.qxd 08-09-2004 09:45 Page 27

Cell C3 in Figure 1.10 simulates the presence or absence of oil at Tract A, so that a 1 in cell C3 denotes the existence of oil at Tract A but a 0 in cell C3 denotes no oil at Tract A. For this model to correctly simulate the oil exploration problem described, the probability of a 1 in cell C3 must depend on whether there are favorable geological strata, as indicated in cell B3. If the value of cell B3 is 1 (denoting favorable strata), then the probability of a 1 (denoting oil) in cell C3 should be 0.3, which could be accomplished by the formula IF(RAND( ) < 0.3,1,0). (Recall that C1 contains a RAND.) But if the value of cell B3 is 0 (unfavorable strata), then the probability of a 1 (oil) in cell C3 should be 0.1, which could be accomplished by the formula IF(RAND( ) < 0.1,1,0). So using nested IF statements, we enter the formula

```
=IF(RAND( )<IF($B$3=1,0.3,0.1),1,0)
```

into cell C3. Similarly, we can simulate the presence or absence of oil at Tract B by copying the same formula into cell D3. It is important to have the same cell $B$3 referenced by both cells C3 and D3 because the probability of finding oil in each tract is influenced by the same (favorable or unfavorable) geological strata that underlie the whole area including these two tracts. But the formulas in cells C3 and D3 contain RANDs that are evaluated independently, representing the other idiosyncratic factors that may determine the presence or absence of oil in each tract.

With three random variables in our model, each of which can take the value 0 or 1, there are 2∗2∗2 = 8 possible outcomes for each simulation. Figure 1.10 contains a table showing how many times each of these outcomes occurred in the 1000 simulations of this model. This table was constructed using some advanced spreadsheet techniques that will be explained later in Section 1.7. (See Figure 1.14.)

Figure 1.11 shows a tree diagram that summarizes the probability assumptions that were used to construct the simulation model in Figure 1.10. (The simulation model is also reconstructed in cells B3:D3 of Figure 1.11.) To read this tree diagram, you need to understand a few basic rules that we will apply in all such probability tree diagrams.

The circles in the diagrams are called *nodes*, and the lines that go out to the right from these nodes are called *branches*. The diagrams should be read from left to right, beginning at the left-hand side with one node, which we may simply call the *root* of the tree.

Each branch in a probability tree is assigned a *label* that identifies it with some event that could occur. At each node, the branches that go out to the right must represent all possible answers to some question about which we are now uncertain. That is, the branches that go out to the right from each node must represent events that are *mutually exclusive* (in the sense that at most one of them can occur) and *exhaustive* (in the sense that at least one of them must occur). The label of each branch is shown directly above the branch.

The two branches that follow from the root of Figure 1.11 represent the two possible answers to the question: Do favorable strata exist in this area?

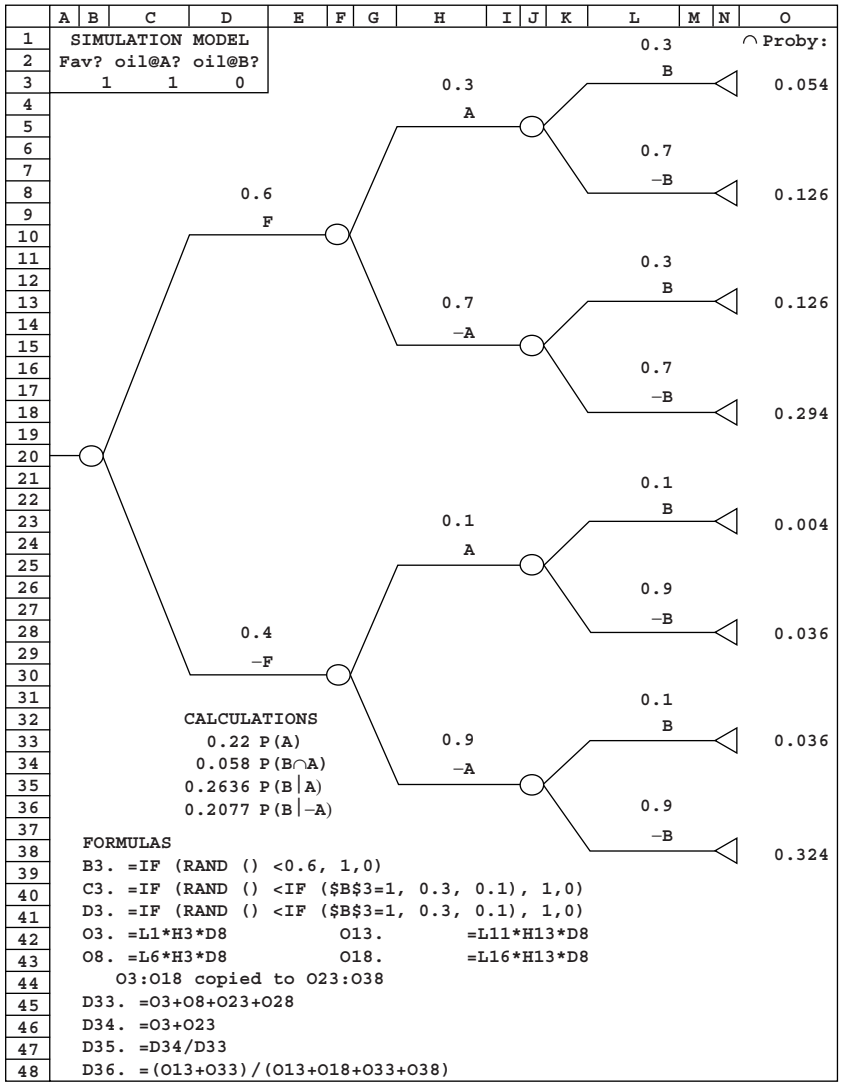


FIGURE **1.11**
**Probability tree for oil exploration example**

Here we use the abbreviation *F* to denote the event that favorable strata exist in this area, and we let *–F* denote the event that favorable strata do not exist. In this abbreviated notation for events, a minus sign may be read as an abbreviation for the word *not*.

Going along either of the branches from the root you come to a second node that is followed by two branches labeled *A* and *–A*. Here *A* denotes the event of that oil can be found in Tract A, so these two branches represent the two possible answers to the question: Can oil be found at Tract A? Similarly, each of the four nodes in the J column is followed by two

Myerson_C01.qxd 08-09-2004 09:45 Page 29

branches labeled $B$ and $-B$, which represent the two possible answers (Yes and No) to the question: Can oil be found at Tract B? Thus, any complete path from left to right through the tree in Figure 1.11 corresponds to one possible combination of values for our simulation model in cells B3:D3. For example, the path from the root to the terminal triangle in cell N8 goes through branches with labels $F$ and $A$ and $-B$, and so it represents the outcome of favorable strata existing with oil in Tract A but not in Tract B, which corresponds to the values (1,1,0) in cells B3:D3.

Each branch in a probability tree is assigned a probability, which we show directly above the label in the diagram. The probability written above any branch always denotes the *conditional probability* of the branch's event given all the events on the path from the root to the node that this branch follows. Thus, for any given node, the probabilities of the various branches that follow this node must always add up to 1, because these branches represent mutually exclusive and exhaustive events.

Of course, the path from the root to itself contains no events. So the probability on a branch that goes out from the root represents the probability of the corresponding event given our current information. In this example, we currently believe that favorable strata have probability 0.6 here, so the probability $P(F) = 0.6$ is shown above the upper branch from the root, while the probability $P(-F) = 1 - 0.6 = 0.4$ appears above the lower branch from the root.

In the upper part of the tree diagram that follows after the $F$ branch, notice that each of the branches with label $A$ or $B$ has the probability 0.3. This is because the conditional probability of finding oil in each tract is 0.3 given favorable strata, that is,

$$P(A|F) = 0.3 \quad \text{and} \quad P(B|F) = 0.3$$

But in the lower part of the tree diagram that follows after the $-F$ branch, each of the branches with label $A$ or $B$ has the probability 0.1, because the conditional probability of finding oil in each tract is 0.1 given unfavorable strata, that is,

$$P(A|-F) = 0.1 \quad \text{and} \quad P(B|-F) = 0.1$$

For any two events $C$ and $D$, we let $C \cap D$ denote the event that $C$ and $D$ are both true, which is called the *intersection* of these two events. So the probability at the topmost $B$ branch that follows the events $F$ and $A$ actually represents $P(B|F \cap A)$, that is, the conditional probability of oil at Tract B given favorable strata and oil at Tract A. Similarly, the probability of the second-highest $B$ branch after $A$ and $-A$ represents $P(B|F \cap -A)$, the conditional probability of oil at Tract B given favorable strata and no oil at Tract A. The equality

$$P(B|F \cap A) = 0.3 = P(B|F \cap -A) = P(B|F)$$

is a result of our assumption that the events of finding oil in the two tracts would be conditionally independent given the existence of favorable strata

Myerson_C01.qxd 08-09-2004 09:45 Page 30

in this area. That is, once we know whether favorable geological strata exist in this area, learning about oil in one tract would not affect our beliefs about the likelihood of finding oil in the other tract.

To explain the initial claim that the tree in Figure 1.11 summarizes the probability assumptions that were used to construct the simulation model in cells B3:D3, the key is that cells in the simulation model have been constructed in the same order as their events appear in the tree. First, in cell B3, we simulate the event of favorable strata using the probability $P(F)$ after the first node in the tree. Then, in cell C3, we simulate the event of oil at Tract A using the conditional probability for this event given the results of the simulation in cell B3, as listed in the second level of branches in the tree (the H column of Figure 1.11). Finally, in cell D3, we simulate the event of oil at Tract B using the conditional probability for this event given the results of the simulation in cell B3 and C3, as listed in the third level of branches in the tree (the L column of Figure 1.11). Of course, the conditional independence of events $A$ and $B$ given the event $F$ implies that the probability of a 1 in cell D3 actually depends only on the simulated value of cell B3, not on cell C3.

For each endpoint of a probability tree, the probability of the intersection of all events on the path from the root to this endpoint can be computed by multiplying all the (conditional) probabilities of all the branches on the path to this endpoint. This probability may be called the *intersection probability* at this endpoint. In Figure 1.11, the intersection probabilities at all endpoints have been computed in the O column. For example, in cell O3 we have

$$P(F \cap A \cap B) = 0.054 = 0.6*0.3*0.3 = P(F)*P(A|F)*P(B|F \cap A)$$

These intersection probabilities in Figure 1.11 are the probabilities of the eight possible outcomes of our simulation model in cells B3:D3. Notice that, as fractions of 1000 simulations, the frequencies found in cells L12:L19 of Figure 1.10 are indeed very close to the actual probabilities that are shown in O3:O38 of Figure 1.11.

The (unconditional) probability of any event can be computed from such a probability tree by adding up the intersection probabilities at all endpoints that follow paths where this event is true. Thus, for example, the probability of finding oil at Tract A is

$$\begin{aligned} P(A) &= P(F \cap A \cap B) + P(F \cap A \cap -B) + P(-F \cap A \cap B) + P(-F \cap A \cap -B) \\ &= 0.054 + 0.126 + 0.004 + 0.036 = 0.22 \end{aligned}$$

as shown in cell D33 of Figure 1.11. Similarly, the probability of finding oil at both tracts is the sum of the intersection probabilities at all endpoints that follow paths where $A$ and $B$ both occur, that is,

$$P(A \cap B) = P(F \cap A \cap B) + P(-F \cap A \cap B) = 0.054 + 0.004 = 0.058$$

as shown in cell D34.

Myerson_C01.qxd 08-09-2004 09:45 Page 31

The conditional probability of any event $C$ given any other event $D$ is the probability of their intersection divided by the probability of the given event $D$. That is,

$$P(C|D) = P(C \cap D)/P(D)$$

This general formula is often called *Bayes' rule*. In a probability tree, this formula tells us that to compute such a conditional probability, we should divide the sum of all endpoint intersection probabilities where both events occur by the sum of all endpoint intersection probabilities where the given event occurs. For example, the conditional probability of finding oil at Tract B given that oil has been found at Tract A (when we do not know about the geological strata) is

$$P(B|A) = P(B \cap A)/P(B) = 0.058/0.22 = 0.2636$$

as shown in cell D35 of Figure 1.11. (Notice that $B \cap A$ and $A \cap B$ both denote the same intersection of events $A$ and $B$.) Similarly, to compute the conditional probability of oil at Tract B given no oil at Tract A, we can divide the sum of intersection probabilities at endpoints that follow both $-A$ and $B$ by the sum of intersection probabilities at all endpoints that follow $-A$, that is,

$$P(B|-A) = P(B \cap -A)/P(-A)$$
$$= (0.126 + 0.036)/(0.126 + 0.294 + 0.036 + 0.324) = 0.2077$$

as computed in cell D36 of Figure 1.11.

Figure 1.12 shows a second probability tree that considers only the questions of whether oil can be found at Tract A and whether oil can be found at Tract B. The question of oil at Tract A (the event $A$ or $-A$) is considered first at the root of this tree, and then the question of oil at Tract B is considered at the pair of nodes that follow after the $A$ and $-A$ branches. The probability at the $A$ branch here is the unconditional probability $P(A) = 0.22$. The probability of the upper $B$ branch, following $A$, is $P(B|A) = 0.2636$. The probability of the lower $B$ branch, following $-A$, is $P(B|-A) = 0.2077$. So all of these branch probabilities in Figure 1.12 can be taken from our calculations of probabilities and conditional probabilities using the first tree in Figure 1.11. For the purposes of computing the probability of any event involving $A$ and $B$, this second probability tree in Figure 1.12 is just as good as the first probability tree in Figure 1.11. For example, cell Y3 of Figure 1.12 computes the endpoint intersection probability

$$P(A \cap B) = P(A)*P(B|A) = 0.22*0.2636 = 0.058$$

which is the same answer that we got from the first tree in cell D34 of Figure 1.12.

Just as the branch probabilities in the first tree were used to construct the simulation model in cells B3:D3 of Figure 1.10, so the branch probabilities in this second tree can be used to construct a second simulation model

Myerson_C01.qxd 08-09-2004 09:45 Page 32

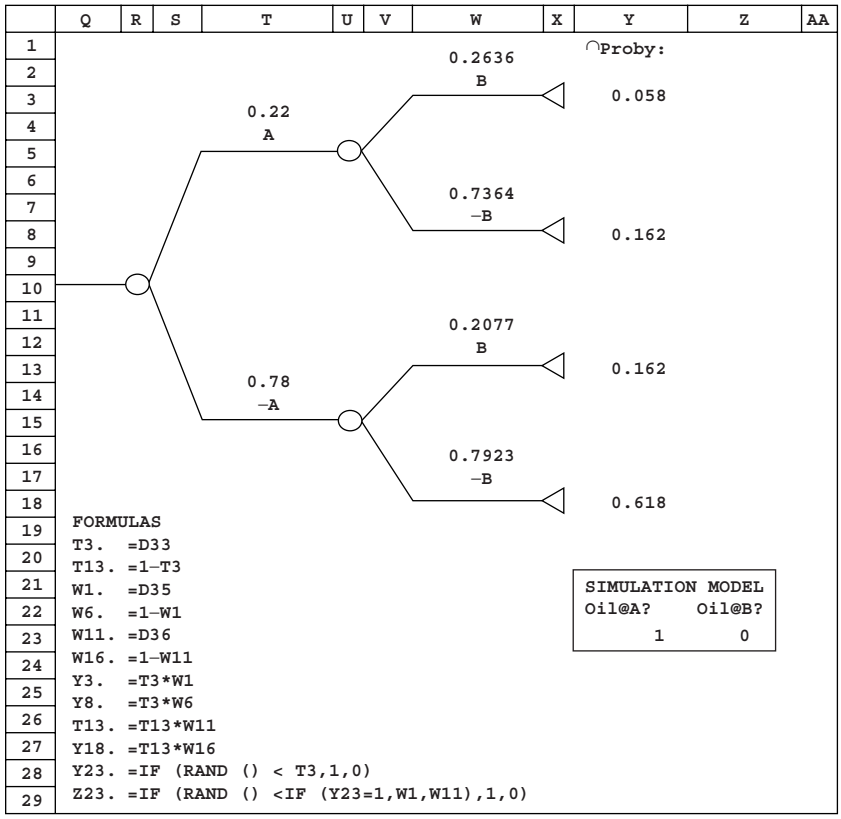


FIGURE **1.12**
**Alternative tree and model for oil exploration example**

of these two tracts, which is shown in cells Y23:Z23 of Figure 1.12. The key to this construction is to consider the simulated events in the same order as they appear in the probability tree. First, in cell Y23, we simulate the answer to the question "Is there oil at Tract A?" so that a value 1 (denoting the answer Yes and the event $A$) occurs with the same probability $P(A)$ that is listed above the $A$ branch in this tree (in cell T3 of the spreadsheet). So the formula in cell Y23 is

=IF(RAND( )<T3,1,0)

Then, in cell Z23,we simulate the answer to the question "Is there oil at Tract B?" so that a value 1 (denoting the answer Yes and the event $B$) occurs with the appropriate probability given the simulated value of cell Y23, either $P(B|A) = 0.2636$ if Y23 = 1, or $P(B|-A) = 0.2077$ if Y23 = 0. These two conditional probabilities appear above the appropriate branches of this tree in cells W1 and W11 of this spreadsheet, and so the formula in cell Z23 is

=IF(RAND( )<IF(Y23=1,W1,W11),1,0)

Myerson_C01.qxd 08-09-2004 09:45 Page 33

If we only care about whether oil can be found at Tracts A and B (and we do not otherwise care about whether the deep geological strata are "favorable" or not), then these two simulation models in Figures 1.10 and 1.12 are just as good, and they can give equally valid estimates of any probability that involves these events $A$ and $B$. In fact, if we were given a large table of simulation data that was generated either by cells C3:D3 in Figure 1.10 or by cells Y23:Z23 in Figure 1.12, there is no statistical test that could distinguish which of these two models generated the data.

So why would we use one model rather than the other? In practice, we should use the model with the parameters that are easier for us to assess. The second model corresponds to the simpler probability tree in Figure 1.12, which only requires us to specify $P(A)$, $P(B|A)$, and $P(B|\text{–}A)$. (The other three branch probabilities can be computed from the basic rule that at each node, the sum of probabilities on branches going out to the right must equal 1.) That is, this model requires us to ask our best expert three questions:

1. How likely is it that oil can be found at Tract A?
2. If you learned that oil could be found at Tract A, then how likely would you consider oil at Tract B?
3. If you learned that oil could be not found at Tract A, then how likely would you consider oil at Tract B?

In an example like this one, it may be critical to recognize that finding oil at Tract A would increase our beliefs about the probability of finding oil at Tract B. So we should not make the error of assuming independence between these two tracts, and we need to ask the latter two questions to differentiate $P(B|A)$ from $P(B|\text{–}A)$.

But although our expert should recognize that finding oil at Tract A would make us somewhat more optimistic about finding oil at Tract B, even an expert might have difficulty quantifying how much more optimistic. In an attempt to analyze this question, the expert might start talking about why oil at one tract should affect our beliefs about the probability of oil at a neighboring tract: "because it is evidence that favorable geological strata may underlie this whole area." Such a statement would suggest that the expert might be more comfortable assessing the parameters of the first model: the probability of favorable strata, the probability of finding oil at any tract given favorable strata, and the probability of finding oil at any tract given unfavorable strata. Notice that with the assumptions of symmetry between the tracts and conditional independence of the tracts given the underlying strata, the first probability tree actually requires only three probability assessments, even though it has many more branches.

Probability trees are a useful way to organize information in such problems, but they become very hard to draw when there are more than a few

Myerson_C01.qxd 08-09-2004 09:45 Page 34

events to consider. For example, suppose we tried to draw a probability tree like Figure 1.11 for our salesperson example. We would begin at the root with two branches to represent the events of Skill = Hi and Skill = Low. Next, instead of drilling for oil at two tracts, we are looking for sales from 20 customers. If we represented each customer by a column of nodes with two branches, we would have to draw a probability tree with over 2,000,000 endpoints. To simplify matters, we might aggregate the results from different customers and simply ask how many sales were achieved out of 20. This will give us a tree with 42 endpoints. After the Skill = Hi branch, there would be a node with 21 branches to represent the outcomes of 0 sales, 1 sale, 2 sales, . . . , 20 sales for a high-skilled salesperson. Similarly, after the Skill = Low branch, there would be another 21 branches to represent the outcomes of 0 sales, 1 sale, 2 sales, . . . , 20 sales for a low-skilled salesperson. But even for this simplified tree, it would be rather difficult to draw all these branches on a page.

Figure 1.13 shows the essential information that would be contained in such a probability tree for our salesperson example, but the lines of the tree diagram are omitted, which saves space (but perhaps reduces clarity). Where the tree would begin with two branches to represent the Skill = Hi and Skill = Low events, we list these events and their probabilities in the A column of Figure 1.13. The probability $P$(Skill = Hi) is entered into cell A3, and the probability $P$(Skill = Low) is entered into cell A27. Next, after each level of skill, the tree would have 21 branches for the possible numbers of sales from 0 to 20. These possible numbers of sales are listed in cells B3:B23 for the Skill = Hi case, and in cells B27:B47 for the Skill = Low case.

To complete the description of the probability tree, we need to specify the conditional probabilities that would be associated with these 42 branches. That is, for each number $k$ from 0 to 20, we need the conditional probabilities $P(k \text{ sales}|\text{Skill} = \text{Hi})$ and $P(k \text{ sales}|\text{Skill} = \text{Low})$. Finding these conditional probabilities is the essence of the problem, and much of our work in Section 1.2 was essentially showing how these conditional probabilities could be estimated from a simulation model. But such questions have been considered by many mathematicians before, and they have developed a theory of *binomial distributions* that can be used to directly compute probabilities in this specific kind of situation. Because the emphasis here is on modeling skills that you can apply to more general kinds of problems, this theory is not stressed. But some of you may have studied binomial distributions in another course, so, briefly, this is how they can be applied here.

Binomial distribution theory tells us that if there will be $n$ independent trials (for some given number $n$), and the probability of a "success" in each trial will be $p$ (for some given probability $p$), and the result of each trial will be independent of all other trials, then the probability of getting exactly $k$ successes in $n$ trials is given by a specific formula that is returned by the Excel function

BINOMDIST(k,n,p,0)

Myerson_C01.qxd 08-09-2004 09:45 Page 35

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | Skill=Hi | [P(Sale)=2/3 in each call] | | | | | |
| 2 | P(Skill) | k | P(Sales=k\|Skill) | Product | | P(Skill=Hi\|Sales=k) | |
| 3 | 0.5 | 0 | 0.0000 | 0.0000 | | 0.0000 | |
| 4 | | 1 | 0.0000 | 0.0000 | | 0.0000 | |
| 5 | | 2 | 0.0000 | 0.0000 | | 0.0000 | |
| 6 | | 3 | 0.0000 | 0.0000 | | 0.0001 | |
| 7 | | 4 | 0.0000 | 0.0000 | | 0.0002 | |
| 8 | | 5 | 0.0001 | 0.0001 | | 0.0010 | |
| 9 | | 6 | 0.0007 | 0.0004 | | 0.0039 | |
| 10 | | 7 | 0.0028 | 0.0014 | | 0.0154 | |
| 11 | | 8 | 0.0092 | 0.0046 | | 0.0588 | |
| 12 | | 9 | 0.0247 | 0.0123 | | 0.2000 | |
| 13 | | 10 | 0.0543 | 0.0271 | | 0.5000 | |
| 14 | | 11 | 0.0987 | 0.0493 | | 0.8000 | |
| 15 | | 12 | 0.1480 | 0.0740 | | 0.9412 | |
| 16 | | 13 | 0.1821 | 0.0911 | | 0.9846 | |
| 17 | | 14 | 0.1821 | 0.0911 | | 0.9961 | |
| 18 | | 15 | 0.1457 | 0.0729 | | 0.9990 | |
| 19 | | 16 | 0.0911 | 0.0455 | | 0.9998 | |
| 20 | | 17 | 0.0429 | 0.0214 | | 0.9999 | |
| 21 | | 18 | 0.0143 | 0.0071 | | 1.0000 | |
| 22 | | 19 | 0.0030 | 0.0015 | | 1.0000 | |
| 23 | | 20 | 0.0003 | 0.0002 | | 1.0000 | |
| 24 | | | | | | | |
| 25 | Skill=Lo | [P(Sale)=1/3 in each call] | | | | | |
| 26 | P(Skill) | k | P(Sales=k\|Skill) | Product | | FORMULAS | |
| 27 | 0.5 | 0 | 0.0003 | 0.0002 | | C3. =BINOMDIST(B3,20,2/3,0) | |
| 28 | | 1 | 0.0030 | 0.0015 | | D3. =$A$3*C3 | |
| 29 | | 2 | 0.0143 | 0.0071 | | C3:D3 copied to C3:D23 | |
| 30 | | 3 | 0.0429 | 0.0214 | | A27. =1-A3 | |
| 31 | | 4 | 0.0911 | 0.0455 | | C27. =BINOMDIST(B27,20,1/3,0) | |
| 32 | | 5 | 0.1457 | 0.0729 | | D27. =$A$27*C27 | |
| 33 | | 6 | 0.1821 | 0.0911 | | C27:D27 copied to C27:D47 | |
| 34 | | 7 | 0.1821 | 0.0911 | | F3. =D3/(D3+D27) | |
| 35 | | 8 | 0.1480 | 0.0740 | | F3 copied to F3:F23 | |
| 36 | | 9 | 0.0987 | 0.0493 | | | |
| 37 | | 10 | 0.0543 | 0.0271 | | | |
| 38 | | 11 | 0.0247 | 0.0123 | | | |
| 39 | | 12 | 0.0092 | 0.0046 | | | |
| 40 | | 13 | 0.0028 | 0.0014 | | | |
| 41 | | 14 | 0.0007 | 0.0004 | | | |
| 42 | | 15 | 0.0001 | 0.0001 | | | |
| 43 | | 16 | 0.0000 | 0.0000 | | | |
| 44 | | 17 | 0.0000 | 0.0000 | | | |
| 45 | | 18 | 0.0000 | 0.0000 | | | |
| 46 | | 19 | 0.0000 | 0.0000 | | | |
| 47 | | 20 | 0.0000 | 0.0000 | | | |

FIGURE **1.13**
**Binomial probability computations for salesperson example**

This (magical) formula is applied in column C of Figure 1.13. The number of trials is $n = 20$. These trials are not independent when the salesperson's skill is unknown, and so the binomial theory could not be applied without knowing the skill. But if we are given that Skill = Hi, then the trials (calls on different customers) become independent with a probability of success (a sale) being 2/3 in each trial. So the probabilities $P(k \text{ sales}|\text{Skill} = \text{Hi})$ are computed, for all the various values of $k$ in B3:B23, by entering the formula

=BINOMDIST(B3,20,2/3,0)

Myerson_C01.qxd 08-09-2004 09:45 Page 36

into cell C3, then copying C3 to C3:C23. Similarly, if we are given that Skill = Low, then the trials also become independent, but now the probability of success is only 1/3 in each trial, and so the probabilities $P(k \text{ sales}|\text{Skill} = \text{Low})$ are computed, for all the various values of $k$ in B27:B47, by entering the formula

```
=BINOMDIST(B27,20,1/3,0)
```

into cell C27, then copying C27 to C27:C47.

The intersection probabilities at the endpoints of this probability tree would just be the first branch's probability (0.5 in each case) multiplied by the binomial conditional probability in the C column of Figure 1.13. These intersection probabilities are computed in the D column of Figure 1.13. For example, cell D12 computes

$$P(\text{Sales} = 9 \cap \text{Skill} = \text{Hi}) = P(\text{Skill} = \text{Hi}) * P(\text{Sales} = 9|\text{Skill} = \text{Hi})$$
$$= 0.5 * 0.0247 = 0.0123$$

Similarly, cell D36 computes

$$P(\text{Sales} = 9 \cap \text{Skill} = \text{Low}) = P(\text{Skill} = \text{Low}) * P(\text{Sales} = 9|\text{Skill} = \text{Low})$$
$$= 0.5 * 0.0987 = 0.0493$$

To compute conditional probability of Skill = Hi given Sales = 9, we divide the probability of the outcome Sales = 9 ∩ Skill = Hi by the sum of probabilities of all outcomes where Sales = 9. That is,

$$P(\text{Skill} = \text{Hi}|\text{Sales} = 9)$$
$$= P(\text{Sales} = 9 \cap \text{Skill} = \text{Hi})/(P(\text{Sales} = 9 \cap \text{Skill} = \text{Hi}) + P(\text{Sales}$$
$$= 9 \cap \text{Skill} = \text{Low}))$$
$$= 0.0123/(0.0123 + 0.0493) = 0.2000$$

This result is computed in cell F12 of Figure 1.13 by the formula =D12/(D12+D36). Copying this formula to D3:D23 yields the correct answer to questions that we were trying to estimate in Sections 1.2 and 1.3. In particular, comparing cells D3:D23 in Figure 1.13 to cells L35:L55 in Figure 1.7, you can see how accurate our simulation estimates were.

It is very nice to have this special BINOMDIST formula that enables us to compute these conditional probabilities exactly. But special formulas require special structures that restrict the class of models that we can consider. In contrast, the methods of simulation analysis can be extended to virtually any case of decision making under uncertainty. This generality is the main reason why this book emphasizes simulation analysis as a

Myerson_C01.qxd 08-09-2004 09:45 Page 37

unifying analytical methodology. Some inexactness of results is a small price to pay for the great generality of practical applications that we will be able to address.

## 1.7 Advanced Spreadsheet Techniques: Constructing a Table with Multiple Inputs

Figure 1.14 shows how the simulation model of Figure 1.10 was analyzed to generate the table of simulation results, which appeared in cells B13:E20 of Figure 1.10 but appears in cells I13:L20 of Figure 1.14. Figure 1.14 looks complicated, but you should recognize that there are really three basic parts of this spreadsheet. First, the range B3:D3 contains the simulation model for our oil exploration example. Second, below the model, we have generated a simulation table that contains data from 1000 recalculations of this model, stored in the range B4:D1003. Only the top 30 rows of data

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | | SIMULATION MODEL | | | | Total | | 1 | 1 | 1 | 1 | | | |
| 2 | | FavStr? | Oil@A? | Oil@B? | | 56 | | Index | Pattern | | | | | |
| 3 | SimTable | 1 | 0 | 1 | | =I1:K1? | | 1 | 1 | 1 | 1 | | | |
| 4 | 0 | 0 | 0 | 0 | | 0 | | 2 | 1 | 1 | 0 | | | |
| 5 | 0.001001 | 1 | 0 | 0 | | 0 | | 3 | 1 | 0 | 1 | | | |
| 6 | 0.002002 | 1 | 0 | 0 | | 0 | | 4 | 1 | 0 | 0 | | | |
| 7 | 0.003003 | 1 | 0 | 1 | | 0 | | 5 | 0 | 1 | 1 | | | |
| 8 | 0.004004 | 0 | 0 | 0 | | 0 | | 6 | 0 | 1 | 0 | | | |
| 9 | 0.005005 | 0 | 0 | 0 | | 0 | | 7 | 0 | 0 | 1 | | | |
| 10 | 0.006006 | 1 | 1 | 1 | | 1 | | 8 | 0 | 0 | 0 | | | |
| 11 | 0.007007 | 1 | 0 | 1 | | 0 | | | | | | | | |
| 12 | 0.008008 | 1 | 0 | 1 | | 0 | | | FavStr? | Oil@A? | Oil@B? | Frequency | | |
| 13 | 0.009009 | 0 | 1 | 1 | | 0 | | 1 | Yes | Yes | Yes | 56 | | |
| 14 | 0.01001 | 1 | 1 | 0 | | 0 | | 2 | Yes | Yes | No | 129 | | |
| 15 | 0.011011 | 1 | 1 | 0 | | 0 | | 3 | Yes | No | Yes | 141 | | |
| 16 | 0.012012 | 1 | 1 | 0 | | 0 | | 4 | Yes | No | No | 284 | | |
| 17 | 0.013013 | 0 | 0 | 0 | | 0 | | 5 | No | Yes | Yes | 3 | | |
| 18 | 0.014014 | 0 | 0 | 0 | | 0 | | 6 | No | Yes | No | 37 | | |
| 19 | 0.015015 | 1 | 0 | 0 | | 0 | | 7 | No | No | Yes | 32 | | |
| 20 | 0.016016 | 1 | 0 | 0 | | 0 | | 8 | No | No | No | 318 | | |
| 21 | 0.017017 | 0 | 0 | 0 | | 0 | | | | | | | | |
| 22 | 0.018018 | 1 | 0 | 1 | | 0 | | FORMULAS | | | | | | |
| 23 | 0.019019 | 0 | 0 | 1 | | 0 | | B3. =IF(RAND()<0.6,1,0) | | | | | | |
| 24 | 0.02002 | 1 | 0 | 0 | | 0 | | C3. =IF(RAND()<IF($B$3=1,0.3,0.1),1,0) | | | | | | |
| 25 | 0.021021 | 1 | 1 | 0 | | 0 | | D3. =IF(RAND()<IF($B$3=1,0.3,0.1),1,0) | | | | | | |
| 26 | 0.022022 | 0 | 0 | 0 | | 0 | | F4. =IF(AND(B4=$I$1,C4=$J$1,D4=$K$1),1,0) | | | | | | |
| 27 | 0.023023 | 1 | 0 | 1 | | 0 | | F4 opiedcto F4:F1003 | | | | | | |
| 28 | 0.024024 | 0 | 0 | 0 | | 0 | | F2. =SUM(F4:F1003) | | | | | | |
| 29 | 0.025025 | 1 | 0 | 0 | | 0 | | I1. =INDEX(I3:I10,$H$1) | | | | I1 copied to I1:K1 | | |
| 30 | 0.026026 | 1 | 1 | 0 | | 0 | | I13. =IF(I1=1,"Yes","No") | | | | I13 copied to J13:K13 | | |
| 31 | 0.027027 | 1 | 1 | 1 | | 1 | | L13. =F2 | | | | H13. =H1 | | |
| 32 | 0.028028 | 1 | 1 | 1 | | 1 | | I14:L20. {=TABLE(,H1)} | | | | | | |

FIGURE **1.14** **Computing frequencies of different outcomes in oil exploration example**

Myerson_C01.qxd 08-09-2004 09:45 Page 38

are shown in Figure 1.14, but the SimTable data continues in the spreadsheet down to row 1003. Third, in columns F through L, we have a data analysis section that serves to extract basic summary statistics from the simulation data. The following paragraphs explain how this analytical section works, because it applies some advanced techniques that you may find useful later.

The cells in the F column of Figure 1.14 are designed to count how many simulation outcomes match the pattern shown in the range I1:K1 (currently 1,1,1). Cell F4 contains the formula

```
=IF(AND(B4=$I$1,C4=$J$1,D4=$K$1),1,0)
```

and cell F4 has been copied to F4:F1003. Here we use Excel's AND function, which takes any number of true-or-false propositions as inputs and returns the value TRUE if all of these inputs are true, but returns the value FALSE if any of these inputs are not true. So if the three values in B4:D4 exactly matched the corresponding values in I1:K1, then the expression

```
AND(B4=$I$1, C4=$J$1, D4=$K$1)
```

would be TRUE, which would make cell F4 take the value 1. But otherwise, if any cell in B4:D4 is different from the corresponding cell in I1:K1, then this AND expression must be FALSE and the value of cell F4 must be 0.

Because the dollar signs in cell F4 of Figure 1.14 make the references to I1:K1 absolute constants in the copying process, each cell in F4:F1003 takes the value 1 if the B:D cells in its row are all the same as the pattern in I1:K1, and it takes the value 0 otherwise. Then cell F2 contains the formula

```
=SUM(F4:F1003)
```

and so it displays the total number of rows where the B:D cells match the pattern in I1:K1. With value 1 in each of the cells of the pattern range I1:K1, the value 56 in cell F2 tells us that among the 1000 simulations stored in this table, there were exactly 56 simulations in which the outcome was favorable strata with oil in both tracts (including rows 10, 31, and 32 shown here).

Now if we change the spreadsheet in Figure 1.14 by putting different combinations of 0s and 1s in the pattern range I1:K1, we can find in cell F2 the total number of times that each possible outcome occurred. These results (with the interpretation that 1 means Yes and 0 means No in our simulation data) are shown in cells I13:L20 of Figure 1.14. For example, cell L14 tells us that there were 129 rows (including rows 14, 15, 16, 25, and 30) where the simulation outcome was (1,1,0), which represents the event that favorable strata exist in this area and there is oil in Tract A but there is no oil in Tract B.

The statistical summary in cells I13:L20 of Figure 1.14 was generated by changing the values in I1:K1. Because there are only $2*2*2 = 8$ possible combinations of 0s and 1s to be tried in this three-cell range, it is not too tedious to make these substitutions by hand. However, you might want to know how this process was accomplished automatically in this spreadsheet.

Myerson_C01.qxd 08-09-2004 09:45 Page 39

The range I3:K10 in Figure 1.14 lists, in its eight rows, all eight possible combinations of 0s and 1s that could be entered into our three-cell pattern I1:K1. (Simtools.xla provides a CombineRows procedure that can be used to automatically generate tables of all combinations, such as we have in I3:K10 here. But these eight combinations here are easy enough to generate by hand, and we will not have further use of this feature, so its discussion is omitted.) Now we want to automatically force the spreadsheet to re-evaluate the total in cell F2 as each of these rows is substituted into the pattern range I1:K1.

When we need to re-evaluate a spreadsheet with different values substituted into a cell, we can do it with a column-input data table. But how can we re-evaluate a spreadsheet with different values being substituted into three cells? The trick is to set up these three cells so that their values are controlled by one cell, and then let that one cell be the column-input cell of our data table.

In Figure 1.14, the cell I1 contains the formula

```
=INDEX(I3:I10,$H$1)
```

and this formula (with its absolute reference to $H$1) has been copied to cells J1 and K1. Given a range of cells in one column as its first parameter and given a number as its second parameter, Excel's INDEX function returns the value of the cell in the given range that is indicated by the given number, counting the top cell as number 1, the second cell as number 2, and so on. So when H1 equals 1, the cells I1:K1 take their values from the corresponding cells in the top row of I3:K10 (that is, I3:K3). But if the value of H1 were changed to 2, then cells I1:K1 would take their values from the second row of I3:K10, which are (1,1,0) in cells I4:K4, and then cell F2 would be counting the number of simulations where favorable strata occurred with oil at Tract A but not at Tract B. So by changing the value of H1 from 1 to 8, we can set the values of I1:K1 equal to each of the possible combinations that are listed in the rows of I3:K10.

Now we are ready to examine how the statistical table in cells I13:L20 of Figure 1.14 was generated. In the top row of this range, the I:K cells repeat the pattern range I1:K1, except that they translate a 1 to a Yes and a 0 to a No. The formula in cell I13 is

```
=IF(I1=1, "Yes", "No")
```

and this formula has been copied to J13 and K13. Cell L13 contains the formula =F2, so L13 simply repeats the statistical count of pattern matches that we computed in cell F2. Then the range I14:L20 below these cells has been filled by a column-input data table, which shows how the values of I13:L13 would change if the corresponding inputs at left (the numbers 2, 3, 4, . . . , 8 in cells H14:H20) were substituted into our index-control cell H1. That is, the cells I14:L20 are filled with the formula {=TABLE(,H1)}, which was created by selecting H13:L20 and then using the Data:Table menu command, with H1 as the column-input cell.

Myerson_C01.qxd 08-09-2004 09:45 Page 40

## 1.8 Using Models

We have introduced several models in this chapter, and we will see many more in the rest of this book. So it may be helpful to take a general look at how models are used.

A model of a complex system is just a simpler and more accessible system that is in some ways like the complex system, so we can try to learn about the complex system's properties by doing experiments with the model. Every model is a simplification, omitting many of the true complexities of the world that it is intended to describe. The only model that could ever be exactly like the real world in every aspect would be the real world itself, and if we denied the relevance of any other models, then we could make no prediction about the results of any decision until it was actually implemented.

This book is about formal quantitative models in spreadsheets. But of course there are many other kinds of models. Even when you think informally about some situation in which you have to make a decision, you are creating a model of this situation in your head as you imagine how your decision might affect the various consequences that you care about. The model in your imagination is a simplification of reality just as a spreadsheet model is, and neither model is necessarily better or more accurate.

One advantage of formal quantitative models is that their details are completely specified in a way that other people can examine and discuss. Indeed, the task of constructing a quantitative model can be divided among various members of a team, giving people responsibility for making the parts of the model that they know best. In contrast, the model of your imagination must be in your mind alone, where it is not so transparent to others, and so it is harder for you to ask friends or teammates to help scrutinize the hidden assumptions of your thoughts. There will always be questionable simplifying assumptions in any model, but in a formal model these assumptions are out in the open for examination by anyone who knows how to read it. When an important applied decision is to be based on the analysis of a formal quantitative model, we should always check to see how our conclusions might change when some of the model's questionable assumptions are changed. This process is called *sensitivity analysis*.

A disadvantage of formal quantitative models is that they may encourage us to focus only on those aspects of the situation that are quantitatively measurable. Furthermore, nothing goes into a formal model until we focus on it and consciously enter it into some part of the model. In contrast, the informal models of your imagination may draw almost effortlessly on your memories from many different experiences in life, possibly including some memories of which you are not even conscious, in which case the results of your thought process may be called intuition. So it is important to compare the results of a formal quantitative model with our intuition about the

Myerson_C01.qxd 08-09-2004 09:45 Page 41

situation. When the results of a formal quantitative model contradict our intuition, we should ask whether there is some important factor that the model has omitted and that is the basis of our intuition. If so, then we must try to extend the formal model to take this factor into account.

In this sense, applied formal modeling work should always be done in a dialog with our intuition. The formal model can help us to see connections that we might not have recognized, but our intuition may point out features that are missing in the model. Thus, although we must accept that our models will always be incomplete simplifications of reality, we must always be prepared to make them more complete (and less simple) when intuition suggests that something important may be missing.

We have seen in this chapter that there may be more than one useful way to model any real situation, and that a decision maker may want to compare different models in terms of how well they describe what is actually known about the real situation. For example, the model of the new salesperson in Figure 1.4 was extended to a more complex model in Figure 1.9, which a manager might consider a better description of this situation. In Chapter 6 (Figure 6.1) we will introduce yet another model that might be an even better fit to the manager's intuitive understanding of the situation. In Figures 1.10 and 1.12 we saw two simple models that are mathematically equivalent, in the sense that they make identical predictions about the probabilities of finding oil in two tracts. Such equivalent models must be considered equally accurate as descriptions of the actual situation, but they may differ in terms of which is more convenient to use.

In any quantitative model there are some parameters that must be assessed or measured from the real situation, and then the model can be used to compute estimates of other quantities that a decision maker might want to know. But different models of the same situation may differ in terms of which quantities are assessed in building the model and which other quantities can be computed from the model. The most convenient model is the one that has parameters easiest for the decision maker to assess and that can be used to compute all the other quantities that are important to the decision maker.

Some parameters in a quantitative model may be very difficult to measure or assess accurately. For example, sellers of some product may not know exactly how many units of this product will be demanded by their customers next year. They may feel that this unknown quantity of demand could be any number within some range of uncertainty. When this range of uncertainty is small, a modeler might just pick some number in this range as a good guess of the quantity and might analyze the model as if there were no uncertainty about this guess. But then the modeler should do sensitivity analysis to see how results might change with other reasonable guesses of this unknown quantity. When the range of uncertainty is large, so that sensitivity analysis suggests that the different reasonable guesses would imply substantially different optimal decisions, then we really need to use models that express our uncertainty.

Myerson_C01.qxd 08-09-2004 09:45 Page 42

As we have seen, uncertainty is expressed in quantitative models by using probabilities. When we feel unable to accurately assess some quantity, like next year's demand, we can instead try to assess the probabilities of different possible values of this quantity. For example, when we cannot say whether next year's demand for our product will be more than 5000 units, we may instead talk about the probability of this demand being more than 5000 units, to describe how likely such outcomes might be. The art of making and using models with probabilities is the subject of this book.

## 1.9 Summary

We focused here on an example of wanting to learn about some unknown quantity (a salesperson's level of skill) by observing other events (the successes and failures in various sales calls) that are influenced by this unknown quantity. We analyzed this problem using spreadsheet simulation models in which one random cell simulates the unknown quantity and other cells simulate the observations that depend on the unknown quantity.

In the context of this problem, we introduced some basic ideas of probability, including: prior probabilities and conditional probabilities $P(A|B)$, independence and conditional independence, and Uniform and Triangular probability distributions. We saw how various conditional probabilities can be used to construct a randomized (Monte Carlo) simulation model and how data from such a model can then be used to estimate other conditional probabilities. In general, the conditional probability of any event $A$ given any other event $B$ satisfies the equation

$$P(A|B) = P(A \cap B)/P(B)$$

This formula implies that to estimate $P(A|B)$ from data of many independent simulations, the number of simulations in which $A$ and $B$ both happened together should be divided by the total number of simulations in which $B$ happened (with or without $A$). A second example about drilling for oil in two tracts was used to illustrate how conditional probabilities can also be computed directly from other conditional probabilities in a probability tree diagram.

This chapter introduced basic keystroke techniques for doing such simulation analysis in spreadsheets using Excel with the Simtools add-in. Excel functions introduced here include RAND, IF, COUNT, SUM, AVERAGE, PERCENTRANK, PERCENTILE, BINOMDIST, AND, and INDEX, and we also introduced the Simtools function TRIANINV. We saw how to get information about these and other technical functions in Excel by using the Insert:Function dialog box. Other basic spreadsheet techniques that were reviewed in this

Myerson_C01.qxd 08-09-2004 09:45 Page 43

chapter include copying and pasting with absolute ($) and relative references in formulas, simulation tables, and column-input data tables.

The RAND( ) function is the heart of all simulation models because every RAND( ) formula returns an independent random variable drawn from the Uniform distribution between 0 and 1. Thus, if $x$ and $y$ are numbers such that $0 \le x \le y \le 1$, then the probability of any RAND( ) being between $x$ and $y$ (after the next recalculation of the spreadsheet) is just the difference $y - x$. That is,

$$P(x \le \text{RAND}( ) \le y) = y - x$$

To compute conditional probabilities from large tables of simulation data, we used a formula-filtering technique in which a column is filled with IF formulas that extract information from a simulation table, returning a nonnumerical value ("..") in data rows that do not match our criterion. The information extracted in such columns can be summarized using statistical functions like COUNT, SUM, AVERAGE, PERCENTRANK, and PERCENTILE, which are designed to ignore nonnumerical entries.

## Exercises

**1.** The Connecticut Electronics Company produces sophisticated electronic modules in production runs of several thousand at a time. It has been found that the fraction of defective modules can be very different in different production runs. These differences are caused by microirregularities that sometimes occur in the electrical current. For a simple first model, we may assume first that there are just two possible values of the defective rate. In about 70% of the production runs, the electric current is regular, in which case every module that is produced has an independent 10% chance of being defective. In the other 30% of production runs, when current is irregular, every module that is produced has an independent 40% chance of being defective.

Testing these modules is quite expensive, so it is valuable to make inferences about the overall rate of defective output based on a small sample of tested modules from each production run.

**(a)** Make a spreadsheet model to study the conditional probability of irregular current given the results of testing 10 modules from a production run. Make a simulation table with data from at least 1000 simulations of your model (where each simulation includes the results of testing 10 modules), and use this table to answer the following questions.

**(b)** What is the probability of finding exactly 2 defective modules among the 10 tested?

**(c)** What is the conditional probability that the current is regular given that 2 defective modules are found among the 10 tested?

**(d)** Make a table showing the estimates of $P$(irregular current| $k$ defectives among 10 tested), for $k$ = 0, 1, 2, . . . . (You may have trouble with some $k$ greater than 7, but the answer in those cases should be clear.)

**2.** Reconsider the Connecticut Electronics problem, but now drop the assumption that each run's defective rate must be one of only two possible values. Managers have observed that due to differences in the electrical current, the defective rate in any production run may be any number between a lower bound of 0% and an upper bound of 50%. They have also observed that the defective rates are more likely to be near 10% than any other single value in this range. So suppose that the defective rate on any production run is drawn from a Triangular distribution with these parameters. We want to quantify what we can infer about the defective rate in the

most recent production run based on the testing of 10 modules from this production run.

**(a)** Before we test any modules, what is the probability of a defective rate less than 0.25 in this production run? For what number $M$ would you say that the defective rate is equally likely to be above or below $M$?

**(b)** Based on data from at least 1000 simulations, what is the probability that we will find exactly 2 defective modules when we test 10 modules from this production run?

**(c)** Suppose that we found exactly 2 defective modules when we tested 10 modules from this production run. Given this observation, what is the conditional probability of a defective rate less than 0.25 in this production run? For what number $M$ would you say that the defective rate in this production run is equally likely to be above or below $M$?

**(d)** How would your answers to part (c) change if we found 0 defective modules when we tested 10 modules from this production run?

**3.** Consider again our basic model of the salesperson who makes 20 sales calls. Before this person makes these calls, we think high skill or low skill is equally likely. If there is high skill, then the probability of making a sale is 2/3 in each call. If there is low skill, then the probability of making a sale is 1/3 in each call. We consider the results of these 20 calls to be conditionally independent given the skill.

**(a)** Suppose we have a policy of promoting the salesperson if there are at least 9 sales in these 20 calls. Based on data from at least 1000 simulations, estimate the following probabilities:

**(a1)** The probability this salesperson will be promoted under this policy

**(a2)** The conditional probability that this person will be promoted under this policy if there is high skill

**(a3)** The conditional probability high skill given that this person is promoted under this policy

**(b)** Make a table showing how the three probabilities that you computed in part (a) would change if the policy were instead to promote the salesperson if there are at least $n$ sales in the 20 calls, for any integer $n$ between 0 and 20.

**4.** Another oil exploration company is negotiating for rights to drill for oil in Tracts C and D, which are near each other (but in a very different area of the world from the Tracts A and B that were discussed in Section 1.6). The company's geological expert has estimated that the probability of finding oil at Tract C is 0.1. The expert has also estimated that the conditional probability of oil at Tract D given oil at Tract C would be 0.4. Finally, it is also estimated that the conditional probability of oil at Tract D given no oil at Tract C would be 0.02.

**(a)** Using the expert's probability assessments, compute the probability of oil at Tract D, and compute the conditional probability of oil at Tract C given oil at Tract D.

**(b)** Suppose that the expert asserts that oil at Tract D is just as likely as oil at Tract C. This assertion suggests that at least one of the three previous assessments must be wrong. Suppose that the expert feels quite confident about the estimate of the probability of C and the estimate of the conditional probability of oil at D given oil at C. To be consistent with this new assertion, how would the conditional probability of oil at D given no oil at C have to be changed?

**5.** Gates & Associates regularly receives proposals from other companies to jointly develop new software programs. Whenever a new proposal is received, it is referred to Gates's development staff for a study to estimate whether the software can be developed as specified in the proposal. In staff reports, proposals are rated as "good" or "not good" prospects. When the development staff evaluates a proposal that is actually feasible, the probability of their giving it a "good" evaluation is 3/4. When the development staff evaluates a proposal that is actually not feasible, the probability of their giving it a "not good" evaluation is 2/3.

Proposals can also be referred to an outside consultant to get a second independent opinion on feasibility of the proposal. In the consultant's reports, proposals are rated as "highly promising" or "not highly promising." When the outside consultant evaluates a proposal that is actually feasible, the probability of it getting a "highly promising" evaluation is 0.95. When the outside consultant evaluates a proposal that is actually not feasible, the probability of it getting a "not highly promising" evaluation is 0.55. Suppose that the reports of the development staff and the outside consultant on a proposal would be conditionally independent, given the actual feasiblity or infeasibility of the proposal.

A new proposal has just come in from Valley Software, a company that has generated some good ideas but also has made more than its share of exaggerated promises. Without any technical evaluation of this new proposal, based only on Valley

Myerson_C01.qxd 08-09-2004 09:45 Page 45

Software's past record, Gates currently figures that the probability of the Valley proposal being feasible is 0.35.

**(a)** Make a probability tree to represent this situation.

**(b)** What is the probability that the staff will report that the Valley proposal is "good"?

**(c)** What is the probability that the outside consultant would report that the Valley proposal is "highly promising"?

**(d)** What is the conditional probability that the Valley proposal is feasible if the staff reports that it is "good"?

**(e)** What is the conditional probability that the outside consultant would report that the Valley proposal is "highly promising" if the staff reports that the proposal is "good"?

**(f)** If the staff reports that the Valley proposal is "good" and the outside consultant also reports that it is "highly promising," then what is the probability that the proposal is feasible?

**(g)** Make a spreadsheet model that simulates this situation, where cell B3 simulates the feasibility of the Valley proposal (1 = feasible, 0 = unfeasible), cell C3 simulates the staff report (1 = good, 0 = not good), and cell D3 simulates the outside consultants report (1 = highly promising, 0 = not highly promising).

**(h)** How might you make a model that is equivalent to the one in part (g), except that the direction of influence between cells B3 and C3 is reversed? (That is, if the formula in cell C3 referred to cell B3 in your answer to part (g), then you should now make cell B3 refer to cell C3 and, to avoid circular references, cell C3 must not refer to cell B3.)

**6.** Random Press is deciding whether to publish a new probability textbook. Let **Q** denote the unknown fraction of all probability teachers who would prefer this new book over the classic textbook that you are now reading. If **Q** were known, then the editor at Random would say that every teacher has a conditionally independent **Q** probability of preferring this new book. Unfortunately, the editor does not know **Q**. In fact, the editor's uncertainty about **Q** can be represented by a uniform distribution over the interval from 0 to 1. That is, up to any given number of digits, the editor thinks that **Q** is equally likely to be any number between 0 and 1. To get more information about **Q**, the editor has paid eight randomly sampled probability teachers to carefully read this new book and compare it to the above-mentioned classic textbook. Three months from now, each of these teachers will send the editor a report indicating whether he or she would prefer this new book.

**(a)** Make a spreadsheet model to simulate the uncertainty about the unknown fraction **Q** and the results that the editor will get from the sample of eight readers. Make a table of data from at least 1000 simulations of your model to answer the following questions.

**(b)** Estimate the conditional probability that **Q** > 0.5 (that is, a majority in the overall population will prefer the new book) given that six out of the eight sampled prefer the new book.

**(c)** Estimate the conditional median value of **Q** given that six out of the eight sampled prefer the new book.

**(d)** Make a table indicating, for each number $k$ from 0 to 8;

**(d1)** The conditional median value for **Q** given that $k$ out of the eight sampled teachers prefer the new book (that is, the number $q$ such that $P(\mathbf{Q} < q | k$ prefer the new book) = 0.50)

**(d2)** The conditional 0.05 percentile for **Q** given that $k$ out of eight prefer the new book (that is, the number $q$ such that $P(\mathbf{Q} < q | k$ prefer the new book) = 0.05)

**(d3)** The conditional 0.95 percentile for **Q** given that $k$ out of eight prefer the new book (that is, the number $q$ such that $P(\mathbf{Q} < q | k$ prefer the new book) = 0.95).

**7.** A student is applying to nine graduate schools. She knows that they are all equally selective, but she figures that there is an independent random element in each school's selection process, so by applying to more schools she has a greater probability of getting into at least one.

She is uncertain about her chances partly because she is uncertain about what kinds of essays might be distinctively impressive, in the sense of being able to attract the attention of admissions officers who read many hundreds of essays per year. Given the rest of her credentials, the student believes that if her application essay was distinctively impressive then, at each school, she would have a probability 0.25 of being admitted, independently of how the other schools responded to her application. But if she knew that her application essay was not distinctively impressive, then she would figure that, at each school, she would have a probability 0.05 of being admitted, independently of how the other schools responded to her application.

Because of her uncertainty about what kind of essay would be distinctively impressive to admissions

Myerson_C01.qxd 08-09-2004 09:45 Page 46

officers, she has actually drafted two different application essays. Essay #1 focuses on her experiences in the Peace Corps after college; essay #2 focuses on her experiences as head of the student laundry services during college. If either essay really is distinctively impressive, then it would have this intrinsic property wherever it was sent. She does not feel sure about the quality of either essay, but she is more optimistic about essay #1. In her beliefs, there is probability 0.5 that essay #1 is a distinctively impressive essay, and there is probability 0.3 that essay #2 is a distinctively impressive essay, each independently of the other.

Although she would be very glad to attend any of these nine schools, the student considers three of the schools to be somewhat less preferred than the other six because of their geographical locations. So she has definitely decided to use essay #1 in the applications to her six more-desired schools. But she is undecided about what to do with the applications for the three less-desired schools: She could use the same essay #1 for these 3 schools as well, or instead she could use essay #2 in her applications to these three schools.

**(a)** Build a simulation model to study this situation. You can use data from at least 1000 simulations to estimate answers to the following questions (b)–(e).

**(b)** What is the probability that she will be rejected by all of her six more-desired schools when she uses essay #1 in each of these applications?

**(c)** What would be the conditional probability of essay #1 being distinctively impressive given that she was rejected by all of the six more-desired schools?

**(d)** If she also used essay #1 in her applications to the three less-desired schools, then what would be her conditional probability of getting accepted by at least one of these three schools with essay #1 given that she was rejected by all of the six more-desired schools?

**(e)** If she instead used essay #2 in the applications to the three less-desired schools, then what would be her conditional probability of getting accepted by at least one of these three schools with essay #2, given that she was rejected by all of the six more-desired schools (with the other essay)?

**8.** Consider again the oil exploration example in Section 1.6.

**(a)** Compute the conditional probability of favorable strata given that neither Tract A nor Tract B has any oil.

**(b)** Compute the conditional probability of favorable strata given that exactly one of the two tracts has oil.

**(c)** Compute the conditional probability of favorable strata given that both of the tracts have oil.

**(d)** Make a probability tree with three levels of branches that represents this oil exploration example, such that the first two levels are the same as in Figure 1.12 and the third-level branches represent the events of favorable strata existing or not existing in this area.

**(e)** In Figure 1.12, what formula could you enter into cell X23 to simulate "Favorable Strata?" so that cells X23:Z23 would be a simulation model that is mathematically equivalent to the simulation model in cells B3:D3 of Figure 1.10? (To make this model in X23:Z23, you may also enter formulas into other blank cells of this spreadsheet, but do not change the formulas in cells Y23 and Z23.)

Myerson_C02.qxd 09-09-2004 10:31 Page 47

# 2

# Discrete Random Variables

In this chapter, we focus on one simple example, but in the context of this example we develop most of the technical concepts of probability theory, statistical inference, and decision analysis that are used throughout the rest of the book. This example is very simple in that it involves only one unknown quantity which has only finitely many possible values. That is, in technical terms, this example involves just one discrete random variable.

With just one discrete random variable, we can make a table or chart that completely describes its probability distribution. Among the various ways of picturing a probability distribution, the most useful in this book will be the *inverse cumulative distribution* chart. After introducing such charts and explaining how to read them, we show how this inverse cumulative distribution can be used to make a simulation model of any random variable.

Next we introduce the two most important summary measures of a random variable's probability distribution: its expected value and standard deviation. These two summary measures can be easily computed for a discrete random variable, but we also show how to estimate them from simulation data. The expected value of a decision maker's payoff will have particular importance throughout this book as a criterion for identifying optimal decisions under uncertainty.

Later we will consider more complex models with many random variables, some of which may have infinitely many possible values. For such complex models, we may not know how to compute expected values and standard deviations directly, but we will still be able to estimate these quantities from simulation data by the methods that are introduced in this chapter. We introduce these methods here with a simple one-variable model because when you first learn to compute statistical estimates from simulation data, it is instructive to begin with a case where you can compare these estimates to the actual quantities being estimated.

Myerson_C02.qxd 09-09-2004 10:31 Page 48

## CASE ■ Superior Semiconductor (Part A)

Peter Suttcliff, an executive vice president at Superior Semiconductor, suspected that the time might be right for his firm to introduce the first integrated T-regulator device using new solid-state technology. This new product seemed the most promising of the several ideas that had been suggested by the head of Superior's industrial products division. So Suttcliff asked his staff assistant Julia Eastmann to work with Superior's business marketing director and the chief production engineer to develop an evaluation of the profit potential from this new product.

According to Eastmann's report, the chief engineer anticipated substantial fixed costs for engineering and equipment just to set up a production line for the new product. Once the production line was set up, however, a low variable cost per unit of output could be anticipated, regardless of whether the volume of output was low or high. Taking account of alternative technologies available to the potential customers, the marketing director expressed a clear sense of the likely selling price of the new product and the potential overall size of the market. But Superior had to anticipate that some of its competitors might respond in this area by launching similar products. To be specific in her report, Eastmann assumed that three other competitive firms would launch similar products, in which case Superior should expect one-fourth of the overall market.

Writing in the margins of Eastmann's report, Suttcliff summarized her analysis as follows:

- Superior's fixed set-up cost to enter the market: $26 million
- Net present value of revenue minus variable costs in the whole market: $100 million
- Superior's predicted market share, assuming three other firms enter: 1/4
- Result: predicted net loss for Superior: ($1 million)

"Your estimates of costs and total market revenues look reasonably accurate," Suttcliff told Eastmann. "But your assumption about the number of other firms entering to share this market with us is just a guess. I can count five other semiconductor firms that might seriously consider competing with us in this market. In the worst possible scenario, all five of these firms could enter the market, although that is rather unlikely. There is no way that we could keep this market to ourselves for any length of time, and so the best possible scenario is that only one other firm would enter the market, although that is also rather unlikely. I would agree with you that the most likely single event is that three other firms would enter to share the market with us, but that event is only a bit more likely than the possibilities of having two other firms enter, or having four other firms enter. If there were only two other entrants, it could change a net loss to a net profit. So there is really a lot of uncertainty about this situation, and your analysis might be more convincing if you did not ignore it."

"We can redo the analysis in a way that takes account of the uncertainty by using a probabilistic model," Eastmann replied. "The critical step is to assess a probability distribution for the unknown number of competitors who would enter this market with us. So I should try to come up with a probability distribution that summarizes the beliefs that you expressed." Then after some thought, she wrote the following table and showed it to Suttcliff:

| **k** | ***Probability that* k *other competitors enter*** |
|---|---|
| 1 | 0.10 |
| 2 | 0.25 |
| 3 | 0.30 |
| 4 | 0.25 |
| 5 | 0.10 |

Suttcliff studied the table. "I guess that looks like what I was trying to say. I can see that your probabilities sum to 1, and you have assigned higher

Myerson_C02.qxd 09-09-2004 10:31 Page 49

probabilities to the events that I said were more likely. But without any statistical data, is there any way to test whether these are really the right probability numbers to use?"

"In a situation like this, without data, we have to use subjective probabilities," Eastmann explained. "That means that we can only go to our best expert and ask him whether he believes each possible event to be as likely as our probabilities say. In this case, if we take you as best expert about the number of competitive entrants, then I could test this probability distribution by asking you questions about your preferences among some simple bets. For example, I could ask you which you would prefer among two hypothetical lotteries, where the first lottery would pay you a $10,000 prize if exactly one other firm entered this market, whereas the second lottery would pay the same $10,000 prize but with an objective 10% probability. Assuming that you had no further involvement with this project, you should be indifferent between these two hypothetical lotteries if your subjective probability of one other firm entering is 0.10, as my table says. If you said that you were not indifferent, then we would try increasing or decreasing the first probability in the table, depending on whether you said that the first or second lottery was preferable. Then we could test the other probabilities in the table by similar questions. But if we change any one probability in my table, then at least one other probability must be changed, because the probabilities of all the possible values of the unknown quantity must add up to 1."

Suttcliff looked again at the table of probabilities for another minute or two, and then he indicated that it seemed to be a reasonable summary of his beliefs.

## 2.1 Unknown Quantities in Decisions Under Uncertainty

Uncertainty about numbers is pervasive in all management decisions. How many units of a proposed new product will we sell in the year when it is introduced? How many yen will a dollar buy in currency markets a month from today? What will be the closing Dow Jones Industrial Average on the last trading day of this calendar year? Each of these numbers is an *unknown quantity*. If our profit or payoff from a proposed strategy depends on such unknown quantities, then we cannot compute this payoff without making some prediction about them.

A common approach to such problems is to assess your best estimate for each of these unknown quantities and use these estimates to compute the bottom-line payoff for each proposed strategy. Under this *method of point estimates*, the optimal strategy is considered to be the one that gives you the highest payoff when all unknown quantities are equal to your best estimates.

But there is a serious problem with this method of point estimates: It completely ignores your uncertainty. In this book, we study ways to incorporate uncertainty into the analysis of decisions. Our basic method will be to

Myerson_C02.qxd 09-09-2004 10:31 Page 50

assess probability distributions for unknown quantities, and then to create random variables that simulate these unknown quantities in spreadsheet simulation models. In the general terminology of decision analysis, the term *random variable* is often taken by definition to mean the same thing as *unknown quantity*. But as a matter of style here, we will generally reserve the term *unknown quantity* for unknowns in the real world, and *random variable* will generally be used for values in spreadsheets that are unknown because they depend on unknown RAND values.

To illustrate these ideas, we consider the Superior Semiconductor case (Part A). In this case, we are faced with a decision about whether our company should introduce a proposed new product. It is estimated that the fixed cost of introducing this new product will be $26 million. The total value of the market (price minus variable unit costs, multiplied by total demand) is estimated to be $100 million. It is also estimated that three other firms will enter this market and share it equally with us. Thus, by the method of point estimates, we get a net profit (in millions of dollars) of $100/(3 + 1) - 26 = -1$, which suggests that this product should not be introduced. But all the quantities in this calculation (fixed cost, value of the market, number of competitive entrants) are really subject to some uncertainty. We will see, however, that when uncertainty is properly taken into account, the new product may be recognized as worth introducing.

The analysis in Part A of this case focuses on just one of these unknowns: the number of entrants. Uncertainty about other quantities (fixed cost, value of the market) is ignored until the end of this chapter, but it will be considered in more detail in Chapter 4. By focusing on just this one unknown quantity for now, we can simplify the analysis as we introduce some of the most important fundamental ideas of probability theory.

## 2.2 Charting a Probability Distribution

We use probability distributions to describe people's beliefs about unknown quantities. When an unknown quantity has only finitely many possible values, we can describe it using a *discrete probability distribution*. (Continuous probability distributions, for unknown quantities with infinitely many possible values, will be discussed in Chapter 4.) A discrete probability distribution can be presented in a table that lists the possible values of the unknown quantity and the probability of each possible value.

In the Superior Semiconductor case, the number of competitors who will enter the market is a quantity that is unknown to the company's decision makers, and they believe that this unknown quantity could be any number from 1 to 5. In our mathematical notation, let **K** denote this unknown number of competitors who will enter this market (we are following a mathematical tradition of representing unknown quantities by

Myerson_C02.qxd 09-09-2004 10:31 Page 51

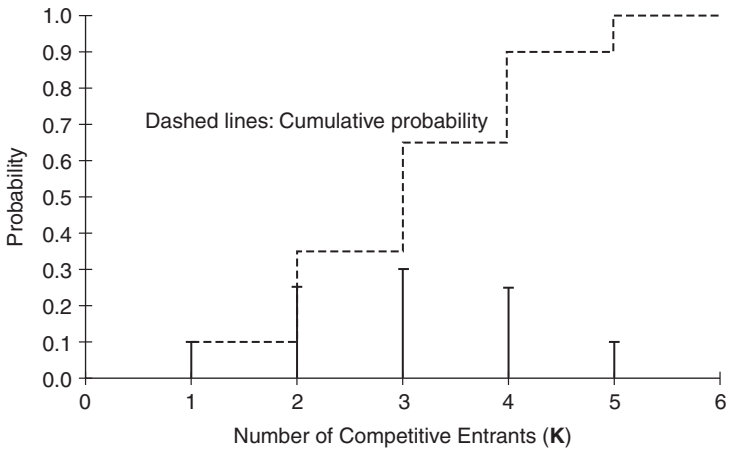


FIGURE **2.1**
**Discrete probability distribution for the number of entrants (K)**

boldface letters). Then the decision maker's beliefs about this unknown quantity **K** are described in the case by a discrete probability distribution such that

$$P(\mathbf{K} = 1) = 0.10, \quad P(\mathbf{K} = 2) = 0.25, \quad P(\mathbf{K} = 3) = 0.30,$$
$$P(\mathbf{K} = 4) = 0.25, \quad P(\mathbf{K} = 5) = 0.10$$

Here, for any number $k$, the mathematical expression $P(\mathbf{K} = k)$ denotes the probability that the unknown quantity **K** is equal to the value $k$. This probability distribution is summarized by a chart in Figure 2.1. Figure 2.1 actually displays this probability distribution in two different ways. The five solid bars show the probabilities of the five points on the horizontal axis that represent possible values of the unknown quantity **K**. Such point probability bars are the most common way of exhibiting a discrete probability distribution. But Figure 2.1 also contains a dashed line that shows cumulative probability values, which we must now explain.

A *cumulative probability* of the unknown quantity **K** at a number $k$ is the probability of **K** being below the value $k$. It is a question of mathematical convention as to whether "cumulative probability of **K** at 2" should be precisely defined as $P(\mathbf{K} < 2)$ or $P(\mathbf{K} \le 2)$, that is, 0.10 or 0.10 + 0.25 = 0.35 in this case. In most books the latter definition is used. But in this book, let us eclectically embrace both definitions and everything in between and define a *cumulative probabilities of* ***K*** *at* $k$ to include $P(\mathbf{K} < k)$ and $P(\mathbf{K} \le k)$ and every number in between. So in this terminology, any number between 0.10 and 0.35 can be called a cumulative probability of **K** at 2 in this example. But notice that such ambiguity only occurs at numbers that have positive probability for the random variable. For example, the cumulative probability of **K** at 2.5 is 0.35, because $P(\mathbf{K} < 2.5) = P(\mathbf{K} \le 2.5) = 0.10 + 0.25 = 0.35$.

The dashed curve in Figure 2.1 shows the *cumulative probabilities* for each number $k$ on the horizontal axis. For any number $k$ between 1 and 2, the height of the dashed cumulative probability curve in Figure 2.1 is 0.10,

Myerson_C02.qxd 09-09-2004 10:31 Page 52

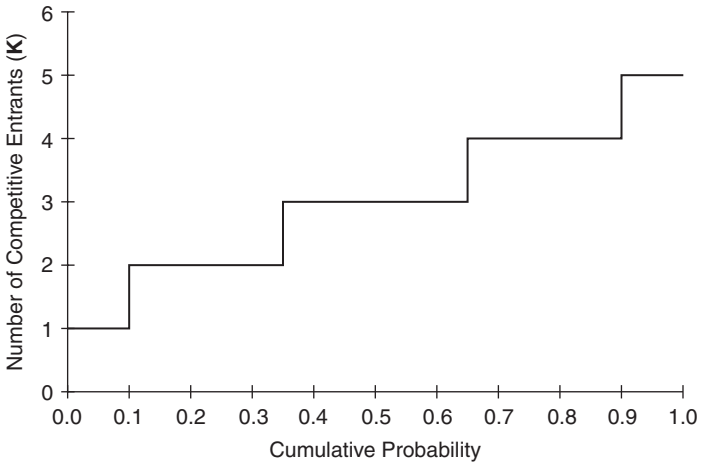


FIGURE **2.2**
**Inverse cumulative probability curve for number of entrants (K)**

because $P(\mathbf{K} < k) = P(\mathbf{K} \le k) = P(\mathbf{K} = 1) = 0.10$ when $1 < k < 2$. For any number $k$ between 2 and 3, the height of the dashed cumulative probability curve in Figure 2.1 is 0.35, because

$$P(\mathbf{K} < k) = P(\mathbf{K} \le k) = P(\mathbf{K} = 1) + P(\mathbf{K} = 2) = 0.10 + 0.25 = 0.35$$
$$\text{when } 2 < k < 3$$

The unknown quantity **K** is sure to be less than 6, and so the height of the cumulative probability curve at 6 is $1 = P(\mathbf{K} < 6)$. The unknown quantity **K** is sure to be not less than 0, and so the height of the cumulative probability curve at 0 is $0 = P(\mathbf{K} < 0)$.

In Figure 2.1, the dashed cumulative probability curve has vertical jumps (representing multiple cumulative probabilities from $P(\mathbf{K} < k)$ to $P(\mathbf{K} \le k)$) exactly where the point probability bars occur, and the height of each vertical jump is the same as the height of the corresponding point probability bar. For example, the dashed cumulative probability curve jumps from 0.10 to 0.35 above the value 2 on the horizontal axis of Figure 2.1, which corresponds to the fact that $P(\mathbf{K} = 2) = 0.25 = 0.35 - 0.10$. So the cumulative probability curve tells us everything about the probability distribution that we could learn from the point probability bars. This observation is important, because we will find that cumulative probability curves are generally more useful for describing probability distributions than point probability bars (which cannot be applied to continuous probability distributions, where there are infinitely many possible values).

Actually, we will find it most useful to invert the cumulative probability distribution, turning the dashed line from Figure 2.1 on its side, with cumulative probabilities on the horizontal axis and possible values of **K** on the vertical axis. Such an *inverse cumulative probability curve* is shown in Figure 2.2. Once you learn how to read it, you can find the discrete probabilities of all possible values of **K** from this inverse cumulative probability curve.

Myerson_C02.qxd 09-09-2004 10:31 Page 53

For example, the inverse cumulative probability curve has height 2 over the interval of probabilities from 0.10 to 0.35, which tells us that the point probability of the value 2 is $P(\mathbf{K} = 2) = 0.35 - 0.10 = 0.25$. The height of the inverse cumulative probability curve goes from 1 to 5 because these are the lowest and highest possible values of the unknown **K**. (Because we will generally draw our cumulative charts in this inverse orientation, it should also be acceptable here to drop the word *inverse* and simply refer to a chart like Figure 2.2 as a cumulative probability chart.)

For any numbers $q$ and $k$, if $q$ is a cumulative probability of an unknown quantity **K** at $k$, then we may also say that $k$ is a *q percentile value* of **K**. So in this example, the 0.2 percentile value of **K** is 2, but any number from 2 to 3 could be called a 0.35 percentile value of **K**.

## 2.3 Simulating Discrete Random Variables

When we say that a random variable in a spreadsheet *simulates* (or represents) some unknown quantity in real life, we mean that any event for this simulated random variable is, from the perspective of our current information and beliefs, just as likely as the same event for the real unknown quantity. If the unknown number of competitive entrants has a probability 0.10 of equaling 1, for example, then the random variable in the spreadsheet should also have probability 0.10 of equaling 1 after the next recalculation of the spreadsheet. For any number $k$, the probability that the random variable will be less than $k$ after the next recalculation should be the same as the probability that the real unknown quantity is less than $k$.

In our spreadsheets, all our random variables are constructed as functions of RAND( ) values. So recall how the RAND( ) function operates: Given any two numbers $x$ and $y$ such that $0 \le x \le y \le 1$, the event that some particular RAND( ) in a spreadsheet formula will take a value between $x$ and $y$ after the next recalculation is equal to the difference $y - x$, that is,

$$P(x \le \text{RAND}( ) \le y) = y - x$$

For example, the probability that a RAND( ) will be between 0 and 0.10 is 0.10, and the probability that the RAND( ) will be between 0.10 and 0.35 is $0.35 - 0.10 = 0.25$.

Let us now try to make a spreadsheet formula that depends on one RAND( ) such that the value of our formula will simulate the unknown quantity **K** from the Superior Semiconductors case. To be specific let us enter =RAND( ) into cell A13, as shown in Figure 2.3. Into another cell we want to enter a formula that depends on the RAND( ) in cell A13 in such a way that its value is 1 with probability 0.10, 2 with probability 0.25, 3 with probability 0.30, 4 with probability 0.25, and 5 with probability 0.10, just like the unknown quantity **K**. One way to get these probabilities is to use

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | A DISCRETE RANDOM VARIABLE FROM THE "SUPERIOR SEMICONDUCTOR" CASE | | | | | | | |
| 2 | | K = (unknown number of competitors entering market). | | | | | | |
| 3 | | Little k = (possible value of big K). | | | | | | |
| 4 | P(K<k) | k | P(K=k) | | | | (for B16) | |
| 5 | 0.00 | 1 | 0.10 | | | | 0 | |
| 6 | 0.10 | 2 | 0.25 | | | | 2 | |
| 7 | 0.35 | 3 | 0.30 | | | | 0 | |
| 8 | 0.65 | 4 | 0.25 | | | | 0 | |
| 9 | 0.90 | 5 | 0.10 | | | | 0 | |
| 10 | | | 1 | | | | | |
| 11 | | | | | | | | |
| 12 | | Simulated value | | | | | | |
| 13 | 0.32748 | 2 | | | | | | |
| 14 | | 2 | | | | | | |
| 15 | | 2 | | | | | | |
| 16 | | 2 | | | | | | |
| 17 | FORMULAS | | | | | | | |
| 18 | A6. =A5+C5 | | | | | | | |
| 19 | A6 copied to A6:A9 | | | | | | | |
| 20 | C10. =SUM(C5:C9) | | | | | | | |
| 21 | A13. =RAND() | | | | | | | |
| 22 | G5. =IF($A$13<A6,B5,0) | | | | | | | |
| 23 | G6. =IF(AND(A6<=$A$13,$A$13<A7),B6,0) | | | | | | | |
| 24 | G6 copied to G6:G8 | | | | | | | |
| 25 | G9. =IF(A9<=$A$13,B9,0) | | | | | | | |
| 26 | B13. =SUM(G5:G9) | | | | | | | |
| 27 | B14. =DISCRINV(A13,B5:B9,C5:C9) | | | | | | | |
| 28 | B15. =IF(A13<A6,B5,IF(A13<A7,B6,IF(A13<A8,B7,IF(A13<A9,B8,B9)))) | | | | | | | |
| 29 | B16. =VLOOKUP(A13,A5:B9,2) | | | | | | | |

FIGURE **2.3**
**Simulation with a discrete probability distribution (four equivalent formulas)**

some formula that returns a value that depends on the RAND( ) in cell A13, as follows:

- The value is 1 when A13 is less than 0.10 (an event having probability 0.10)
- The value is 2 when A13 is between 0.10 and 0.10 + 0.25 = 0.35 (having probability 0.25)
- The value is 3 when A13 is between 0.35 and 0.35 + 0.30 = 0.65 (having probability 0.30)
- The value is 4 when A13 is between 0.65 and 0.65 + 0.25 = 0.90 (having probability 0.25)
- The value is 5 when A13 is greater than 0.90 (having probability 1 – 0.90 = 0.10)

Notice that we have gotten all the probabilities we want with a collection of disjoint intervals that exhaust all the possible RAND( ) values because the probabilities of all possible values of **K** sum to 1. This is why all the possible values in any discrete probability distribution must have probabilities that sum to exactly 1.

Figure 2.3 shows several different ways to create an Excel formula that has these properties. The possible values and corresponding probabilities for **K** are listed in the ranges B5:B9 and C5:C9 of Figure 2.3. The (low) cumulative probabilites of the various possible values are computed from

Myerson_C02.qxd 09-09-2004 10:31 Page 55

this probability distribution in cells A5:A9 of Figure 2.3. Cell G5 contains the formula

=IF($A$13<A6,B5,0)

So cell G5 equals 1 (B5) if the RAND( ) in A13 satisfies the inequality A13 < A6 = 0.10, and G5 equals 0 otherwise. Cell G6 contains the formula

=IF(AND(A6<=$A$13,$A$13<A7),B6,0)

Excel's AND function returns the logical value TRUE if all of its parameters are TRUE; otherwise AND returns the logical value FALSE. So by the IF function in this formula, cell G6 equals 2 (B6) when the RAND( ) in A13 satisfies the inequalities $0.10 = A6 \le A13 < A7 = 0.35$, and G6 equals 0 otherwise. (Excel uses <= to denote less than or equal to, and uses >= to denote greater than or equal to.) The other formulas in cells G5:G9 are constructed similarly so that one of these cells will equal the possible value of **K** that would be designated by the RAND( ) in cell A13 under the preceding rule, while all the other cells in G5:G9 will equal 0. Thus, the formula =SUM(G5:G9) in cell B13 returns the random variable that we wanted.

Simtools provides a function called DISCRINV to accomplish this same calculation more easily. As you can verify by consulting the Insert-Function dialog box, the DISCRINV function takes three parameters. The first parameter (which is called "randprob" in the Insert-Function dialog box) should simply be a RAND( ). The second parameter (called "values") should be a range that lists the possible values of our discrete random variable. The third parameter (called "probabilities") should be another range that has the same size as the values range and that lists the corresponding probabilities of these values. Then the formula

DISCRINV(RAND( ), values, probabilities)

returns a random variable that has values and discrete probabilities as listed in these ranges. In this spreadsheet, the formula =DISCRINV(A13,B5:B9,C5:C9) in cell B14 returns a random variable that depends on the RAND( ) in A13 according to the rule that we described earlier, and so the value of cell B14 is always the same as cell B13 for any value of the RAND( ) in cell A13.

The cell B15 contains a rather ugly formula that does the job in one line:

=IF(A13<A6,B5,IF(A13<A7,B6,IF(A13<A8,B7,IF(A13<A9,B8,B9))))

This formula is not recommended. First, when a formula is so complicated your chances of typing it incorrectly are very high. Second, this method cannot be applied to larger problems, because Excel may refuse to consider formulas where functions are nested more than eight parentheses deep. Cell B16 illustrates another, more elegant way to compute the same random variable using the Excel function VLOOKUP.

At this point you may wonder why the Simtools function for simulating a discrete random variable should be called DISCRINV. The letters DISC obviously come from the word *discrete*, but what does INV signify? To see

Myerson_C02.qxd 09-09-2004 10:31 Page 56

the answer, notice that this function has been designed to return the value 1, 2, 3, 4, or 5, depending on the value of the RAND( ) that is its first parameter, and the points where the function's value changes are 0.10, 0.10 + 0.25 = 0.35, 0.35 + 0.30 = 0.65, and 0.65 + 0.25 = 0.90. You have seen a function like this before: It is the inverse cumulative function shown in Figure 2.2. Indeed, the INV in our function's name is short for *inverse cumulative*.

In general, any random variable can be simulated by processing a RAND( ) through the inverse cumulative probability function. If a function G(•) is the inverse cumulative probability function of an unknown quantity **X,** then, for any numbers $k$ and $q$, we get $G(q) < k$ when $q < P(\mathbf{X} < k)$. So for any number $k$, the random variable $G$(RAND( )) will be less than $k$ when the RAND( ) is in the interval from 0 to $P(\mathbf{X} < k)$, which occurs with probability $P(\mathbf{X} < k)$. So $G$(RAND( )) has the same probability distribution as **X**.

The bottom-line quantity of interest in the Superior Semiconductor case is profit. If we assume that the $100 million total value of the market will be shared equally by Superior Semiconductor and its **K** competitors, then, after deducting $26 million of fixed costs, Superior Semiconductor's profit (in millions of dollars) should depend on the number of competitors **K** according to the formula

$$\text{Profit} = 100/(1 + \mathbf{K}) - 26$$

So if **K** equals 1, then profit is 100/(1 + 1) – 26 = 24 ($million), but if **K** equals 5, then profit is 100/(1 + 5) – 26 = –9.33. So profit is also an unknown quantity, with a discrete probability distribution as shown in the following table:

| *Competitors* | *Profit* | *Probability* |
|---|---|---|
| 1 | 24 | 0.10 |
| 2 | 7.33 | 0.25 |
| 3 | –1 | 0.30 |
| 4 | –6 | 0.25 |
| 5 | –9.33 | 0.10 |

Figure 2.4 shows two attempts to make a model in which the number of competitors and the resulting profit for Superior Semiconductor are both simulated.

The range B16:E18 in Figure 2.4 contains Model 2, which illustrates one of the most common errors that students make in simulation modeling. In Model 2, the number of competitors and the profit are simulated in cells B18 and E18, respectively, with two DISCRINV formulas that depend on separate RANDs. The result is that profit is independent of the number of competitors in Model 2, which is wrong. For example, Figure 2.4 shows a realization of these random variables such that the simulated number of competitors in B18 is 4 and the simulated profit in E18 is 24 ($million). But if there were really four competitive entrants, then Superior Semiconductor's profit would be 100/(1 + 4) – 26 = –6.

The range B12:E14 in Figure 2.4 contains Model 1, which is done correctly. In this model, the number of competitors is simulated by a

Myerson_C02.qxd 09-09-2004 10:31 Page 57

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | "SUPERIOR SEMICONDUCTOR" CASE | | | | 26 | FixedCost | |
| 2 | | | | | 100 | MarketValue | |
| 3 | Let K = (unknown number of competitors entering market) | | | | | | |
| 4 | | k | P(K=k) | | Profit | (in $millions) | |
| 5 | | 1 | 0.10 | | 24.00 | | |
| 6 | | 2 | 0.25 | | 7.33 | | |
| 7 | | 3 | 0.30 | | -1.00 | | |
| 8 | | 4 | 0.25 | | -6.00 | | |
| 9 | | 5 | 0.10 | | -9.33 | | |
| 10 | | | 1 | sum | | | |
| 11 | | | | | | | |
| 12 | | Model 1 (correct) | | | | | |
| 13 | | #Competitors entering | | | Profit | | |
| 14 | | 2 | | | 7.33 | | |
| 15 | | | | | | | |
| 16 | | Model 2 (WRONG!!!) | | | | | |
| 17 | | #Competitors entering | | | Profit | | |
| 18 | | 4 | | | 24 | | |
| 19 | | | | | | | |
| 20 | FORMULAS FROM RANGE A1:F18 | | | | | | |
| 21 | E5. =$E$2/(1+B5)-$E$1 | | | | | | |
| 22 | E5 copied to E5:E9 | | | | | | |
| 23 | C10. =SUM(C5:C9) | | | | | | |
| 24 | B14. =DISCRINV(RAND(),B5:B9,$C$5:$C$9) | | | | | | |
| 25 | E14. =$E$2/(1+B14)-$E$1 | | | | | | |
| 26 | B18. =DISCRINV(RAND(),B5:B9,$C$5:$C$9) | | | | | | |
| 27 | E18. =DISCRINV(RAND(),E5:E9,$C$5:$C$9) | | | | | | |

FIGURE **2.4**
**Making a simple simulation model of competitors and profit**

DISCRINV formula in cell B14, and the profit is simulated in cell E14 as a function of the simulated number of competitors by the formula =$E$2/(1+B14)–$E$1 (where E2 contains the market value 100 and E1 contains the fixed cost 26). Thus, Model 1 always displays the correct relationship between the number of competitors and the profit.

In general, a simulation model is a good representation of a real situation if our uncertainty about the next recalculated values of the random variables in the model (that is, what they will be after we next press [F9]) is the same as our uncertainty about the corresponding unknown quantities in the real situation. Fancy formulas in a spreadsheet obviously cannot be asked to magically return the actual values of real-world quantities that we are unable to observe or measure by other means. Instead, what we must ask of our simulation models is that their formulas should express our beliefs about the real unknown quantities, in the sense that our beliefs about the next recalculated values of these formulas are the same as our beliefs about the real unknown quantities.

## 2.4 Expected Value and Standard Deviation

We have seen how a probability distribution can be used to describe our beliefs about an unknown quantity that has finitely many possible values and how to represent such a probability distribution by a chart or simulation

Myerson_C02.qxd 09-09-2004 10:31 Page 58

model. But when there are many possible values, the probability distribution may be quite complicated. In such cases, we may want to describe the overall pattern of a probability distribution by a few summary numbers that people could interpret more easily than some complicated chart or some simulated random variable that jumps around whenever [F9] is pressed.

There are many formulas that people have used to generate summary measures of probability distributions (expected value, median, mode, standard deviation, mean absolute error, etc.). Each of these formulas has some drawbacks and limitations, because it is impossible to perfectly summarize everything we want to know about every probability distribution by just a couple of simple numbers. But two summary measures have been found particularly useful and will be emphasized throughout this book: the *expected value* and the *standard deviation*.

The *expected value* or *mean* of an unknown quantity **X** may be denoted by $E(\mathbf{X})$ or $\mu_{\mathbf{X}}$, and it is defined by the formula

$$E(\mathbf{X}) = \mu_{\mathbf{X}} = \sum_x P(\mathbf{X} = x) * x$$

where the summed terms include $P(\mathbf{X} = x) * x$ for all numbers $x$ that are possible values of the unknown quantity **X**. For example, in the Superior Semiconductor case, the expected value of the unknown number of competitors, **K**, is

$$E(\mathbf{K}) = 0.10*1 + 0.25*2 + 0.30*3 + 0.25*4 + 0.10*5 = 3$$

Similarly, if we let $\mathbf{Y} = 100/(1 + \mathbf{K}) - 26$ denote the profit in this case, then the expected value of profit is

$$E(\mathbf{Y}) = 0.10*24 + 0.25*7.33 + 0.30*(-1) + 0.25*(-6) + 0.10*(-9.33)$$
$$= 1.5$$

These calculations are illustrated in cells B13 and F13 of Figure 2.5, where the Excel function SUMPRODUCT is used. When "range1" and "range2" denote two ranges that have the same numbers of rows and columns in a spreadsheet, the Excel formula SUMPRODUCT(range1, range2) multiplies the values of each pair of corresponding cells in these ranges (starting with the top-left cell of range 1 multiplied by the top-left cell in range 2) and then adds up all of these multiplicative products. Thus, when the possible values of **K** are listed in the range B5:B9 and the corresponding probabilities are listed in C5:C9, the expected number of competitors $E(\mathbf{K})$ can be returned by the formula

```
=SUMPRODUCT(B5:B9,$C$5:$C$9)
```

in cell B13. Similarly, when the corresponding profit levels are computed in the range F5:F9, the expected profit $E(\mathbf{Y})$ can be returned by the formula

```
=SUMPRODUCT(F5:F9,$C$5:$C$9)
```

in cell F13 of Figure 2.5.

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | DISCRETE RANDOM VARIABLES FROM THE "SUPERIOR SEMICONDUCTOR" CASE | | | | | | | |
| 2 | | K = (unknown number of competitors entering market). | | | | | | |
| 3 | | Little k = (possible value of big K). | | | | | | |
| 4 | | k | P(K=k) | (k-E(K))^2 | | Y=Profit | | |
| 5 | | 1 | 0.1 | 4 | | 24 | | |
| 6 | | 2 | 0.25 | 1 | | 7.333333 | | |
| 7 | | 3 | 0.3 | 0 | | -1 | | |
| 8 | | 4 | 0.25 | 1 | | -6 | | |
| 9 | | 5 | 0.1 | 4 | | -9.33333 | | |
| 10 | | | 1 | | | | | |
| 11 | | | | | | | | |
| 12 | Mean or E(K) | | Stdev(K) | | | E(Y) | Stdev(Y) | |
| 13 | | 3 | 1.14018 | 1.140175 | | 1.5 | 9.31695 | |
| 14 | | | | | | | | |
| 15 | FORMULAS FROM RANGE A1:G13 | | | | | | | |
| 16 | D5. =(B5-$B$13)^2 | | | | | | | |
| 17 | D5 copied to D5:D9 | | | | | | | |
| 18 | F5. =100/(1+B5)-26 | | | | | | | |
| 19 | F5 copied to F5:F9 | | | | | | | |
| 20 | C10. =SUM(C5:C9) | | | | | | | |
| 21 | B13. =SUMPRODUCT(B5:B9,$C$5:$C$9) | | | | | | | |
| 22 | C13. =STDEVPR(B5:B9,$C$5:$C$9) | | | | | | | |
| 23 | D13. =SUMPRODUCT(D5:D9,C5:C9)^0.5 | | | | | | | |
| 24 | F13. =SUMPRODUCT(F5:F9,$C$5:$C$9) | | | | | | | |
| 25 | G13. =STDEVPR(F5:F9,$C$5:$C$9) | | | | | | | |

FIGURE **2.5**
**Expected values and standard deviations of discrete random variables**

The expected value is interpreted as a measure of the center of a probability distribution. There will generally be some possible values that are higher than the expected value and other possible values that are lower. Notice, however, that the expected value of an unknown quantity is not necessarily itself a possible value of the unknown quantity. In this example, the expected value of **K** ($E(\mathbf{K}) = 3$) happens to be a possible value of **K**. But the expected value of **Y** ($E(\mathbf{Y}) = 1.5$) is not among the possible values of the unknown profit **Y**. So the term *expected* is being used here in a technical sense that may be different from common English usage of the word. In each case, however, the expected value could be reasonably described (in some intuitive sense) as "near the center" of the possible values of the unknown quantity.

Notice also that the expected profit in this example is different from the profit that occurs at the expected number of competitors, even though profit here is a function that depends on the number of competitors. When the number of competitors **K** is 3, which is $E(\mathbf{K})$, the corresponding profit is $\mathbf{Y} = 100/(1 + 3) - 26 = -1$, but $E(\mathbf{Y}) = 1.5$. More generally, if **X** is an unknown quantity and $f(\bullet)$ is any function, the function $f$ evaluated at $E(\mathbf{X})$ may be different from the expected value of the unknown quantity $f(\mathbf{X})$, that is,

$f(E(\mathbf{X}))$ may be different from $E(f(\mathbf{X}))$

The only case where $f(E(\mathbf{X}))$ is guaranteed to equal $E(f(\mathbf{X}))$ is when $f$ is a linear function.

Myerson_C02.qxd 09-09-2004 10:31 Page 60

The erroneous assumption that $f(E(\mathbf{X}))$ ought to be the same as $E(f(\mathbf{X}))$ has been called the *fallacy of averages*. This error seems to arise in people's minds because, after computing the expected value of an unknown quantity, there is a temptation to simplify the world by assuming that the unknown quantity will be equal to its expected value. For example, in the Superior Semiconductor case, we might be tempted to assume that the number of competitors will be three for sure, in which case the profit would be 100/(1 + 3) – 26 = –1 (a loss of $1 million). But when the uncertainty is not assumed away, we actually get a positive expected profit of $1.5 million. This positive expected value of profits should seem intuitively reasonable when you notice that the positive profits (24 and 7.33) that could be generated by a number of competitors less than three are significantly larger in absolute value than the equally likely negative profits (–9.33 or –6) that could be generated by a number of competitors greater than three. This result is in turn caused by a kind of nonlinearity in our profit function (where a decrease of **K** below 3 would increase profit by more than a similar increase of **K** above 3 would decrease profit).

Now that we have the expected value as a summary measure of the center of a probability distribution, we should also want some summary measure of the spread of a probability distribution to say something about what kinds of deviations from this expected value are likely to occur. The most useful summary measure of spread is the standard deviation.

The *standard deviation* of a random variable **X** may be denoted by Stdev(**X**) or $\sigma_{\mathbf{X}}$, and it is defined by the formula

$$\text{Stdev}(\mathbf{X}) = \sigma_{\mathbf{X}} = (E((\mathbf{X} - \mu_{\mathbf{X}})\text{^}2))\text{^}0.5 = (\sum_x P(\mathbf{X} = x)*(x - \mu_{\mathbf{X}})\text{^}2)\text{^}0.5$$

(Here, $\mu_{\mathbf{X}} = E(\mathbf{X})$, and the summation ($\sum_x$) includes all numbers $x$ that are possible values of the unknown quantity **X**. The symbol "^" is used here and in Excel to indicate exponentiation, so 3^2 = $3^2$ = 9, for example.) In words, the standard deviation of **X** may be defined as the square root of the expected squared deviation of **X** from its mean. If we dropped the square root (^0.5) from this definition, then we get the definition of the *variance* of **X**, which is the expected value of the squared deviation of **X** from its mean

$$\textit{Var}(\mathbf{X}) = E((\mathbf{X} - \mu_{\mathbf{X}})\text{^}2) = \sum_x P(\mathbf{X} = x)*(x - \mu_{\mathbf{X}})\text{^}2 = (\text{Stdev}(\mathbf{X}))\text{^}2$$

In the Superior Semiconductor example, the unknown number of competitors **K** has standard deviation

$$(0.10*(1 - 3)\text{^}2 + 0.25*(2 - 3)\text{^}2 + 0.30*(3 - 3)\text{^}2 + 0.25*(4 - 3)\text{^}2 + 0.10*(5 - 3)\text{^}2)\text{^}0.5 = 1.14$$

To see how these calculations may be done in a spreadsheet, look at cells D5:D9 and D13 in Figure 2.5. Recall that the possible values of **K** are listed in cells B5:B9, the corresponding probabilities are listed in C5:C9, and the

Myerson_C02.qxd 09-09-2004 10:31 Page 61

expected value or mean of **K** is listed in cell B13 of Figure 2.5. So, entering the formula

=(B5–$B$13)^2

into cell D5, and copying D5 to D5:D9, we get the possible squared deviations of **K** from its mean listed in cells D5:D9. Then the formula SUMPRODUCT(D5:D9,C5:C9) returns the expected squared deviation of **K** from its mean. This expected squared deviation of a random variable from its mean is called the *variance* of the random variable. The standard deviation is the square root of the variance, and it is returned in cell D13 by the formula

=SUMPRODUCT(D5:D9,C5:C9)^0.5

To simplify these calculations, we may also use the Simtools function STDEVPR. This function is designed to compute standard deviations from discrete probability distributions. (Be careful not to be confuse Simtools's STDEVPR function with Excel's functions STDEV and STDEVP, which are used to compute sample standard deviations from sample data. The PR in STDEVPR is short for "PRobability distribution.") STDEVPR takes two parameters, which should be ranges that have the same size (the same number of rows and the same number of columns). The first parameter, called "values" in the Insert-Function dialog box, should be a range that lists the possible values of some discrete random variable. The second parameter, called "probabilities" in the Insert-Function dialog box, should be a range that lists the corresponding probabilities of these values. Then the formula STDEVPR(values, probabilities) returns the standard deviation of the discrete random variable. For example, the formula

=STDEVPR(B5:B9,$C$5:$C$9)

in cell C13 of Figure 2.5 returns the same standard deviation (1.14) that we computed in cell D13.

If you have never seen a standard deviation before, then be aware that learning to interpret standard deviations takes time. There is no simple rule about how to interpret them. The most that can be said now is that the unknown quantity generally has some substantial probability of being more than 1 standard deviation above or below its expected value, but it is generally very unlikely to be more than 3 standard deviations above or below its expected value. Thus, for example, if you only told me that $E(\mathbf{K})$ = 3.0 and $Stdev(\mathbf{K})$ = 1.14 for some unknown quantity **K**, then I would infer that the probability of **K** being above 3.0 + 1.14 = 4.14 or below 3.0 – 1.14 = 1.86 was substantial (this probability is actually 0.20 in our example), but the probability of **K** being above 3.0 + 3∗1.14 = 6.42 or below 3.0 – 3∗1.14 = –0.42 was very small (this probability is actually 0 in our example). If these vague inferences were not enough for me, then I would ask you to show me the (inverse) cumulative chart for this unknown quantity.

Myerson_C02.qxd 09-09-2004 10:31 Page 62

Two formulas that you should know about are how multiplication by a nonrandom constant $c$ and addition of a nonrandom constant $d$ would affect expected values and standard deviations. If **X** is a random variable but $c$ and $d$ are numbers that are not random, then

$$E(c*\mathbf{X} + d) = c*E(\mathbf{X}) + d$$
$$Stdev(c*\mathbf{X} + d) = |c| * Stdev(\mathbf{X})$$

(In the second formula, $|c|$ denotes the absolute value of $c$, that is, $|c| = c$ if $c \geq 0$, while $|c| = -c$ if $c < 0$.) For example, let the random variable **R** denote Superior Semiconductor's total revenue in dollars (not millions) from the new product, ignoring the fixed costs. This revenue number **R** in dollars and the profit **Y** in millions of dollars are obviously related by the formula $\mathbf{R} = 1{,}000{,}000*\mathbf{Y} + 26{,}000{,}000$. So knowing that $E(\mathbf{Y}) = 1.5$ and $Stdev(\mathbf{Y}) = 9.317$, we can compute

$$E(\mathbf{R}) = 1{,}000{,}000*E(\mathbf{Y}) + 26{,}000{,}000 = 27{,}500{,}000$$
$$Stdev(\mathbf{R}) = 1{,}000{,}000*Stdev(\mathbf{Y}) = 9{,}317{,}000$$

## 2.5 Estimates from Sample Data

In complex decision problems we will want to estimate the expected values, standard deviations, and cumulative distributions of unknown quantities that we can study only by simulation. In this section we show how such estimates can be generated. It will be helpful to begin studying these techniques in the context of the simple Superior Semiconductor example, where we know how to compute explicitly the numbers that we are trying to estimate from simulation, because this simplicity will enable you to get a hands-on feel for the accuracy of these simulation methods.

Figure 2.6 shows a table of simulations of the unknown number of competitors **K** from the Superior Semiconductor case. The possible values of **K** are listed in cells A3:A7, and their corresponding probabilities are listed in cells B3:B7. So, of course, we can compute the expected value and standard deviation of **K** directly from this probability distribution, as shown in cells D7 and E7. But let us pretend that we did not know how to do these computations and instead try to work with simulation data.

A random variable that simulates the unknown quantity **K** has been created in cell B14 of Figure 2.6 by the formula

```
=DISCRINV(RAND( ),A3:A7,B3:B7)
```

Then data from 401 independent simulations of this random variable have been entered below into cells B15:B415 of this spreadsheet by selecting the range A14:B415 and using the command sequence Tools:SimTools: SimulationTable. (Simtools's SimulationTable command also entered the

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | Probability distribution of K (#competitors entering market) | | | | | | | |
| 2 | k | P(K=k) | | | | | | |
| 3 | 1 | 0.1 | | | | | | |
| 4 | 2 | 0.25 | | | | | | |
| 5 | 3 | 0.3 | | | | | | |
| 6 | 4 | 0.25 | | E(K) | Stdev(K) | | | |
| 7 | 5 | 0.1 | | 3 | 1.140175 | | | |
| 8 | | | | | | | | |
| 9 | | 2.975062 | E(K) estimated from SimTable | | | | | |
| 10 | | 1.095161 | Stdev(K) estimated from SimTable | | | | | |
| 11 | | 401 | Sample size n | | | | | |
| 12 | | | | | | | | |
| 13 | | Sim'd K | | 1.095161 | (SumSqDevs/(n-1))^0.5 | | | |
| 14 | SimTable | 2 | | Squared deviations from sample mean | | | | |
| 15 | 0 | 4 | | 1.050497 | | | | |
| 16 | 0.0025 | 3 | | 0.000622 | FORMULAS FROM RANGE A1:E15 | | | |
| 17 | 0.005 | 2 | | 0.950747 | D7. =SUMPRODUCT(A3:A7,B3:B7) | | | |
| 18 | 0.0075 | 2 | | 0.950747 | E7. =STDEVPR(A3:A7,B3:B7) | | | |
| 19 | 0.01 | 1 | | 3.900871 | B9. =AVERAGE(B15:B415) | | | |
| 20 | 0.0125 | 5 | | 4.100373 | B10. =STDEV(B15:B415) | | | |
| 21 | 0.015 | 4 | | 1.050497 | B11. =COUNT(B15:B415) | | | |
| 22 | 0.0175 | 2 | | 0.950747 | D13. =(SUM(D15:D415)/(B11-1))^0.5 | | | |
| 23 | 0.02 | 1 | | 3.900871 | B14. =DISCRINV(RAND(),A3:A7,B3:B7) | | | |
| 24 | 0.0225 | 1 | | 3.900871 | D15. =(B15-$B$9)^2 | | | |
| 25 | 0.025 | 4 | | 1.050497 | D15 copied to D15:D415 | | | |

FIGURE **2.6**
**Estimating an expected value and standard deviation from simulation data**

label SimTable into cell A14 and filled the range A15:A415 with a percentile index consisting of 401 equally spaced numbers from 0 to 1.)

To remind us how many data points we have in our simulation data (401), the formula =COUNT(B15:B415) has been entered into cell B11. Notice that the data range, as defined here, does not include the original random variable in cell B14. If we had included it, the statistics that we compute would change slightly every time the spreadsheet was recalculated. In our mathematical formulas, the number of independent data points that we are using to compute our statistics will be commonly denoted by the letter *n*.

The best estimate of $E(\mathbf{K})$ that can be computed from a sample of independent simulated values of **K** is the average or *sample mean* of this simulation data. So to estimate $E(\mathbf{K})$ in Figure 2.6, the formula

=AVERAGE(B15:B415)

has been entered into cell B9. Excel's AVERAGE function, of course, just sums the *n* numerical values in our data range and divides this sum by *n*. You can see that the value returned in this case (2.975) is reasonably close to the actual expected value (3).

The *law of large numbers* is a mathematical theorem which asserts that when we have a very large number of values drawn independently from a fixed probability distribution, the average of these values is very likely to be very close to the expected value of the probability distribution. (This is just an informal description of this theorem. Its formal mathematical statement gives precise meanings to the phrases "very large," "very likely," and "very close.") To see why this law is true, consider any discrete probability

Myerson_C02.qxd 09-09-2004 10:31 Page 64

distribution like the one for **K** here. If we let $m_k$ denote the number of times that the value $k$ occurs in our data range and let $n$ denote the size of the whole data range, then the average of the data range is

$$(\sum_k k*m_k)/n = \sum_k k*(m_k/n)$$

where the summation ($\sum_k$) includes all numbers $k$ that are possible values of the random variable. But when we generate hundreds of independent samples from the probability distribution of **K**, we should anticipate that the relative frequency $m_k/n$ of each possible value $k$ should be close to its probability. That is, $m_k/n$ should be quite close to $P(\mathbf{K} = k)$. Thus, the average of the data range should be close to the expected value

$$E(\mathbf{K}) = \sum_k k*P(\mathbf{K} = k)$$

For example, in the spreadsheet shown in Figure 2.6, the value 1 actually occurs 41 times among the 401 data values in the range B15:B415 (most of which are not shown in the figure). So $m_1/n = 41/401 = 0.102$, which is not far from $0.10 = P(\mathbf{K} = 1)$. The other frequencies in this data set happen to be such that

$$\begin{aligned}(\sum_k k*m_k)/n &= (1*41 + 2*92 + 3*134 + 4*104 + 5*30)/401 \\ &= \sum_k k*m_k/n = 1*41/401 + 2*92/401 + 3*134/401 \\ &\quad + 4*104/401 + 5*30/401 \\ &= 1*0.102 + 2*0.229 + 3*0.334 + 4*0.259 + 5*0.075 = 2.975\end{aligned}$$

which is the value shown in cell B9 of Figure 2.6. Matching the actual relative frequencies with the theoretical probabilities, you can see why this average is close to the true expected value

$$E(\mathbf{K}) = \sum_k k*P(\mathbf{K} = k) = 1*0.10 + 2*0.25 + 3*0.30 + 4*0.25 + 5*0.10 = 3$$

To estimate the standard deviation from sample data, we use the Excel function STDEV. For example, cell B10 of Figure 2.6 contains the formula

```
=STDEV(B15:B415)
```

which returns the value 1.095. You can see that this estimate is not far from the true standard deviation shown in cell E7 (1.140). Cells D13:D415 in Figure 2.6 have been set up to show you how Excel's STDEV function really works. Recall that the standard deviation of a random variable is the square root of its variance, and its variance is the expected squared deviation of this random variable from its expected value. So to estimate the standard deviation of a random variable from sample data, Excel's STDEV function must begin by estimating the expected value, which it does by computing the sample average, just as we have done in cell B9. Next, the STDEV function computes the squared deviation of each value in the data range from this sample average. These squared deviations have been computed in Figure 2.6 by entering the formula

```
=(B15–$B$9)^2
```

Myerson_C02.qxd 09-09-2004 10:31 Page 65

into cell D15, and then copying D15 to cells D15:D415. Next, you might think that we should estimate the variance (which is the expected squared deviation) by computing the average of these squared deviations. But statisticians have recommended that at this step we should instead compute a modified average in which the sum of the squared deviations is divided by $n - 1$ instead of $n$. (Computing deviations from the sample average instead of the true expected value tends to slightly reduce the average of these squared deviations, which could cause a downward bias in our estimates of the variance. To correct for this downward bias, statisticians have recommended dividing by $n - 1$ in the variance estimator.) In Figure 2.6, this modified average of the squared deviations would be returned by the formula

SUM(D15:D415)/(B11 – 1)

(Recall that the sample size $n$ is in cell B11.) Finally, to reverse the squaring of the deviations, STDEV takes the square root of this sample variance estimate. Taking the square root is the same as raising to the 0.5 power, so the formula

=(SUM(D15:D415)/(B11–1))^0.5

has been entered into cell D13 in Figure 2.6. You can verify that the result of this formula is exactly the value returned by the STDEV function (1.095). This estimated standard deviation computed by STDEV from sample data is called the *sample standard deviation*.

The law of large numbers can be extended to these estimated sample standard deviations. That is, when we have a very large number of values drawn independently from a fixed probability distribution, the sample standard deviations of these values is very likely to be very close to the actual standard deviation of the probability distribution.

Notice that Microsoft has provided STDEV as a built-in Excel function, because practical statistical work often requires that standard deviations be estimated from data. But the STDEVPR function is provided only by the Simtools.xla add-in because, in practice, standard deviations are only rarely computed from the probability distributions that define them.

To estimate the inverse cumulative distribution of a random variable from data in a simulation table we only need to sort the simulation data and plot the sorted values on the vertical ($Y$) axis of an $XY$ scatter chart, with the corresponding percentile index values from the left column of the simulation table plotted on the horizontal ($X$) axis of the chart. Figure 2.7 shows such a chart, made from the simulation table in Figure 2.6. To make Figure 2.7, first select the simulated data range B15:B415 and then enter the command sequence Data:Sort:Ascending. (When Excel's "sort warning" dialog box interrupts the task, choose the "continue" option.) Next, select the range A15:B415, which includes the percentile index numbers that go from 0 to 1 as well as the now-sorted data range, and then enter the command sequence Insert:Chart:XY-scatter (selecting the option to show lines

Myerson_C02.qxd 09-09-2004 10:31 Page 66

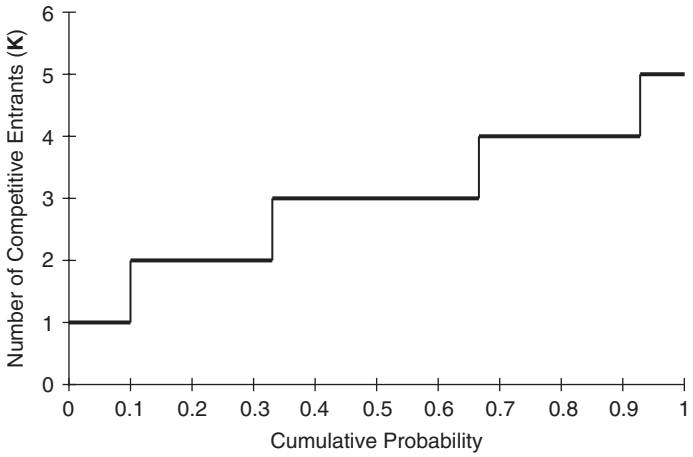


FIGURE **2.7**
**Estimate of an inverse cumulative distribution from simulation data**

and point markers) to create the chart as shown in Figure 2.7. Notice that Figure 2.7 is indeed a good approximation to the actual inverse cumulative probability chart shown previously as Figure 2.2.

## 2.6 Accuracy of Sample Estimates

When we estimate an expected value by computing the average of sample data, we need to know something about how accurate this estimate is likely to be. Of course, the average of several random variables is itself a random variable that has its own probability distribution. Figure 2.8 shows a spreadsheet designed to help you learn about how sample averages behave as random variables. Like other figures in this chapter, Figure 2.8 begins with the probability distribution for the unknown quantity **K** from the Superior Semiconductor case, with the possible values listed in cells A3:A7 and their corresponding probabilities listed in cells B3:B7. Then, to create 30 independent random variables with this probability distribution, the formula

=DISCRINV(RAND( ), $A$3:$A$7, $B$3,$B$7)

has been entered into every cell in the range A15:C24. The average or sample mean of these 30 random variables is entered into cell E15 by the formula

=AVERAGE(A15:C24)

You should make a spreadsheet like Figure 2.8 and recalculate it many times, watching the sample mean in cell E15. If you watch any individual cell in the sample range A15:C24, you will see it jump around to all the

Myerson_C02.qxd 09-09-2004 10:31 Page 67

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | Probability distribution of K (#competitors entering market) | | | | | | | |
| 2 | k | P(K=k) | | | | | | |
| 3 | 1 | 0.1 | | | | | | |
| 4 | 2 | 0.25 | | | | | | |
| 5 | 3 | 0.3 | | | | | | |
| 6 | 4 | 0.25 | | | | | | |
| 7 | 5 | 0.1 | | Calculated from probability distribution: | | | | |
| 8 | | 1 | | | 3 | E(K) | | |
| 9 | | | | | 1.140175 | Stdev(K) | | |
| 10 | | 30 | Sample size n | | 0.208167 | Stdev of sample mean | | |
| 11 | | | | | | | | |
| 12 | Normal random variable with same E & Stdev as sample mean: | | | | | | | 2.726072 |
| 13 | | | | | | | | |
| 14 | Simulated sample, size 30 | | | Calculated from the sample: | | | | |
| 15 | 3 | 4 | 3 | | 3.033333 | Sample mean or average | | |
| 16 | 3 | 4 | 4 | | 1.188547 | Sample standard deviation | | |
| 17 | 3 | 2 | 1 | | 0.216998 | Est'd stdev of sample mean | | |
| 18 | 1 | 5 | 5 | | | | | |
| 19 | 3 | 2 | 3 | 95% confidence interval for E(K) | | | | |
| 20 | 3 | 2 | 5 | | 2.608017 | 3.458649 | | |
| 21 | 5 | 4 | 3 | E(K) actually in the interval? | | | | |
| 22 | 3 | 1 | 1 | | TRUE | | | |
| 23 | 3 | 4 | 2 | | | | | |
| 24 | 3 | 3 | 3 | | | | | |
| 25 | | | | | | | | |
| 26 | FORMULAS FROM RANGE A1:H24 | | | | | | | |
| 27 | E8. =SUMPRODUCT(A3:A7,B3:B7) | | | | E15. =AVERAGE(A15:C24) | | | |
| 28 | E9. =STDEVPR(A3:A7,B3:B7) | | | | E16. =STDEV(A15:C24) | | | |
| 29 | B8. =SUM(B3:B7) | | | | E17. =E16/(B10^0.5) | | | |
| 30 | B10. =COUNT(A15:C24) | | | | E20. =E15-1.96*E17 | | | |
| 31 | E10. =E9/(B10^0.5) | | | | F20. =E15+1.96*E17 | | | |
| 32 | H12. =NORMINV(RAND(),E8,E10) | | | | E22. =AND(E20<E8,E8<F20) | | | |
| 33 | A15. =DISCRINV(RAND(),$A$3:$A$7,$B$3:$B$7) | | | | | | | |
| 34 | A15 copied to A15:C24 | | | | | | | |

FIGURE **2.8**
**A spreadsheet for studying the properties of a sample mean**

integer values from 1 to 5. But when you watch cell E15, the average of all 30 cells in this sample, you will see it vary much less widely around 3, almost never going below 2.2 or above 3.8. A remarkable mathematical fact called the Central Limit Theorem tells us much more about the way this average varies.

Before stating the Central Limit Theorem, I must introduce the idea of a Normal probability distribution. In probability theory, the phrase "Normal probability distribution" is used in a technical sense, emphasized by capitalizing it, and refers to a particular collection of mathematical probability distributions that have some important properties. For any two numbers $\mu$ and $\sigma$ such that $\sigma > 0$, there is precisely defined Normal probability distribution that has mean (or expected value) $\mu$ and standard deviation $\sigma$. We will discuss such Normal distributions at length in Chapter 4, but for now it is enough to introduce them by citing a few basic facts about them.

In an Excel spreadsheet, you can make a random variable that has a Normal probability distribution with mean $\mu$ and standard deviation $\sigma$ by the formula

NORMINV(RAND( ), $\mu$, $\sigma$)

Myerson_C02.qxd 09-09-2004 10:31 Page 68

Once you know how to make a spreadsheet cell that simulates a given probability distribution, you can learn anything that anybody might want to know about this distribution by simulating it many times in a spreadsheet. So if you want to know what a "Normal distribution with mean 100 and standard deviation 20" is like, you should simply copy the formula

=NORMINV(RAND( ),100,20)

into a large range of cells and watch how the values of these cells jump around whenever you press the recalc key [F9].

Here are a few other useful formulas about Normal distributions. If a random variable **X** has a Normal probability distribution with mean μ and standard deviation σ (where μ and σ are given numbers such that σ > 0), then

$$P(\mathbf{X} < \mu) = 0.5 = P(\mathbf{X} > \mu),$$
$$P(\mu - \sigma < \mathbf{X} < \mu + \sigma) = 0.683$$
$$P(\mu - 1.96*\sigma < \mathbf{X} < \mu + 1.96*\sigma) = 0.95$$
$$P(\mu - 3*\sigma < \mathbf{X} < \mu + 3*\sigma) = 0.997$$

That is, a Normal random variable is equally likely to be above or below its mean, it has probability 0.683 of being less than 1 standard deviation away from its mean, and it has probability 0.997 (almost sure) of being less than 3 standard deviations of its mean. For constructing 95% confidence intervals, we will use the fact that a Normal random variable has probability 0.95 of being within 1.96 standard deviations from its mean.

Now we are ready for the remarkable *Central Limit Theorem*, which tells us Normal distributions can be used to predict the behavior of sample averages:

> Consider the average of $n$ random variables that are drawn independently from a probability distribution with expected value μ and standard deviation σ. This average, as a random variable, has expected value μ, has standard deviation σ/($n$^0.5), and has a probability distribution that is approximately Normal.

For example, cell E15 of Figure 2.8 contains the average of $n$ = 30 independent random variables drawn from a probability distribution that has expected value μ = 3 and standard deviation σ = 1.14. So the Central Limit Theorem tells us that this sample mean should behave like a random variable that has a Normal distribution where μ = 3 is the mean and σ/($n$^0.5) = 1.14/(30^0.5) = 0.208 is the standard deviation. Such a Normal random variable is entered into cell H12 of Figure 2.8. If you watch cell H12 and cell E15 through many recalculations, the only difference in their pattern of behavior that you should observe is that the average in cell E15 is always a multiple of 1/30. To show more precisely that the sample average in cell E15 has a probability distribution very close to that of the Normal random variable in cell H12, you could make a simulation table containing several

Myerson_C02.qxd 09-09-2004 10:31 Page 69

hundred independently recalculated values of each of these random variables. Then, by separately sorting each column in this simulation table, you could make a chart that estimates the inverse cumulative distribution of each random variable. These two curves should be very close.

This Central Limit Theorem is the reason why, of all the formulas that people could devise for measuring the center and the spread of probability distributions, the expected value and the standard deviation have been the most useful for statistics. Other probability distributions that have the same expected value 3 and standard deviation 1.14 could be quite different in other respects, but the Central Limit Theorem tells us that an average of 30 independent samples from any such distribution would behave almost the same. (For example, try the probability distribution in which the possible values are 2, 3, and 7, with respective probabilities $P(2) = 0.260$, $P(3) = 0.675$, and $P(7) = 0.065$.)

Now suppose that we did not know the expected value of **K**, but we did know that its standard deviation was $\sigma = 1.14$, and we knew how to simulate **K**. Then we could look at any average of $n$ independently simulated values and we could assign 95% probability to the event that our sample average does not differ from the true expected value by more than $1.96*\sigma/(n\text{^}0.5)$. That is, if we let $\mathbf{Y}_n$ denote the average of our $n$ simulated values, then the interval from $\mathbf{Y}_n - 1.96*\sigma/(n\text{^}0.5)$ to $\mathbf{Y}_n + 1.96*\sigma/(n\text{^}0.5)$ would include the true $E(\mathbf{K})$ with probability 0.95. This interval is called a 95% confidence interval. With $n = 30$ and $\sigma = 1.14$, the *radius* $r$ (that is, the distance from the center to either end) of this 95% confidence interval would be

$$r = 1.96*\sigma/(n\text{^}0.5) = 1.96*1.14/(30\text{^}0.5) = 0.408$$

If we wanted the radius of our 95% confidence interval around the sample mean to be less than some number $r$, then we would need to increase the size of our sample so that

$$1.96*\sigma/(n\text{^}0.5) < r, \qquad \text{and so } n > (1.96*\sigma/r)\text{^}2$$

For example, to make the radius of our 95% confidence interval smaller than 0.05, the sample size $n$ must be

$$n > (1.96*\sigma/r)\text{^}2 = (1.96*1.14/0.05)\text{^}2 = 1997$$

Now consider the case where we know how to simulate an unknown quantity but we do not know how to calculate its expected value or its standard deviation. In this case, where our confidence interval formula calls for the unknown probabilistic standard deviation $\sigma$, we must replace it by the sample standard deviation that we compute from our simulation data. If the average of $n$ independent simulations is $X$ and the sample standard deviation is $S$, then our estimated 95% confidence interval for the true expected value is from $X - 1.96*S/(n\text{^}0.5)$ to $X + 1.96*S/(n\text{^}0.5)$, where the quantity $S/(n\text{^}0.5)$ is our estimated standard deviation of the sample average. In Figure 2.8, for example, the sample standard deviation $S$ is computed in cell E16 by the formula =STDEV(A15:C24), the sample size $n$ is computed

Myerson_C02.qxd 09-09-2004 10:31 Page 70

in cell B10 by the formula =COUNT(A15:C24), the quantity S/*n*^0.5) is computed in cell E17 by the formula =E16/(B10^0.5), and then a 95% confidence interval for $E(\mathbf{K})$ is calculated in cells E20 and F20 by the formulas

=E15–1.96*E17 and =E15+1.96*E17

(Recall that E15 is the sample average.)

To say that the interval from E20 to F20 is a 95% confidence interval for the true expected value is to say that the true expected value of 3 (in cell E8) should be between these two numbers 95% of the time when the spreadsheet in Figure 2.8 is recalculated many times independently. You can verify this by watching cell E22 while recalculating. Cell E22 contains the formula

=AND(E20<E8,E8<F20)

and so it should read TRUE about 95% of the time (about 19 times out of 20). (Recall that the Excel function AND(statement1,statement2) returns the value TRUE if statement1 and statement2 are both TRUE, and returns FALSE otherwise.)

So based only on the 30 independent simulations shown in Figure 2.8, the formulas in cells E20 and F20 give us a 95% confidence interval from 2.608 to 3.458, that is 3.033 ± 0.425, for $E(\mathbf{K})$. If the radius 0.425 seems too large, then we could get a narrower 95% confidence interval by using a larger sample. If we only knew this estimated sample standard deviation of 1.188, we could estimate that the radius of our 95% confidence interval could be reduced below 0.05 if the sample size *n* were increased so that 1.96*1.188/(*n*^0.5) were less than 0.05, which happens when *n* is greater than (1.96*1.188/0.05)^2 = 2170.

It is important to understand the difference between cells E16 and E17 in Figure 2.8. The value of cell E16 is the sample standard deviation STDEV(A15:C24). So E16 is our statistical estimate of the standard deviation of any one random cell in the range A15:C24. The value of cell E17 is the sample standard deviation in E16 divided by the square root of the sample size. So E16 is our estimate of the standard deviation of the random cell E15, which is average of all the random cells in the range A15:C24. When we recalculate the 30 simulated values in cells A15:C24, the sample average tends to vary less than any one cell in the sample range, because when one cell is relatively high there is usually some other relatively low cell to cancel it out. That is why the standard deviation of the sample average is smaller than the standard deviation of any one cell in the sample, by a factor of 1/*n*^0.5.

**NOTE** (You might wonder whether we should broaden our 95% confidence intervals when we use an estimated sample standard deviation instead of the true standard deviation. The answer is that we should, but the adjustments are relatively minor unless the sample size is small. The corrected formula is based on something that statisticians call a T distribution. When the sample size *n* is small, we should replace the

Myerson_C02.qxd 09-09-2004 10:31 Page 71

constant 1.96 in our 95% confidence formulas by the value of the Excel formula TINV(0.05,$n$ – 1). When $n$ is 30, this value is 2.045. As the sample size $n$ increases, the value of TINV(0.05,$n$ – 1) rapidly approaches 1.96, and so we will not worry about this T adjustment in this book.)

Another application of the Central Limit Theorem can also be used to tell us something about the accuracy of our statistical estimates of points on the inverse cumulative probability curve. Suppose that **X** is a random variable with a probability distribution that we know how to simulate, but we do not know how to directly compute its cumulative probability curve. For any number $y$, let $\mathbf{Q}_n(y)$ denote the percentage of our sample that is less than $y$, when we get a sample of $n$ independent values drawn from the probability distribution of **X**. Then we should use $\mathbf{Q}_n(y)$ as our estimate of $P(\mathbf{X} < y)$. But how good is this estimate? When $n$ is large, $\mathbf{Q}_n(y)$ is a random variable with an approximately Normal distribution, its expected value is $P(\mathbf{X} < y)$, and its standard deviation is $(P(\mathbf{X} < y)*(1 - P(\mathbf{X} < y))/n)$^0.5. But $(p*(1 - p))$^0.5 ≤ 0.5 for any probability $p$, and so the standard deviation of $\mathbf{Q}_n(y)$ is always less than 0.5/($n$^0.5), and multiplying this standard deviation by 1.96 yields a number slightly less than 1/$n$^0.5. Thus, around any point $(\mathbf{Q}_n(y),y)$ in an estimated inverse cumulative probability curve like Figure 2.7, we could put a horizontal confidence interval over the cumulative probabilities from $\mathbf{Q}_n(y) - 1/n$^0.5 to $\mathbf{Q}_n(y) + 1/n$^0.5, and this interval would have a probability greater than 95% of including the true cumulative probability at $y$. When $n$ is 400, for example, the radius of this cumulative probability interval around $\mathbf{Q}_n(y)$ is 1/$n$^0.5 = 1/20 = 0.05. If we wanted to reduce the radius of this 95% confidence interval below 0.02, then we would increase the sample size to 1/0.02^2 = 2500.

## 2.7 Decision Criteria

We have used the Superior Semiconductor case as an example to introduce basic ideas of probability and statistics. But after all these ideas have been introduced, we are still left with a decision problem. On the basis of our analysis of the uncertainty in this situation, should we recommend that Superior Semiconductor introduce the new T-regulator product or not? To answer this question, we need some fundamental assumption about what determines an optimal decision under uncertainty. The assumption that we will usually apply in this course is the *criterion of expected value maximization* (or the expected value criterion).

Any quantitative decision analysis must involve some numerical measure of payoff such that increasing the decision-maker's payoff is considered an improvement. For most economic decision problems, net monetary returns or profit (or the present discounted value of future profit) may be

Myerson_C02.qxd 09-09-2004 10:31 Page 72

identified as the payoff that the decision maker wants to increase whenever possible. But in situations with uncertainty, we may not know whether a particular decision (like that of introducing the new T-regulator product) would increase or decrease the decision-maker's payoff. The criterion of expected value maximization asserts that, among the various alternatives that are available to a decision maker, the optimal decision is the one that yields the highest expected value of the decision-maker's payoff.

So if we take Superior Semiconductor's profit to be the measure of "payoff" in this decision problem, then the optimal decision can be identified simply by computing the expected value of Superior Semiconductor's profit from the proposed new product. As we have seen, this expected value is

$$E(\text{Profit}) = 0.10*24 + 0.25*7.33 + 0.30*(-1) + 0.25*(-6) + 0.10*(-9.33) = 1.5$$

The alternative of not introducing the new product generates an expected profit of 0, of course. So by the criterion of expected value maximization, Superior Semiconductor should introduce the new product, because $1.5 > 0$.

The expected value formula has many good properties to recommend it as a criterion for decision making under uncertainty. It takes account of all possible outcomes in a sensible way, and it is more sensitive to outcomes that are more likely. The argument for expected value maximization is particularly compelling in games that can be repeated. If we know that we will repeat a given type of decision problem many times, with new payoffs from each repetition being added to the payoffs from previous rounds but with the new outcome being determined independently each time, then a strategy of choosing the alternative that yields the highest expected value will almost surely maximize our long-term total payoff by the law of large numbers.

This expected value criterion may be interpreted to mean that all we should care about is the expected value of some appropriately measured payoff. But this interpretation can lead to trouble if the words "of payoff" are forgotten. If you thought that our expected value criterion meant that we should only care about the expected number of competitors, then you would act as though the number of competitors would be three, in which case profit would be –1, and your recommendation would be to not introduce the new product. The error here is to compute the expected value of the wrong random variable (not payoff) and then to try to compute an expected payoff from it according to the fallacy of averages (as discussed in Section 2.4).

Figure 2.9 shows a decision analysis of the Superior Semiconductor case. In cell E14 the expected value in profit is calculated directly from the probability distribution. Under the expected value criterion, this positive expected value in E14 tells us that we should recommend the new product. But to illustrate what we would do if we could not compute expected profit directly from the probability distribution, the spreadsheet also contains a table of 401 independent simulations of the unknown profit, and the average of these profits is exhibited in cell B21 as an estimator of the expected profit. Even if we could not see the true expected value in cell E14, our

Myerson_C02.qxd 09-09-2004 10:31 Page 73

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | SUPERIOR SEMICONDUCTOR (A) | (B) | | | 26 | FixedCost | | |
| 2 | | | | | 100 | MarketValue | | |
| 3 | Let K = (unknown number of competitors entering market) | | | | | | | |
| 4 | Probability distribution of K | | | | | | | |
| 5 | | k | P(K=k) | | Profit ($millions) if k=#compet'rs | | | |
| 6 | | 1 | 0.10 | | 24.00 | | | |
| 7 | | 2 | 0.25 | | 7.33 | Measuring risk tolerance | | |
| 8 | | 3 | 0.30 | | -1.00 | | High$ | 20.00 |
| 9 | | 4 | 0.25 | | -6.00 | | Low$ | -10.00 |
| 10 | | 5 | 0.10 | | -9.33 | Assessed value$ | | 2.00 |
| 11 | | | | | | RiskTolerance | | 36.4891 |
| 12 | Computations from probability distribution | | | | | | | |
| 13 | | E(K) | | | E(Profit) | CertaintyEquivalent | | |
| 14 | | 3 | | | 1.5 | | | 0.423906 |
| 15 | | | | | | | | |
| 16 | Simulation model | | | | | | | |
| 17 | | K | | | Profit | | | |
| 18 | | 5 | | | -9.33 | | | |
| 19 | | | | | | | | |
| 20 | Statistics from simulations: | | | | | | | |
| 21 | | 1.83458 | E(Profit) est'd from SimTable | | | | | |
| 22 | | 9.611946 | Stdev(Profit) est'd from SimTable | | | | | |
| 23 | | -9.33333 | 5% cumulative-probability level | | | | | |
| 24 | | 401 | Sample size | | | Estimated CE | | 0.696039 |
| 25 | | | | 95% conf.intl for E(Profit) | | | | |
| 26 | Sim'd profit | | | | 0.8937847 | 2.775376 | | |
| 27 | SimTable | -9.33 | | | | | | |
| 28 | 0 | -6 | FORMULAS | | | | | |
| 29 | 0.0025 | -9.33333 | E6. =$E$2/(1+B6)-$E$1 | | | | | |
| 30 | 0.005 | -6 | E6 copied to E6:E10 and E18 | | | | | |
| 31 | 0.0075 | -6 | B14. =SUMPRODUCT(B6:B10,$C$6:$C$10) | | | | | |
| 32 | 0.01 | -1 | E14. =SUMPRODUCT(E6:E10,$C$6:$C$10) | | | | | |
| 33 | 0.0125 | -6 | B18. =DISCRINV(RAND(),B6:B10,C6:C10) | | | | | |
| 34 | 0.015 | -6 | B27. =E18 | | | | | |
| 35 | 0.0175 | -1 | B21. =AVERAGE(B28:B428) | | | | | |
| 36 | 0.02 | -1 | B22. =STDEV(B28:B428) | | | | | |
| 37 | 0.0225 | -1 | B23. =PERCENTILE(B28:B428,0.05) | | | | | |
| 38 | 0.025 | 7.333333 | B24. =COUNT(B28:B428) | | | | | |
| 39 | 0.0275 | -1 | E26. =B21-1.96*B22/(B24^0.5) | | | | | |
| 40 | 0.03 | 24 | F26. =B21+1.96*B22/(B24^0.5) | | | | | |
| 41 | 0.0325 | -1 | H11. =RISKTOL(H8,H9,H10) | | | | | |
| 42 | 0.035 | -1 | H14. =CEPR(E6:E10,C6:C10,H11) | | | | | |
| 43 | 0.0375 | -9.33333 | H24. =CE(B28:B428,H11) | | | | | |

FIGURE **2.9**
**Analysis of Superior Semiconductor case**

simulation data is strong enough to support reasonable confidence that the expected value is greater than 0 because we find a positive lower bound (0.8937) in the 95% confidence interval for expected profit that is computed in cells E26 and F26.

But now, having advocated the expected value criterion, it must be admitted that it is often not fully satisfactory as a basis for decision making. In practice, people often prefer decision alternatives that yield lower expected profits when the alternatives that yield higher expected profits are also more risky. People who feel this way are *risk averse*. Because most people express attitudes of risk aversion at least some of the time, a serious decision analysis should go beyond simply reporting expected payoff values and should also report some measures of the risks associated with the alternatives that are being considered.

Myerson_C02.qxd 09-09-2004 10:31 Page 74

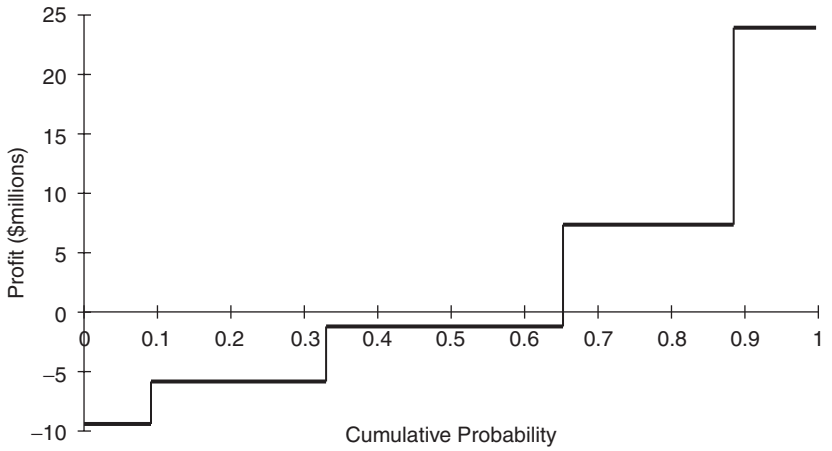


FIGURE **2.10**
**Cumulative risk profile from simulation data**

As we have seen, the standard deviation is often used as a measure of the spread of likely outcomes of an unknown quantity, and so the standard deviation of profit may be used as a measure of risk. Thus, cell B22 in Figure 2.9 estimates the standard deviation of profit from the proposed new product in this case. The large size of this sample standard deviation ($9.61 million, much larger than the expected value) is a strong indication that this new product should be seen as very risky.

Another measure of risk that has gained popularity in recent years is called *value at risk*. The value at risk is defined to be the level of net profit that has some small, prespecified cumulative probability, often taken to be 0.05, so that the probability of profit being below this level is not more than 1/20. So cell B23 in Figure 2.9 estimates the profit level that has 5% cumulative probability from our simulation data in B28:B428, using the formula

=PERCENTILE(B28:B428,0.05).

The *cumulative risk profile* for a decision may be defined as the inverse cumulative probability distribution of the payoff that would result from this decision. Figure 2.10 shows the cumulative risk profile for the decision to introduce the new product in this case. This cumulative risk profile was made from the simulation table in Figure 2.9. (First the simulated profit data in cells B28:B428 were sorted by Excel's Data:Sort command, and then these sorted profit values were plotted on the vertical axis of an *XY* chart, with the percentile index in cells A28:A428 plotted on the horizontal axis.) Notice that this cumulative risk profile contains all the information about the value at risk for any probability level. By definition, the value at risk for the cumulative probability level 0.05 is just the height of the cumulative risk profile above 0.05 on the horizontal cumulative probability axis. So the cumulative risk profile may give the most complete overall picture of the risks associated with a decision.

Myerson_C02.qxd 09-09-2004 10:31 Page 75

More generally, the limitations of the expected value maximization as a criterion for decision making have been addressed by two important theories of economic decision making: *utility theory* and *arbitrage pricing theory*. Utility theory is about decision making by risk-averse individuals. Arbitrage pricing theory is about decision making for publicly held corporations, which may be owned by people who have different attitudes toward risks and different beliefs about the probabilities of various events. Utility theory will be discussed in Chapter 3, and arbitrage pricing theory will be discussed in Sections 8.3 and 8.4. But each of these theories turns out to be mathematically equivalent to a simple extension of the expected value criterion. In utility theory, the expected value criterion is extended by introducing a new way of measuring payoffs, called a *utility function*, that takes account of the decision-maker's personal willingness to take risks. In arbitrage pricing theory, the expected value criterion is extended by introducing a new way of measuring probabilities, called *market-adjusted probabilities*, that takes account of asset pricing in financial markets. Thus, the techniques that we have developed in this chapter for estimating expected values will also be applicable to more sophisticated theories of individual and corporate decision making.

**NOTE** (For a preview of the results of utility theory, you can look in column H of Figure 2.9. In Part B of the Superior Semiconductor case, the decision maker will remark that a simple gamble that could generate either a profit of $20 million or a loss of $10 million for Superior Semiconductor, each with probability 1/2, may be just as good for the company as $2 million for sure. Based on this assessment, cell H11 in Figure 2.9 applies a Simtools function called RISKTOL to compute a measure of the company's "risk tolerance." Then, with this measure of the company's risk tolerance, cell H14 applies a Simtools function called CEPR to compute a risk-adjusted "certainty equivalent" value of the T-regulator project using the probability distribution of its profits. Cell H24 applies another Simtools function called CE to estimate this same certainty equivalent value using simulation data. The meanings of these mysterious functions and quantities will be discussed at length in Chapter 3.)

## 2.8 Multiple Random Variables

We have been considering a simple example with just one random variable because this simplicity allowed us to compare our first estimates from simulation data to the actual quantities being estimated. But now that we understand how simulation works, we can begin applying it to more interesting problems that have many random variables, where probabilities, expected

Myerson_C02.qxd 09-09-2004 10:31 Page 76

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | #Competitors | | | Dev't cost | | | Total market | | |
| 2 | value | proby | | value | proby | | value | proby | |
| 3 | 1 | 0.1 | | 20 | 0.2 | | 70 | 0.3 | |
| 4 | 2 | 0.25 | | 26 | 0.5 | | 100 | 0.4 | |
| 5 | 3 | 0.3 | | 30 | 0.2 | | 120 | 0.2 | |
| 6 | 4 | 0.25 | | 34 | 0.1 | | 150 | 0.1 | |
| 7 | 5 | 0.1 | | | 1 | | | 1 | |
| 8 | | 1 | | | | | | | |
| 9 | SIMULATION MODEL | | | | | | | | |
| 10 | #Competitors | | | Development cost | | | Total market$ | | |
| 11 | 3 | | | 20 | | | 100 | | |
| 12 | | Profit | | | | | | | |
| 13 | SimTabl | 5 | | Ests from SimTable | | | 95% Conf | Intl for EProfit | |
| 14 | 0 | -22.333 | | 1.11477 | E(Profit) | | -0.0069 | 2.23647 | |
| 15 | 0.002 | -20 | | 12.8097 | Stdev(Proft) | | | | |
| 16 | 0.004 | -20 | | -14.333 | 5%ile Profit | | | Sample size | |
| 17 | 0.006 | -20 | | 0.582 | P(Profit<0) | | | 501 | |
| 18 | 0.008 | -20 | | FORMULAS | | | | | |
| 19 | 0.01 | -20 | | B8. =SUM(B3:B7) | | | | | |
| 20 | 0.012 | -18.333 | | E7. =SUM(E3:E6) | | | E7 copied to H7 | | |
| 21 | 0.014 | -18.333 | | A11. =DISCRINV(RAND(),A3:A7,B3:B7) | | | | | |
| 22 | 0.016 | -18.333 | | D11. =DISCRINV(RAND(),D3:D6,E3:E6) | | | | | |
| 23 | 0.018 | -18.333 | | G11. =DISCRINV(RAND(),G3:G6,H3:H6) | | | | | |
| 24 | 0.02 | -18.333 | | B13. =G11/(1+A11)-D11 | | | | | |
| 25 | 0.022 | -18.333 | | D14. =AVERAGE(B14:B514) | | | | | |
| 26 | 0.024 | -16.5 | | D15. =STDEV(B14:B514) | | | | | |
| 27 | 0.026 | -16.5 | | D16. =PERCENTILE(B14:B514,0.05) | | | | | |
| 28 | 0.028 | -16.5 | | D17. =PERCENTRANK(B14:B514,0) | | | | | |
| 29 | 0.03 | -16.5 | | H17. =COUNT(B14:B514) | | | | | |
| 30 | 0.032 | -16 | | G14. =D14-1.96*D15/H17^0.5 | | | | | |
| 31 | 0.034 | -16 | | H14. =D14+1.96*D15/H17^0.5 | | | | | |
| 32 | 0.036 | -16 | | | | | | | |
| 33 | 0.038 | -16 | | | | | | | |
| 34 | 0.04 | -16 | | | | | | | |
| 35 | 0.042 | -16 | | | | | | | |
| 36 | 0.044 | -16 | | | | | | | |
| 37 | 0.046 | -16 | | | | | | | |
| 38 | 0.048 | -14.333 | | | | | | | |
| 39 | 0.05 | -14.333 | | | | | | | |
| 40 | 0.052 | -14.333 | | | | | | | |
| 41 | 0.054 | -14 | | | | | | | |
| 42 | 0.056 | -14 | | | | | | | |
| 43 | 0.058 | -14 | | | | | | | |
| 44 | 0.06 | -14 | | | | | | | |
| 45 | 0.062 | -14 | | Data in B14:B514 is sorted for chart. | | | | | |
| 46 | 0.064 | -14 | | Chart plots (A14:A514,B14:B514). | | | | | |

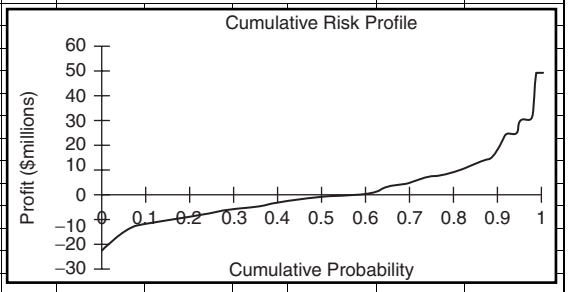


FIGURE **2.11**
**A simulation model with three random variables affecting profit**

values, and standard deviations may be very difficult to calculate from probability distributions, so that simulation analysis becomes our best technique.

For example, Figure 2.11 shows an analysis of a more complicated version of the Superior Semiconductor case where the development cost and total market value are also unknown quantities. Instead of assuming that development cost **D** is $26 million for sure, it is assumed now that the development cost could be $20, 26, 30, or 34 million with probabilities 0.2, 0.5, 0.2, and 0.1, respectively. Instead of assuming that the total market value **M** is $100 million for sure, it is assumed now that the development cost could be $70, 100, 120, or 150 million with probabilities 0.3, 0.4, 0.2, and 0.1, respectively. The number of competitors **K**, as before, could be

Myerson_C02.qxd 09-09-2004 10:31 Page 77

1, 2, 3, 4, or 5, with probabilities 0.1, 0.25, 0.3, 0.25, and 0.1, respectively. The profit depends on these quantities by the formula $\mathbf{Y} = \mathbf{M}/(1 + \mathbf{K}) - \mathbf{D}$ because we assume that the total market value will be divided equally among Superior Semiconductor and its **K** competitors.

Cells A11, D11, and G11 in Figure 2.11 contain random variables (made with DISCRINV) that simulate the number of competitors, the development cost, and the total market value for Superior Semiconductor's new T-regulator device. The spreadsheet uses an assumption that these unknown quantities are independent (that is, learning about any one of them would not influence our beliefs about the others) because the three random variables in cells A11, D11, and G11 are independently driven by their own separate RANDs. Then profit is calculated in cell B13 by the formula =G11/(1+A11)–D11.

Profit data from 501 simulations of this model are stored in B14:B514. The sample average is calculated in D14, and a 95% confidence interval for the true expected profit is calculated around this sample average in cells G14 and H14. Based on this analysis, it appears that the strategy of introducing the new T-regulator product can be recommended under the expected value criterion, but a larger simulation table may be needed to be more confident about the positivity of expected profit. The risks appear even greater in this model, as evidenced by a higher standard deviation (in cell D15) and a lower 0.05 cumulative probability profit level (in cell D16) compared to the analogous statistics in Figure 2.9.

A fuller picture of the probability distribution of profit in this example is offered by the cumulative risk profile, which plots the sorted simulation data in B14:B514 on the vertical axis against the simulation table's percentile index (in A14:A514) on the horizontal axis. The Excel functions PERCENTILE and PERCENTRANK are also used in Figure 2.11 to extract numerical information about this distribution. To estimate the cumulative probability of profit at \$0, the fraction of the simulated profits in B14:B514 that are less than \$0 is computed in cell D17 by the formula

```
=PERCENTRANK(B14:B514,0)
```

To estimate the profit value that has cumulative probability 0.05, cell D16 finds the value that is greater than 5% of the simulated profits using the formula

```
=PERCENTILE(B14:B514,0.05)
```

## 2.9 Summary

In this chapter we focused on a simple decision problem involving an unknown quantity that has only finitely many possible values. In this context we introduced some basic concepts for describing discrete probability

Myerson_C02.qxd 09-09-2004 10:31 Page 78

distributions: the expected value or mean of the distribution, the standard deviation and the variance, and cumulative probability charts. We saw how to make a random variable with any given probability distribution by using the inverse cumulative probability function with a RAND( ) as input. The Simtools function DISCRINV was introduced to facilitate such simulations.

We then introduced techniques for estimating expected values, standard deviations, and cumulative probabilities from simulation data, using the law of large numbers for assurance that these estimates are very likely to be quite accurate if the sample size is very large. For a more precise assessment of the accuracy of the sample average as an estimate for an unknown expected value, we introduced Normal distributions and the Central Limit Theorem. We learned that a sample average, as a random variable, has a standard deviation that is inversely proportional to the square root of the sample size. We then saw how to compute a 95% confidence interval for the expected value of a random variable, using simulation data.

Finally, criteria for optimal decision making were discussed, beginning with the basic concept of expected value maximization. The expected value of monetary income or some other suitably measured payoff quantity was recommended as the best single number to guide decision making under uncertainty. The standard deviation of payoff, the value at risk (for some prespecified cumulative probability level), and the entire cumulative risk profile were recommended as also worth reporting in a decision analysis, to better describe the levels of risk entailed by different decision alternatives.

Excel functions used in this chapter included AND, NORMINV, STDEV, and SUMPRODUCT. Simtools functions introduced in this chapter included DISCRINV and STDEVPR. We also used the Data:Sort and Insert:Chart:XY-Scatter commands to make inverse cumulative charts from simulation data.

## Exercises

**1.** Let **X** denote an unknown quantity that has three possible values: 2, 3, and 7, and suppose that their probabilities are $P(\mathbf{X} = 2) = 0.260$, $P(\mathbf{X} = 3) = 0.675$, $P(\mathbf{X} = 7) = 0.065$.
Let **Y** denote another unknown quantity that has three possible values: –1, 3, and 4, and suppose that their probabilities are $P(\mathbf{Y} = -1) = 0.065$, $P(\mathbf{Y} = 3) = 0.675$, $P(\mathbf{Y} = 4) = 0.260$.

**(a)** Compute $E(\mathbf{X})$, $Stdev(\mathbf{X})$, $E(\mathbf{Y})$, and $Stdev(\mathbf{Y})$.

**(b)** According to the Central Limit Theorem, an average of 36 random variables drawn from the probability distribution of **X** should have approximately what probability distribution? (Be sure to specify the mean and standard deviation.)

**(c)** In a spreadsheet, make a simulation table that tabulates values of five random variables as follows:

The first is a single cell that simulates **X**

The second is a single cell that simulates **Y**

The third is an average of 36 cells independently drawn from the probability distribution of **X**

The fourth is an average of 36 cells independently drawn from the probability distribution of **Y**

The fifth is a single random cell drawn from the probability distribution that you predicted in (b)

Myerson_C02.qxd 09-09-2004 10:31 Page 79

Include at least 400 data rows in your simulation table. (This calculation may take a few minutes on older computers.)

**(d)** Using your simulation table in (c), compute the sample *mean* and *standard deviation* for each of the five random variables, and make an *XY* chart that estimates the (inverse) *cumulative distribution* for these five random variables. (*Hint on charting keystrokes:* You can separately sort each of your five columns of simulation data, then select the percentile index and five sorted data columns in the simulation table, and insert an *XY* chart.)

**2.** In a simulation table with data from 400 independent simulations of a random variable **W**, the sample mean is 220.12, and the sample standard deviation is 191.63.

**(a)** Estimate the standard deviation of the sample mean when the sample size is 400.

**(b)** Based on this data, compute a 95% confidence interval for the true expected value of this random variable, $E(\mathbf{W})$.

**(c)** Suppose we want to make a new table of simulation data that will generate a 95% confidence interval for $E(\mathbf{W})$ with a radius of about 5. How large should this new simulation table be? (That is, how many independent simulations should it include?)

**3.** What is the discrete probability distribution of the random variable that would be generated by each of the following Excel formulas? Check your answer by a large simulation.

**(a)** =IF(RAND( )>0.3,2,0)+IF(RAND( )>0.6,3,0)

**(b)** =IF(AND(0.3<RAND( ),RAND( )<0.4),1,0)

**(c)** =IF(RAND( )<0.6,IF(RAND( )<0.5,1,2),3)

**(d)** How would your answers change if we entered =RAND( ) into cell A1 and we replaced every RAND( ) in the preceding formulas by a reference to cell A1?

**4.** Acme Widget Company has substantial uncertainty about many factors that will affect its profit from selling widgets next year. Acme's director of marketing estimates that the total demand for widgets (sold by all firms in the market) may be 60,000 or 70,000 or 80,000 widgets next year, and the probabilities of these three possibilities are 0.2, 0.6, and 0.2, respectively. Acme's share of this total market for widgets next year may be 0.15 or 0.20 or 0.25 or 0.30, with probabilities 0.2, 0.3, 0.3, and 0.2, respectively. The price of widgets next year may be \$90 or \$100 or \$110 per widget, with probabilities 0.2, 0.7, and 0.1, respectively.

Suppose that after the assembly line is set up, Acme can produce its widgets as customers order them, and so Acme's production quantity will equal its demand. Acme's costs can be separated into two parts: a fixed cost of setting up the assembly line, and a variable cost per widget produced. Acme's production manager estimates that the fixed cost may be \$450,000 or \$500,000 or \$550,000 or \$600,000, each with probability 0.25. The variable cost per widget may be either \$50 per widget, with probability 0.6, or \$60 per widget, with probability 0.4.

**(a)** Make a spreadsheet simulation model to represent this situation, assuming that these unknown quantities are independent.

**(b)** Generate a large table of simulated profits, and compute a 95% confidence interval for Acme's expected profit from widget production next year. Make sure that the simulation table is large enough that the radius of your 95% confidence interval for Acme's expected profit is less than \$10,000. (That is, get a 95% confidence interval of the form $m \pm r$, where $r < 10{,}000$.)

**(c)** Using the simulation data from part (b) also estimate:

**(i)** The standard deviation of Acme's profit next year

**(ii)** The probability of Acme's profit next year being negative

**(iii)** The median level of Acme's profit next year

**(iv)** The level of Acme's profit next year that has cumulative probability 0.75

**(d)** Based on the simulation data from part (b), make a chart showing the cumulative risk profile for Acme's profit next year.

Myerson_C02.qxd 09-09-2004 10:31 Page 80

Myerson_C03.qxd 09-09-2004 10:32 Page 81

# 3

# Utility Theory with Constant Risk Tolerance

In the previous chapter, we introduced the criterion of maximizing the expected monetary value of profit or income for decision makers as a basis for determining their optimal decisions under uncertainty. But this simple expected value formula, although intuitive and straightforward to compute, may not accurately describe everyone's willingness to take risks. So in this chapter, we introduce the basic ideas of utility theory, which is a more sophisticated method for computing optimal decisions that can take account of a decision maker's personal willingness to accept risks.

The truth is that people's attitudes toward risk taking can be very complex. How you feel about a gamble that could substantially increase or decrease your income may depend on many personal details about how such changes of income would affect your lifestyle. So a decision theory that has enough degrees of freedom to accurately fit everybody's personal feelings about risk taking may have to be very complex and difficult to apply.

Thus, a practical theory of risk preferences should include some parameter or parameters that allow us to take account of the differences among people's attitudes toward risk taking; but to be tractable, the theory must also make some simplifying assumptions that reduce the number of such risk-tolerance parameters that we have to measure. For these reasons, we will emphasize here a technical assumption called *constant risk tolerance*, which implies that a decision maker's attitudes toward risk taking can be summarized by a single "risk-tolerance" parameter. With this assumption, we will see how to measure a decision maker's risk tolerance, and how to use this risk-tolerance measurement to compute an appropriate "certainty equivalent" value for any gamble that the decision maker might face.

Section 3.3 considers more general utility functions, and Sections 3.4 and 3.5 are concerned with the technical derivations of the mathematical formulas that are used in the theory of utility with constant risk tolerance. These sections will not be used elsewhere in the book and can be omitted by readers who have less interest in such matters.

Myerson_C03.qxd 09-09-2004 10:32 Page 82

## CASE ■ Superior Semiconductor (Part B)

As she was looking over the current draft of her decision analysis report on the proposed T-regulator product, Eastmann became concerned about whether she was using the right criteria to evaluate risks for Superior Semiconductor. So she went to Suttcliff's office to ask him one more question.

"The prescription of a complete decision analysis may depend on our company's willingness to take financial risks," Eastmann began. "For bottom-line recommendations, I have been using an expected value criterion which says that a project is worth undertaking if it increases the expected value of profits for our company. But when people are risk averse, they sometimes want to use other criteria that put more weight on the downside risks. To get your assessment of our corporate willingness to take risks, I should ask you about other simple hypothetical situations. So suppose that we had to choose between two options. Under the first option, Superior Semiconductor would be equally likely to either gain $20 million or lose $10 million, each with probability 1/2. Under the second option, Superior would gain $4 million for sure. Notice that the first option has an expected value of $5 million, but it is risky. Given this choice, would you recommend that Superior should take the risky first option, or the safe second option?"

Suttcliff thought for a bit and replied, "You are essentially asking me about whether we should sell the risky option for a price of $4 million. I get paid to make hard decisions about what risks are worth taking. In a situation like this, I might well say that the $4 million was a good price and recommend that we sell our risky option."

Eastmann probed further. "What if the price were somewhat lower, like $3 million or $2 million? What is the least amount for which you would recommend selling the risky option? If we were actually committed to this gamble and could only get out of it at some cost, would you ever recommend paying something to avoid such a risk?"

Suttcliff frowned. "No, an equal chance of gaining $20 million or losing $10 million would be a gamble worth taking, if we had to. They would have to offer us at least $2 million before I would even consider selling it."

# 3.1 Taking Account of Risk Aversion: Utility Analysis with Probabilities

In the decision analysis literature, decision makers are called *risk neutral* if they are willing to base their decisions purely on the criterion of maximizing the expected value of their monetary income. The criterion of maximizing expected monetary value is so simple to work with that it is often used as a convenient guide to decision making even by people who are not perfectly risk neutral. But in many situations, people feel that comparing gambles only on the basis of expected monetary values would take insufficient account of their aversion to risks.

For example, imagine that you had a lottery ticket that would pay you either $20,000 or $0, each with probability 1/2. If you are risk neutral, then you should be unwilling to sell this ticket for any amount of money less than its expected value, which is $10,000. But many risk-averse people might be very glad to exchange this risky lottery for a certain payment of $9000.

Myerson_C03.qxd 09-09-2004 10:32 Page 83

Given any such lottery or gamble that promises to pay an amount of money that will be drawn randomly from some probability distribution, a decision maker's *certainty equivalent* (abbreviated *CE*) of this gamble is the lowest amount of money-for-certain that the decision maker would be willing to accept instead of a gamble. That is, saying that \$7000 is your certainty equivalent of the lottery that would pay you either \$20,000 or \$0, each with probability 1/2, means that you would be just indifferent between having a ticket to this lottery or having \$7000 cash in hand.

In these terms, risk-neutral people are those whose certainty equivalent of any gamble is just equal to its expected monetary value (abbreviated *EMV*). People are *risk averse* if their certainty equivalent of a gamble is less than the gamble's expected monetary value. The difference between the expected monetary value of a gamble and a risk-averse decision maker's certainty equivalent of the gamble is called the decision maker's *risk premium* (abbreviated *RP*) for the gamble. Thus,

$$RP = EMV - CE$$

So if a lottery paying \$20,000 or \$0, each with probability 1/2, is worth \$7000 to you, then your risk premium for this lottery is \$10,000 – \$7000 = \$3000.

When you have a choice among various gambles, you should choose the one for which you have the highest certainty equivalent because it is the one that is worth the most to you. But when a gamble is complicated, you may find it difficult to assess your certainty equivalent for it. The great appeal of the risk-neutrality assumption is that by identifying your certainty equivalent with the expected monetary value, it makes your certainty equivalent something that is straightforward to compute or estimate by simulation. So what we need now is to find more general formulas that risk-averse decision makers can use to compute their certainty equivalents for complex gambles and monetary risks.

A realistic way of calculating certainty equivalents must include some way of taking account of a decision maker's personal willingness to take risks. The full diversity of formulas that a rational decision maker might use to calculate certainty equivalents is described by a branch of economics called *utility theory*. Utility theory generalizes the principle of expected value maximization in a simple but very versatile way. Instead of assuming that people want to maximize their expected monetary values, utility theory instead assumes that each individual has a personal utility function that assigns a utility value to every possible monetary income level that the individual might receive, such that the individual always wants to maximize the expected value of his or her utility. For example, suppose that you have to choose between two gambles, where the random variable **X** denotes the amount of money that you would get from the first gamble and the random variable **Y** denotes the amount of money that you would get from the second gamble. A risk-neutral decision maker would prefer the first gamble if $E(\mathbf{X}) > E(\mathbf{Y})$. But according to utility theory, when $U(x)$ denotes

Myerson_C03.qxd 09-09-2004 10:32 Page 84

your "utility" for getting any amount of money $x$, you should prefer the first gamble if $E(U(\mathbf{X})) > E(U(\mathbf{Y}))$. (Recall from Chapter 2 that, when $\mathbf{X}$ has a discrete distribution, the expected value operator $E(\bullet)$ can be written

$$E(U(\mathbf{X})) = \sum_x P(\mathbf{X} = x) * U(x)$$

where the summation range includes every number $x$ that is a possible value of the random variable $\mathbf{X}$.) Furthermore, your certainty equivalent $CE$ of the gamble that will pay the random monetary amount $\mathbf{X}$ should be the amount of money that gives you the same utility as the expected utility of the gamble. Thus, we have the basic equation

$$U(CE) = E(U(\mathbf{X}))$$

Utility theory can account for risk aversion, but it also is consistent with risk neutrality or even risk-seeking behavior, depending on the shape of the utility function. In 1947, Von Neumann and Morgenstern gave an ingenious argument to show that any consistent rational decision maker should choose among risky gambles according to utility theory. Since then, decision analysts have developed techniques to assess individuals' utility functions. Such assessment can be difficult, because people have difficulty thinking about decisions under uncertainty and because there are so many possible utility functions. But we may simplify the process of assessing a decision maker's personal utility function if we can assume that the utility function is in some mathematically natural family of utility functions.

For practical decision analysis, the most convenient utility functions to use are those that have a special property called *constant risk tolerance*. Constant risk tolerance means that if we change a gamble by adding a fixed additional amount of money to the decision maker's payoff in all possible outcomes of the gamble, then the certainty equivalent of the gamble should increase by this same amount. This assumption of constant risk tolerance is very convenient in practical decision analysis. One nice consequence of constant risk tolerance is that it allows us to evaluate independent gambles separately. If you have constant risk tolerance, then when you are going to earn money from two independent gambles, your certainty equivalent for the sum of the two independent gambles is just the sum of your certainty equivalent for each gamble by itself. That is, the lowest price at which you would be willing to sell a gamble is not affected by having other independent gambles. This independence property only holds if you have one of these constant-risk-tolerance utility functions.

If a risk-averse decision maker's preferences over gambles satisfy this assumption of constant risk tolerance, then the decision maker must have a utility function $U$ in a simple one-parameter family of functions that are defined by the mathematical formula

$$U(x) = -\text{EXP}(-x/r)$$

where the parameter $r$ is called the *risk-tolerance constant*. (Here EXP is a standard function in Excel that will be discussed at length in Chapter 4.)

Myerson_C03.qxd 09-09-2004 10:32 Page 85

Simtools includes a special function called UTIL that can be used in place of this formula. That is, for any amount $x$ of monetary income and for any positive risk-tolerance constant $r$, Simtools's UTIL function returns the utility value

$$\text{UTIL}(x, r) = -\text{EXP}(-x/r)$$

Thus, if we can assume that a decision maker has constant risk tolerance, then we only need to measure this one risk-tolerance parameter $r$ for the decision maker. By asking the decision maker to subjectively assess the personal certainty equivalent for one simple gamble, we can get enough information to compute the risk tolerance that accounts for this personal assessment. Then, once we have found an appropriate risk tolerance for this decision maker, we will be able to use it to estimate the decision maker's expected utility and certainty equivalent for any gamble that we can simulate, no matter how complex.

It is natural that this numerical measure of "risk tolerance" seem mysterious at first. The meaning of these risk-tolerance numbers will become clearer as you get practical experience using them.

In Part B of the Superior Semiconductor case, the decision maker expresses attitudes toward risk that are incompatible with the assumption of risk neutrality. In particular, he says that a gamble where the Superior Semiconductor Company could either gain \$20 million or lose \$10 million, each with probability 1/2, might be worth only \$2 million to the company, even though the expected monetary value of this gamble is \$5 million. That is, Superior Semiconductor's certainty equivalent for this lottery is subjectively assessed at \$2 million (less than the expected monetary value by a risk premium of \$3 million). But the expected utility of a lottery must be equal to the utility of its certainty equivalent. If Superior Semiconductor acts as a decision maker with a risk tolerance constant equal to $r$, then the gamble paying either \$20 million or –\$10 million, each with probability 0.5, would have expected utility

$$0.5*\text{UTIL}(20, r) + 0.5*\text{UTIL}(-10, r)$$

while a payoff of \$2 million would have utility

$$\text{UTIL}(2, r)$$

Thus, the only way to fit the decision maker's expressed preferences for risk taking here is to find the value of $r$ that satisfies the equation

$$0.5*\text{UTIL}(20, r) + 0.5*\text{UTIL}(-10, r) = \text{UTIL}(2, r)$$

This equation is satisfied when $r = 36.49$, as shown in cells B4 and C4 of Figure 3.1. Thus, the risk aversion expressed here by the vice president of Superior Semiconductor in Part B is consistent with constant risk tolerance only when the risk tolerance parameter is $r$ = \$36.49 million. This utility function $U(x) = \text{UTIL}(x, 36.49)$, for monetary incomes $x$ from –20 to 70 (in millions of dollars), is plotted in Figure 3.1.

Myerson_C03.qxd 09-09-2004 10:32 Page 86

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | RiskTolerance | | | | | | | |
| 2 | 36.49 | | | | | | | |
| 3 | | EU | U(CE) | | | | | |
| 4 | | -0.9466642 | -0.9466642 | | | | | |
| 5 | | | | | | | | |
| 6 | $X | UTIL(X,RT) | UINV(U,RT) | | | | | |
| 7 | -20 | -1.7299782 | -20 | | | | | |
| 8 | -15 | -1.5084485 | -15 | | | | | |
| 9 | -10 | -1.3152864 | -10 | | | | | |
| 10 | -5 | -1.1468593 | -5 | | | | | |
| 11 | 0 | -1 | -0 | | | | | |
| 12 | 5 | -0.8719465 | 5 | | | | | |
| 13 | 10 | -0.7602907 | 10 | | | | | |
| 14 | 15 | -0.6629328 | 15 | | | | | |
| 15 | 20 | -0.578042 | 20 | | | | | |
| 16 | 25 | -0.5040217 | 25 | | | | | |
| 17 | 30 | -0.4394799 | 30 | | | | | |
| 18 | 35 | -0.383203 | 35 | | | | | |
| 19 | 40 | -0.3341325 | 40 | | | | | |
| 20 | 45 | -0.2913457 | 45 | FORMULAS FROM RANGE A6:C26 | | | | |
| 21 | 50 | -0.2540378 | 50 | A2. =RISKTOL(20,-10,2) | | | | |
| 22 | 55 | -0.2215074 | 55 | B4. =0.5*UTIL(20,A2)+0.5*UTIL(-10,A2) | | | | |
| 23 | 60 | -0.1931426 | 60 | C4. =UTIL(2,A2) | | | | |
| 24 | 65 | -0.16841 | 65 | B7. =UTIL(A7,$A$2) | | | | |
| 25 | 70 | -0.1468445 | 70 | C7. =UINV(B7,$A$2) | | | | |
| 26 | | | | B7:C7 copied to B7:C25. | | | | |
| 27 | | | | Chart plots (A7:A25,B7:B25). | | | | |

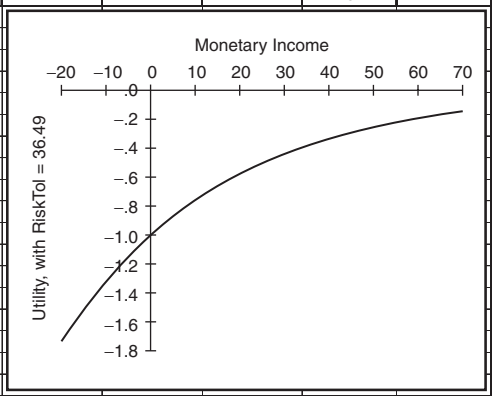


FIGURE **3.1**
**Utility function with constant risk tolerance for Superior Semiconductor (B)**

This technique can be extended to give us a practical way of assessing a risk tolerance constant for any risk-averse decision maker. We ask the decision maker to consider a simple binary lottery that will pay either some high monetary amount, $*H*, or some low monetary amount, $*L*, each with probability 0.5, and we can ask the decision maker to assess the monetary amount $*C* such that she would be indifferent between playing this lottery and getting $*C* with certainty. Then the decision maker's expressed preferences can be consistent with utility theory and constant risk tolerance only if the risk-tolerance constant *r* satisfies

$$0.5*\text{UTIL}(H, r) + 0.5*\text{UTIL}(L, r) = \text{UTIL}(C, r)$$

With Simtools, the value of *r* that solves this equation can be automatically calculated by the function RISKTOL(*H*, *L*, *C*). For example, RISKTOL(20, –10, 2) is equal to 36.49, as shown in cell A2 of Figure 3.1.

**NOTE** (You could also use Excel's GoalSeek tool to solve this utility equation. To see how, consider a spreadsheet like Figure 3.1 with =0.5*UTIL(20,A2) + 0.5*UTIL(–10,A2) in cell B4 and =UTIL(2,A2) in cell C4. But now suppose that any positive number was entered into cell A2 instead of the RISKTOL formula, and the formula =B4 –C4 was entered into cell C1. Then we could find the risk tolerance that makes *EU* = *U*(*CE*) with Excel's Tools:GoalSeek command by asking it to set this cell C1 to the value 0 by changing cell A2.)

Myerson_C03.qxd 09-09-2004 10:32 Page 87

The risk tolerance constant is a number measured in the same units that we are using to measure money. In our analysis of the Superior Semiconductor case, we are measuring all monetary quantities in millions of dollars, and so the numerical value

RISKTOL(20,–10,2) = 36.49

in cell A2 of Figure 3.1 should be read as "36.49 million dollars." For example, if we did our accounting in thousands of dollars instead of millions, then the decision maker could equivalently say that a lottery paying either 20000 $thousand or –10000 $thousand with equal probability would be as good as a certain 2000 $thousand for the company, and RISKTOL should yield the equivalent risk tolerance

RISKTOL(20000, –10000, 2000) = 36490 thousand dollars

To assess your own personal degree of risk tolerance, try thinking about a lottery ticket that would pay you either $20,000 or $0, each with probability 1/2. If you are risk averse, then you would be willing to sell that ticket for a price slightly below its expected value of $10,000. But then think about lower prices until you find the certainty-equivalent price at which you would be just indifferent between taking this amount of money (with no risk) or playing the risky lottery. For example, you might decide that this lottery ticket would be worth just $7000 to you, and then (assuming that you have constant risk tolerance) your risk tolerance must be RISKTOL (20000, 0, 7000) = $15,641. A more risk-averse individual whose certainty equivalent for this same lottery ticket was only $5000 would have a lower risk tolerance of RISKTOL(20000, 0, 5000) = $8,205. A more risk-tolerant individual whose certainty equivalent for this same lottery ticket was $9900 would have a much higher risk tolerance, RISKTOL(20000, 0, 9500) = $99,833. Lower levels of risk-tolerance would correspond to certainty equivalents that are farther below the expected monetary value and closer to the worst possible outcome. Higher levels of risk tolerance correspond to certainty equivalents that are closer to the expected monetary value. A risk-neutral individual, whose certainty equivalent for this lottery would be the expected monetary value $10,000, has an infinite risk tolerance, which RISKTOL (20000, 0, 10000) represents with a "dividing by zero" error message.

Decision analysts like to assess risk tolerance by asking about simple gambles that involve just two equally likely outcomes, each with probability 1/2, because such simple binary gambles are easiest for people to think about. So the basic idea is to ask the decision maker to assess his or her certainty equivalent for such a simple binary gamble, compute a risk tolerance with Simtools's RISKTOL function, and then apply that risk tolerance constant to evaluate the more complex risks that really confront the decision maker.

But which high and low prizes should we ask the decision maker to think about? In theory, if the decision maker is perfectly rational and satisfies our assumption of constant risk tolerance, then it would not matter which high and low prizes we asked about: the assessed certainty equivalent would

Myerson_C03.qxd 09-09-2004 10:32 Page 88

always yield the same risk tolerance parameter. But in practice, we may get different risk tolerances when we ask about different binary lotteries, because the assumption of constant risk tolerance is not a perfect fit for most people's preferences. So it is best to assess risk tolerances by using hypothetical binary lotteries where the levels and spread of the two prizes are similar to the levels and spread of the possible outcomes in the actual decision situation that we are studying. So in Part B of the Superior Semiconductor case, the decision analyst may reasonably ask the decision maker to think about a hypothetical binary lottery that would generate either some profit of $20 million or a loss of $10 million, each with probability 1/2, because these values are nice round figures (to be easier to think about) and are close to the best and worst outcomes of $24 million and –$9.33 million that could occur in the discrete Superior Semiconductor case.

After we have measured a decision maker's risk tolerance, we can use it to compute an expected utility for any monetary gamble, as shown in Figure 3.2. Here the appropriate risk tolerance in this case is computed (using the RISKTOL function) in cell A2. The possible profit outcomes of the T-regulator project are listed in cells A6:A10 of Figure 3.2, and the corresponding probabilities are listed in cells B6:B10. Then the corresponding utility values are calculated in cells D6:D10, by entering

```
=UTIL(A6,$A$2)
```

into cell D6 and then copying D6 to D6:D10. (The equivalent exponential formula =–EXP(–A6/$A$2) is entered into cell F6 for comparison.) So the expected utility of this gamble for Superior Semiconductor is computed from the given probabilities and utilities by the formula

```
=SUMPRODUCT(D6:D10,$B$6:$B$10)
```

in cell D12. We find here that the expected utility of this gamble has the numerical value –0.9884, which is very hard to interpret. There is no significance to its negativity. In fact, our UTIL function with constant risk tolerance always yields negative utility values, even for very good prizes. For example, the utility of $24 million is shown in cell D6 to be –0.5180. So how do we interpret these expected utilities?

The answer is that we cannot interpret any one expected utility by itself, but we can only understand it in comparison with expected utilities of other alternatives that the decision maker could choose. In any such comparison, utility theory predicts that the decision maker should prefer the gamble with the higher expected utility. For example, the alternative of terminating the T-regulator product and getting $0 profit from it would have utility UTIL(0, 36.49) = –1 in this scale (as shown in cell B11 of Figure 3.1). The expected utility –0.9884 of developing the T-regulator product is higher than –1, and so the decision maker at Superior Semiconductor should prefer to develop the T-regulator product, even though it is risky.

To avoid the confusion of trying to interpret expected utility numbers, we should always convert them back into monetary units by asking what

Myerson_C03.qxd 09-09-2004 10:32 Page 89

| | A | B | C | D | E | F | G | H | I | J |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | DecisionMaker's RiskTolerance | | | | | From: | High$ | 20.00 | | |
| 2 | 36.49 | | | | | | Low$ | -10.00 | | |
| 3 | | | | | | Assessed CE | | 2.00 | K | |
| 4 | CALCULATING CERTAINTY EQUIVALENT OF DISCRETE LOTTERY | | | | | | | | 26.00 | |
| 5 | $Prize | Proby | | Util(Prize,RT) | | Exponential formula | | | Prize+K | Util |
| 6 | 24.00 | 0.1 | | -0.5180 | | -0.5180 | | | 50.00 | -0.2540 |
| 7 | 7.33 | 0.25 | | -0.8179 | | -0.8179 | | | 33.33 | -0.4011 |
| 8 | -1.00 | 0.3 | | -1.0278 | | -1.0278 | | | 25.00 | -0.5040 |
| 9 | -6.00 | 0.25 | | -1.1787 | | -1.1787 | | | 20.00 | -0.5780 |
| 10 | -9.33 | 0.1 | | -1.2915 | | -1.2915 | | | 16.67 | -0.6333 |
| 11 | EMV | CE | | EU | CE from EU | | RP | | CE(Prize+K) | EU |
| 12 | 1.50 | 0.42 | | -0.9884 | 0.42 | 0.42 | 1.08 | | 26.42 | -0.4847 |
| 13 | | | | | | | | | | |
| 14 | ESTIMATES FROM SIMULATION: | | | | | | Calculations from utilities: | | | |
| 15 | | 2.00 | | Expected Monetary Value | | | EU | -0.9799 | | |
| 16 | | 0.74 | | Certainty Equivalent | | | CE | 0.74 | | |
| 17 | | 1.26 | | Risk Premium | | | Stdev(U) | 0.24 | | |
| 18 | | $Prize | | | | | Count | 401 | | |
| 19 | SimTabl | -1.00 | | Utility | | | 95% confidence intervals: | | | |
| 20 | 0 | 7.33 | | -0.8179 | | | EU | -1.0030 | -0.9569 | |
| 21 | 0.0025 | -9.33 | | -1.2915 | | | CE | -0.11 | 1.61 | |
| 22 | 0.005 | -1.00 | | -1.0278 | | | | | | |
| 23 | 0.0075 | 7.33 | | -0.8179 | FORMULAS FROM RANGE A1:J30 | | | | | |
| 24 | 0.01 | 7.33 | | -0.8179 | A2. =RISKTOL(H1,H2,H3) | | | | | |
| 25 | 0.0125 | -6.00 | | -1.1787 | D6. =UTIL(A6,$A$2) | | | D6 copied to D6:D10 | | |
| 26 | 0.015 | -6.00 | | -1.1787 | F6. =-EXP(-A6/$A$2) | | | F6 copied to F6:F10 | | |
| 27 | 0.0175 | 7.33 | | -0.8179 | A12. =SUMPRODUCT(A6:A10,$B$6:$B$10) | | | | | |
| 28 | 0.02 | -9.33 | | -1.2915 | B12. =CEPR(A6:A10,B6:B10,A2) | | | | | |
| 29 | 0.0225 | -9.33 | | -1.2915 | D12. =SUMPRODUCT(D6:D10,$B$6:$B$10) | | | | | |
| 30 | 0.025 | -1.00 | | -1.0278 | E12. =UINV(D12,$A$2) | | | F12. =-A2*LN(-D12) | | |
| 31 | 0.0275 | 7.33 | | -0.8179 | G12. =A12-B12 | | | | | |
| 32 | 0.03 | -1.00 | | -1.0278 | I6. =A6+$I$4 | | | I6 copied to I6:I10 | | |
| 33 | 0.0325 | 24.00 | | -0.5180 | J6. =UTIL(I6,$A$2) | | | J6 copied to J6:J10 | | |
| 34 | 0.035 | 7.33 | | -0.8179 | J12. =SUMPRODUCT(J6:J10,$B$6:$B$10) | | | | | |
| 35 | 0.0375 | 24.00 | | -0.5180 | I12. =UINV(J12,$A$2) | | | | | |
| 36 | 0.04 | 7.33 | | -0.8179 | B19. =DISCRINV(RAND(),A6:A10,B6:B10) | | | | | |
| 37 | 0.0425 | -9.33 | | -1.2915 | B15. =AVERAGE(B20:B420) | | | | | |
| 38 | 0.045 | -6.00 | | -1.1787 | B16. =CE(B20:B420,A2) | | | B17. =B15-B16 | | |
| 39 | 0.0475 | -1.00 | | -1.0278 | D20. =UTIL(B20,$A$2) | | | D20 copied to D20:D420 | | |
| 40 | 0.05 | 7.33 | | -0.8179 | H15. =AVERAGE(D20:D420) | | | | | |
| 41 | 0.0525 | -1.00 | | -1.0278 | H16. =UINV(H15,A2) | | | | | |
| 42 | 0.055 | -1.00 | | -1.0278 | H17. =STDEV(D20:D420) | | | H18. =COUNT(D20:D420) | | |
| 43 | 0.0575 | -1.00 | | -1.0278 | H20. =H15-1.96*H17/H18^0.5 | | | | | |
| 44 | 0.06 | -1.00 | | -1.0278 | I20. =H15+1.96*H17/H18^0.5 | | | | | |
| 45 | 0.0625 | -1.00 | | -1.0278 | H21. =UINV(H20,$A$2) | | | I21. =UINV(I20,$A$2) | | |

FIGURE **3.2**
**Analysis of Superior Semiconductor (B) with constant risk tolerance**

sure amount of money would also yield this same expected utility. This amount of money is then the certainty equivalent of the lottery. Given any positive risk-tolerance constant $r$, and given any expected utility value $u$, the monetary amount $x$ that would have utility $u$ is returned by the Simtools function UINV($u$, $r$), which is defined by the formula

$$\text{UINV}(u, r) = -r * \text{LN}(-u)$$

(LN is a standard function in Excel that will be discussed in Chapter 4.) So with this function,

$$\text{if } x = \text{UINV}(u, r) \text{ then } u = \text{UTIL}(x, r)$$

Myerson_C03.qxd 09-09-2004 10:32 Page 90

In this example, we should ask what amount of profit would have the same utility value –0.9884, and the answer of $0.42 million is computed in cell E12 of Figure 3.2 by the formula

=UINV(D12,$A$2)

(The defining formula for the UINV function is applied in cell F12, for comparison.) You can check that UTIL(0.42, 36.49) is indeed very close to –0.9884. Thus, $0.42 million is Superior Semiconductor's certainty equivalent for the risky monetary returns that it can get from this T-regulator project, given the degree of risk tolerance that has been expressed by the vice president in Part B of the case. Superior Semiconductor's risk premium for this T-regulator gamble is then $EMV - CE = 1.50 - 0.42 = \$1.08$ million, as shown in cell G12 of Figure 3.2.

So after assessing an appropriate risk-tolerance constant for our decision maker, there are three steps in the utility analysis of any gamble with discrete probabilities. First we convert monetary payoffs into utility values, using the UTIL function. Then we compute the expected utility value for this gamble, using SUMPRODUCT with the given probabilities and utility values. Finally we convert this expected utility back to a monetary certainty equivalent using the UINV function.

Simtools includes a shortcut function called CEPR (for computing Certainty Equivalents for monetary gambles with discrete PRobabilities) that does this whole three-step analysis at once. The three required parameters of the CEPR function are: first a range of cells listing the possible monetary income amounts, second a range of cells containing the corresponding discrete probabilities of these incomes, and third the decision maker's risk-tolerance constant. (The Insert:Function dialog box also lists two optional parameters for CEPR, but we will ignore these here.) For the example in Figure 3.2, the possible income values are in cells A6:A10, their corresponding probabilities are in cells B6:B10, and the risk tolerance is in cell A2, so the formula

=CEPR(A6:A10, B6:B10, A2)

in cell B12 returns the same certainty equivalent that we got by utility analysis in cell E12. In addition to being easy to use, CEPR also has the advantage that it lets us see the results of utility analysis without having to look at those hard-to-interpret utility numbers. But we should look again at the utility graph in Figure 3.1 to understand how such utility analysis can implement our intuitive notions of risk aversion.

Looking at the graph of the utility function in Figure 3.1, you should notice that the higher monetary income levels are associated with higher utility levels, but the rate of utility increase becomes slower as income increases. This shape of the utility function means that an extra dollar when income is low would increase utility by more than an extra dollar when income is high. It is this greater sensitivity to income changes when poor that makes people risk averse. For example, relative to a basic alternative

Myerson_C03.qxd 09-09-2004 10:32 Page 91

of earning (say) $5 million from some safe investment, the utility impact of a $15 million decrease of monetary income (taking net income down to –$10 million) would be strictly greater than the utility impact of a $15 million increase (taking monetary income up to $20 million). That is,

$$(-0.872) - (-1.315) = 0.433 > 0.294 = (-0.578) - (-0.872)$$

Thus, the expected utility of a lottery ticket that pays either $20 million or –$10 million with equal probability is strictly less than the utility of the sure $5 million

$$0.5*(-1.315) + 0.5*(-0.578) = -0.947 < -0.872$$

So when we compare alternatives by such expected utilities, we find that a sure amount of money always looks strictly better than a risky lottery that has the same expected monetary value. If we had assessed a different level of risk tolerance (in cell A2), then the utility curve in Figure 3.1 would have a somewhat different shape, becoming straighter with higher risk tolerance. But the pattern of decreasing slope of utility at higher levels of monetary income would hold with any positive risk tolerance.

The constant-risk-tolerance property is illustrated in cells I3:J12 of Figure 3.2, which show what happens if the prizes in our lottery are all increased by some constant amount of money (here $26 million, given in cell I4). Cell I12 computes the certainty equivalent of the lottery that offers these increased prizes (in I6:I10) with the original probabilities (in B6:B10). Comparing I12 and B12, you can see that increasing the prizes by 26 has increased the certainty equivalent by the same amount. Thus, because we are using a UTILity function that has the special property of constant risk tolerance, computing the certainty equivalent of the risky gross revenues ($26.42 million) and then subtracting the known cost ($26 million) yields the same answer ($0.42 million) as computing the certainty equivalent of net profits.

You should know that constant risk tolerance does not imply that certainty equivalents would vary proportionally with the size of the gamble. For example, if we doubled all the monetary prizes here (in cells A6:A10), then the certainty equivalent (in cell B12) would not double. In fact, it would decrease to –$0.91 million with the given risk tolerance, because doubling the magnitude of the risks would cause the constant-risk-tolerant decision maker's risk premium to more than double.

Because risk tolerances are hard to assess, it is a good idea to do sensitivity analysis in the risk-tolerance parameter and see how much it would have to change before the recommendations of our decision analysis would substantially change. With any positive risk tolerance, the certainty equivalent of a gamble will always be between the lowest possible monetary outcome (here, –9.33) and the expected monetary value (here, 1.5). A greater risk tolerance would yield a higher certainty equivalent for any gamble. A very small positive risk tolerance would yield a certainty equivalent close to the lowest possible outcome, but a very large risk tolerance would yield a certainty equivalent close to the expected monetary value.

Myerson_C03.qxd 09-09-2004 10:32 Page 92

Using the spreadsheet shown in Figure 3.2, you can verify that the T-regulator project in this example would have a certainty equivalent (in cell B12 or E12) greater than $0 as long as the risk-tolerance constant in cell A2 is greater than $25.02 million. You can also check that the computed risk tolerance in A2 would be greater than $25.02 million for Superior Semiconductor as long as the decision maker's subjectively assessed certainty equivalent for the given simple binary lottery in cell H3 is greater than $0.75 million.

## 3.2 Utility Analysis from Simulation Data

In the preceding section, we learned how to use utility analysis to compute certainty equivalents for simple gambles where our monetary income will be drawn from a given discrete probability distribution. In more complex gambles, we may not be able to find such a simple discrete probability distribution, but we may still be able to construct an appropriate simulation model. So in this section, we study techniques for estimating an expected utility value and certainty equivalent of any complex gamble that we can describe by a simulation model.

By the law of large numbers, we can estimate the expected utility of any gamble by computing the average of our utility values for the results of many independent simulations of this gamble. Then we can estimate the certainty equivalent by applying the UINV function (with the appropriate risk tolerance) to this estimated expected utility.

As in Section 2.5 of Chapter 2, it may be helpful to illustrate these techniques in the context of the simple Superior Semiconductor (B) example, where we know how to explicitly compute the numbers that we are trying to estimate from simulation, because this simplicity will enable you to get a hands-on feel for the accuracy of these simulation methods. The potential profit from this project is simulated in cell B19 of Figure 3.2, and data from 401 independent simulations of Superior Semiconductor's profit in this model are stored in the range B20:B420 of this spreadsheet. With the assessed risk tolerance $36.49 million in cell A2, the corresponding utility values for each of the simulated profits are computed by entering

```
=UTIL(B20,$A$2)
```

into cell D20 and then copying D20 to D20:D420. Then our estimate of the expected utility from this gamble is computed in cell H15 by the formula

```
=AVERAGE(D20:D420)
```

As usual, because utility values have no clear interpretation, we should convert this expected utility value to its monetary certainty equivalent by the UINV function. So our estimate of Superior Semiconductor's certainty

Myerson_C03.qxd 09-09-2004 10:32 Page 93

equivalent for this gamble (here, $0.74 million) is computed in cell H16 of Figure 3.2 by the formula

=UINV(H15,A2)

Simtools has a shortcut function called CE that computes such estimated certainty equivalents directly from simulated profit data. For this example, the formula

=CE(B20:B420,A2)

has been entered into cell B16, where it returns the same estimated certainty equivalent value of $0.74 million, exactly equal to the estimated certainty equivalent value computed by UINV in cell H16. In general, when the CE function is given any range of monetary incomes as its first parameter and any positive risk tolerance as its second parameter, it automatically does the following three-step calculation. First, it computes the utility values with the given risk tolerance for each of the monetary incomes. Second, it computes the average of these utility values. Third, applying UINV with the given risk tolerance, it computes the certainty equivalent that corresponds to this average of the utilities, and it returns this certainty equivalent. Thus, the value of this CE function is the monetary payoff that should be considered just as good as a random draw from the distribution of the simulated monetary payoffs in the range specified by CE's first parameter, when the decision maker has the constant risk tolerance specified by CE's second parameter.

So the CE shortcut function allows us to do utility analysis without looking at those hard-to-interpret utility numbers. We can use RISKTOL to compute an appropriate risk tolerance number from the subjectively assessed certainty equivalent of any simple binary lottery with two equally likely prizes. Then we can estimate a certainty equivalent value for any gamble that we can simulate by applying the CE function to a large sample of simulated monetary incomes with the appropriate risk tolerance from RISKTOL listed as CE's second parameter.

Thus, based on the simulation data in this spreadsheet (of which only a small part is shown in Figure 3.2), we have found in cells B16 and H16 an estimated certainty equivalent of $0.74 million, which is somewhat higher than the actual certainty equivalent ($0.42 million) that we were able to compute from the underlying probability distribution. To get a more accurate estimate, we could make a larger simulation table, with data from more independent simulations.

When we want to make a 95% confidence interval for the certainty equivalent based on such simulation data, we have to go back to the underlying utility values. A 95% confidence interval for the true expected utility can be computed around the average of our simulated utility values by adding and subtracting the amount 1.96 times the standard deviation of the utility values divided by the square root of the sample size. In this example, with simulated utility values contained in the range D20:D420, the

Myerson_C03.qxd 09-09-2004 10:32 Page 94

lower and upper bounds of a 95% confidence interval for the true expected utility can be computed by the formulas

=AVERAGE(D20:D420) –1.96*STDEV(D20:D420)/COUNT(D20:D420)^0.5

=AVERAGE(D20:D420) +1.96*STDEV(D20:D420)/COUNT(D20:D420)^0.5

These lower and upper bounds for expected utility are shown in cells H20 and I20 in Figure 3.2. Then the corresponding lower and upper bounds for the true certainty equivalent are computed by applying UINV to these lower and upper bounds for expected utility, using the formulas

=UINV(H20,$A$2) and =UINV(I20,$A$2)

in cells H21 and I21 of Figure 3.2. So for this simulation data, the bounds of this 95% confidence interval for the certainty equivalent are –$0.11 million and $1.61 million. Of course, we actually know the true certainty equivalent value for this simple example ($0.42 million, computed in cells B14 and E14), and we can see that it is inside this 95% confidence interval. But if we only had the simulation model, then we could say that this simulation data provides strong evidence of the certainty equivalent being in the interval between –$0.11 million and $1.61 million. A larger table of simulation data is apparently needed before we can narrow this interval and support (at the 95% confidence level) a decisive conclusion that the T-regulator project's certainty equivalent is greater than $0.

## 3.3 The More General Assumption of Linear Risk Tolerance

The most common difficulty with the assumption of constant risk tolerance is that many people would become more risk tolerant as their wealth increased. For example, you might prefer $0 to a gamble in which you would either gain $5000 or lose $4000, each with probability 1/2, even though this gamble has a positive expected monetary value of $500. But now suppose you were suddenly given $100,000. The assumption of constant risk tolerance says that you should still prefer not to accept this gamble, which would leave you either with $105,000 or with $96,000 more than you had when you woke up this morning, when the alternative is to just keep the $100,000. But many people might feel that their concern about the risks of a few thousand dollars would become less urgent when their basic wealth went up and might be willing to take the gamble after receiving the $100,000 windfall. Such people have *increasing risk tolerance*.

Such increasing risk tolerance can be admitted into our analysis only if we consider a more general class of utility functions. A mathematically

Myerson_C03.qxd 09-09-2004 10:32 Page 95

convenient class of utility functions that are commonly used for such purposes is the class of utility functions with *linear risk tolerance*. These utility functions have two risk-tolerance parameters, a *constant r* and a *slope s*. When the slope is 0, then the utility function is the same as a utility function with constant risk tolerance *r*. But when the slope *s* is a positive number, then we get a new kind of utility function. The general formulas are somewhat complicated, but the Simtools functions UTIL, UINV, and CE allow us to do decision analysis with linear risk tolerance when an optional RiskTolSlope parameter is specified. Omitting this optional slope parameter, as we do throughout the rest of this book, is equivalent to specifying a slope equal to 0. (*Warning:* Do not try to model linear risk tolerance by making the risk tolerance constant a linear function of the payoff! Utility functions do not work this way, and doing so would generate nonsensical results.)

Suppose that we have such a decision maker who has linear risk tolerance with constant *r* and slope *s*. For any dollar amount *x*, if this decision maker considers some gambles for which all monetary payoffs are close to *x*, then the certainty equivalents and risk premiums for these gambles would be similar to those that we would compute by assuming a constant risk tolerance equal to $r + s*x$. In this sense, the linear formula $r + s*x$ may be called the decision maker's *local risk tolerance* for monetary payoffs near *x* dollars.

If you want to know more about the exact meaning of local risk tolerance for monetary payoffs near *x*, consider a small number *d*, and consider a gamble that will pay either $x + d$ or $x - d$, each with probability 1/2. If *d* is small, then the certainty equivalent of this gamble will be close to its expected monetary value *x*, but we should look carefully at the risk premium, which is the difference between the certainty equivalent and the expected monetary value. If the decision maker had constant risk tolerance *r*, then this risk premium would be very close to $0.5*(d\text{^}2)/r$ when *d* is small. More generally, if the decision maker has linear risk tolerance with constant *r* and slope *s*, then the risk premium for this $x \pm d$ gamble would be $0.5*(d\text{^}2)/(r + s*x)$. So according to this formula, the risk premium for small gambles around *x* should be approximately proportional to the gamble's variance, which is the square of its standard deviation. This formula also implies that for any small number *d*, the local risk tolerance $(r + s*x)$ for monetary payoffs near *x* should be approximately

$$0.5*(d\text{^}2)/RP(x \pm d)$$

where $RP(x \pm d)$ denotes the decision maker's risk premium for a gamble that pays either $x + d$ or $x - d$, each with probability 1/2. Such calculations of local risk tolerance are illustrated in Figure 3.3.

The assumption of linear risk tolerance probably fits many people's risk preferences better than the narrower assumption of constant risk tolerance. But the assumption of linear risk tolerance is less mathematically convenient, for two reasons. First, the problem of estimating two risk-tolerance

Myerson_C03.qxd 09-09-2004 10:32 Page 96

| | A | B | C | D | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | | | d | | | RiskTolConst | | | RiskTolSlope | | |
| 2 | | | 10 | | | 10000 | | | 0.2 | | |
| 3 | | | | | | | | | | | |
| 4 | | | $Prizes,Pr=1/2 | | | ConstantRiskTolerance | | | LinearRiskTolerance | | |
| 5 | X | | X+d | X-d | | RP | 0.5*(d^2)/RP | | RP | 0.5*(d^2)/RP | |
| 6 | 1000 | | 1010 | 990 | | 0.005 | 10000 | | 0.0049 | 10200 | |
| 7 | 5000 | | 5010 | 4990 | | 0.005 | 10000 | | 0.00455 | 11000 | |
| 8 | 10000 | | 10010 | 9990 | | 0.005 | 10000 | | 0.00417 | 12000 | |
| 9 | 20000 | | 20010 | 19990 | | 0.005 | 10000 | | 0.00357 | 14000 | |
| 10 | 50000 | | 50010 | 49990 | | 0.005 | 10000 | | 0.0025 | 20000 | |
| 11 | | | | | | | | | | | |
| 12 | FORMULAS | | | | | | | | | | |
| 13 | C6. =A6+$C$2 | | | | | | | | | | |
| 14 | D6. =A6-$C$2 | | | | | | | | | | |
| 15 | F6. =A6-CE(C6:D6,$F$2) | | | | | | | | | | |
| 16 | G6. =0.5*($C$2^2)/F6 | | | | | | | | | | |
| 17 | I6. =A6-CE(C6:D6,$F$2,$I$2) | | | | | | | | | | |
| 18 | J6. =0.5*($C$2^2)/I6 | | | | | | | | | | |
| 19 | C6:J6 copied to C6:J10 | | | | | | | | | | |

FIGURE **3.3**
**Estimating local risk tolerances, with constant or linear risk tolerance**

parameters (the constant and the slope) for a decision maker is significantly more complicated than the estimation of one risk-tolerance constant. Second, the assumption of linear risk tolerance cannot be applied when the risk tolerance becomes negative. For example, an assumption of linear risk tolerance with constant $r = 10{,}000$ and slope $s = 0.2$ cannot be applied if any monetary payoffs are below –50,000. Payoffs below this amount would generate infinitely negative utilities, yielding an error in Excel.

## 3.4 Advanced Technical Note on Utility Theory

John Von Neumann and Oskar Morgenstern showed in 1947 how simple rationality assumptions could imply that any decision makers' risk preferences should be consistent with utility theory. In this section, we discuss a simplified version of their argument, which justifies our use of utility theory.

To illustrate the logic of the argument, let us begin with an example. Suppose that we want to study a decision maker's preferences among gambles that involve possible prizes between (say) $0 and $5000. We might start by asking whether he would prefer to get $1000 for sure or a lottery ticket that would pay $5000 with probability 0.20 and would pay $0 otherwise. If the decision maker were risk neutral, he would express indifference between these two alternatives (as they both have expected monetary value of $1000), but suppose that our decision maker expresses a clear preference for the sure $1000 over the lottery ticket. Next, we might ask whether he would prefer the sure $1000 or the lottery ticket that pays either $5000 or $0 if its probability of paying $5000 were increased to 0.30. Suppose that

Myerson_C03.qxd 09-09-2004 10:32 Page 97

the decision maker now responded that he would be willing to give up a sure $1000 for such a lottery ticket. So somewhere between 0.20 and 0.30 there should be some number $p$ such that this decision maker would be indifferent between getting $1000 for sure and getting a lottery ticket that pays $5000 with probability $p$ and pays $0 with probability $1 - p$. Suppose we ask the decision maker to think about such lotteries and tell us what probability $p$ would make him just indifferent, and he tells us (perhaps after some long thought) that he would be just indifferent between $1000 for sure and a lottery ticket that pays $5000 with probability 0.27 and pays $0 otherwise.

Next, we might ask this decision maker a similar question about $2000 for sure: "For what probability $p$ would you be indifferent between getting $2000 for sure and a lottery ticket that would pay $5000 with probability $p$ and $0 otherwise?" Suppose that in answer to this question, the decision maker says that he would be just indifferent between $2000 for sure and a lottery ticket that pays $5000 with probability 0.50 and pays $0 otherwise.

Now consider the following two gambles:

Gamble 1 pays $2000 with probability 0.5 and pays $1000 otherwise.

Gamble 2 pays $5000 with probability 0.4 and pays $0 otherwise.

Given only the preceding information about this decision maker, should we be able to predict which of these two gambles he should prefer?

The answer to this question is Yes. If his preferences are logically consistent, he should prefer Gamble 2 here. This is why. Suppose we offered to exchange Gamble 1 for the following two-stage gamble: At stage 1 we toss a fair coin, and if the coin comes up Heads (which has probability 0.5), then we give him a lottery ticket that pays $5000 with probability 0.50 and pays $0 otherwise; but if the coin comes up Tails, then we give him a lottery ticket that pays $5000 with probability 0.27 and pays $0 otherwise. He already told us that he would be indifferent between $2000 for sure and the first of these lottery tickets and he would be indifferent between $1000 for sure and the second of these lottery tickets. So he should be just indifferent between Gamble 1 and this two-stage gamble. But in this two-stage gamble, his ultimate prize is either $5000 or $0, and his probability of getting $5000 is

$$0.5*0.50 + 0.5*0.27 = 0.385$$

(This calculation could be done with a probability tree, as in Section 1.6.) Gamble 2 also pays either $5000 or $0, and its probability of paying $5000 is 0.4, which is greater than 0.385. So the decision maker should rationally prefer Gamble 2 over the two-stage gamble, which he should think is just as good as Gamble 1. Thus, he should prefer Gamble 2 over Gamble 1 if he is logically consistent. Figure 3.4 shows a spreadsheet illustrating this argument.

This argument can be generalized to show that logically consistent decision makers should always satisfy utility theory. To formulate this general argument, we first need to introduce some basic notation. Let us consider a decision maker who may have to choose among a variety of possible

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | The derivation of Utility Theory begins by choosing High and Low Prizes. | | | | | | | |
| 2 | 5000 | High Prize$ | | 0 | Low Prize$ | | | |
| 3 | For other prizes $X in between these prizes, we ask the Decision Maker: | | | | | | | |
| 4 | What probability p would make $X as good as a lottery that pays | | | | | | | |
| 5 | the High Prize with probability p, the Low Prize with probability 1-p? | | | | | | | |
| 6 | We can then make a table of the DM's responses. | | | | | | | |
| 7 | | | | | | | | |
| 8 | Given | DM assesses | | So the DM considers the Prize$ as good as | | | | |
| 9 | Prize$ | p=P(High$): | | this equivalent High-Low lottery: | | | | |
| 10 | 0 | 0 | | 0 | | | | |
| 11 | 920 | 0.25 | | 0 | | | | |
| 12 | 1000 | 0.27 | | 0 | | | | |
| 13 | 2000 | 0.5 | | 5000 | | | | |
| 14 | 3300 | 0.75 | | 5000 | | | | |
| 15 | 5000 | 1 | (rand) | 5000 | | | | |
| 16 | | | 0.321601 | | | | | |
| 17 | Now consider any other lotteries. | | | | | | | |
| 18 | | | | | | | | |
| 19 | LOTTERY 1 | | | | | | | |
| 20 | Prize$ | Prob'y | | P(Get High$ in H-L-equiv) | | | | |
| 21 | 1000 | 0.5 | | 0.27 | | | | |
| 22 | 2000 | 0.5 | | 0.5 | | | | |
| 23 | Simulations of Lottery 1 and an equivalent High-Low lottery: | | | | | | | |
| 24 | Lottery 1 | | H-L-equiv lottery | | P(Get High$ in H-L-equiv) | | | |
| 25 | 2000 | | | 5000 | 0.385 | | | |
| 26 | | | | | | | | |
| 27 | LOTTERY 2 | | | | | | | |
| 28 | Prize$ | Prob'y | | P(Get High$ in H-L-equiv) | | | | |
| 29 | 0 | 0.6 | | 0 | | | | |
| 30 | 5000 | 0.4 | | 1 | | | | |
| 31 | Simulations of Lottery 2 and an equivalent High-Low lottery: | | | | | | | |
| 32 | Lottery 2 | | H-L-equiv lottery | | P(Get High$ in H-L-equiv) | | | |
| 33 | 0 | | | 0 | 0.4 | | | |
| 34 | | | | | | | | |
| 35 | The decision-maker, if rational, should prefer Lottery 1 over Lottery 2 | | | | | | | |
| 36 | if the value of cell E28 is greater than the value of cell E37. | | | | | | | |
| 37 | Now let the DM's assessed equivalent probabilities of the high prize | | | | | | | |
| 38 | in cells B10:B15 be called "Utilities" of the prizes in A12:A17. | | | | | | | |
| 39 | Then the decision maker, if rational should always prefer the lottery | | | | | | | |
| 40 | that has the higher Expected Utility (in cell E25 or cell E33) | | | | | | | |
| 41 | FORMULAS FROM RANGE A1:H40 | | | | | | | |
| 42 | C15. =RAND() | | | | | | | |
| 43 | D10. =IF($C$15<B10,$A$2,$D$2) | | | | D10 copied to D10:D15 | | | |
| 44 | D21. =VLOOKUP(A21,$A$10:$B$15,2,0) | | | | D21 copied to D21:D22,D29:D30 | | | |
| 45 | A25. =DISCRINV(RAND(),A21:A22,B21:B22) | | | | | | | |
| 46 | D25. =VLOOKUP(A25,$A$10:$D$15,4,0) | | | | | | | |
| 47 | E25. =SUMPRODUCT(B21:B22,D21:D22) | | | | A25:E25 copied to A33:E33 | | | |

FIGURE **3.4**
**A spreadsheet illustrating the derivation of utility theory**

monetary lotteries or gambles. A gamble may be denoted here by a random variable (in boldface) that represents the unknown monetary value that the decision maker would be paid from this gamble if it was accepted. Given such a decision maker and given any two gambles **X** and **Y**, we may write

$$\mathbf{X} \succ \mathbf{Y}$$

to denote the proposition that "the decision maker would strictly prefer the gamble **X** over the gamble **Y**." Similarly, we may write

$$\mathbf{X} \sim \mathbf{Y}$$

Myerson_C03.qxd 09-09-2004 10:32 Page 99

to denote the proposition that "the decision maker would be indifferent between the gambles **X** and **Y**."

Utility theory says that each decision maker's preferences among all possible monetary gambles can be represented by some utility function $U(\bullet)$ such that, for any gambles **X** and **Y**,

$$\mathbf{X} \succ \mathbf{Y} \quad \text{when } E(U(\mathbf{X})) > E(U(\mathbf{Y}))$$

and

$$\mathbf{X} \sim \mathbf{Y} \quad \text{when } E(U(\mathbf{X})) = E(U(\mathbf{Y}))$$

Thus, according to utility theory, once we know a decision maker's utility function, then we can predict the decision maker's preferences over all possible monetary gambles. According to utility theory, then, different people may have different preferences for taking and avoiding risks because they have different utility functions.

To keep the argument simple, let us only consider lotteries where the outcome will be selected from some finite set of monetary prizes, which we may denote by

$$\{W_1, W_2, \ldots, W_n\}$$

Let $W_1$ denote the best of these possible outcomes, and let $W_n$ denote the worst of these possible outcomes. Now, for any other possible prize $W_i$ in this set, let us ask the decision maker:

> If you were comparing the alternatives of (1) getting $W_i$ dollars for sure, and (2) a binary lottery in which you will get either the best prize $W_1$ or the worst prize $W_n$, then what probability of getting the best prize $W_1$ in this binary lottery would make you just indifferent between these two alternatives.

Obviously, if the probability of the best prize were 1, then the binary lottery would be better; but if the probability of the best prize were 0, then the sure $W_i$ would be better. So there should exist some probability of getting the best prize such that the decision maker would be just indifferent between these two alternatives.

Given the decision maker's answer to this question, let $\mathbf{Z_i}$ denote such a random variable that will be equal to either $W_1$ or $W_n$ with the probability distribution that the decision maker considers just as good as $W_i$ for sure. That is, let $\mathbf{Z_i}$ be such that

$$P(\mathbf{Z_i} = W_1) + P(\mathbf{Z_i} = W_n) = 1 \quad \text{and} \quad \mathbf{Z_i} \sim W_i$$

Notice that we must have

$$P(\mathbf{Z_n} = W_1) = 0$$

because the worst prize, $W_n$, would be worse than any lottery that gave any positive probability of the best prize. Similarly, we must have

$$P(\mathbf{Z_1} = W_1) = 1$$

Now consider any two general lotteries **X** and **Y** that have outcomes drawn from this finite set of possible prizes $\{W_1, W_2, \ldots, W_n\}$. Since the decision maker is indifferent between each $W_i$ and the corresponding lottery $\mathbf{Z_i}$, he should not care if we substitute a ticket to the lottery $\mathbf{Z_i}$ wherever he would have gotten the prize $W_i$. Repeating this substitution argument for each $i$ from 1 to $n$, we conclude that the lottery **X** should be just as good, for this decision maker, as a two-stage lottery in which the first stage gives a ticket to a lottery $\mathbf{Z_1}$ or $\mathbf{Z_2}$ or . . . or $\mathbf{Z_n}$, respectively, with the probabilities $P(\mathbf{X} = W_1)$, $P(\mathbf{X} = W_2)$, . . ., $P(\mathbf{X} = W_n)$. In this two-stage lottery, the final outcome will be either the best prize $W_1$ or the worst prize $W_n$, and the probability of getting the best prize $W_1$ is

$$P(\mathbf{X} = W_1)*P(\mathbf{Z_1} = W_1) + P(\mathbf{X} = W_2)*P(\mathbf{Z_2} = W_1) + P(\mathbf{X} = W_n)*P(\mathbf{Z_n} = W_1)$$

So the lottery **X** should be just as good, for this decision maker, as a lottery that pays either the best prize $W_1$ or the worst prize $W_n$, where the probability of the best prize is

$$\sum_{i=1}^{n} P(\mathbf{X} = W_i)*P(\mathbf{Z_i} = W_1)$$

By a similar argument, the decision maker should be indifferent lottery **Y** should be just as good, for this decision maker, as a lottery that pays either the best prize $W_1$ or the worst prize $W_n$, where the probability of the best prize is

$$\sum_{i=1}^{n} P(\mathbf{Y} = W_i)*P(\mathbf{Z_i} = W_1)$$

Between these two binary lotteries where the only possible outcomes are $W_1$ or $W_n$, the better one is obviously the one with the higher probability of the best prize $W_1$. So the decision maker should think that **X** is better than **Y** if

$$\sum_{i=1}^{n} P(\mathbf{X} = W_i)*P(\mathbf{Z_i} = W_1) > \sum_{i=1}^{n} P(\mathbf{Y} = W_i)*P(\mathbf{Z_i} = W_1)$$

The trick now is to let $U(W_i) = P(\mathbf{Z_i} = W_1)$. Then we have just shown that

$$\mathbf{X} \succ \mathbf{Y} \quad \text{whenever } \sum_{i=1}^{n} P(\mathbf{X} = W_i)*U(W_i) > \sum_{i=1}^{n} P(\mathbf{Y} = W_i)*U(W_i)$$

That is, the better gamble should always be the one with the higher expected utility. Thus, we have found a way of assigning utility values to prizes such that the decision maker, if he is logically consistent, should always prefer the lottery with the higher expected utility.

You may notice that the utility numbers $U(X_i)$ that are generated by this argument are all between 0 and 1. However, there can also be valid utility functions in which the utility numbers are negative, or bigger than 1. In fact, we can multiply a utility function by any positive constant and add any other constant and we will still have have an equally valid utility function:

**Fact 1** Suppose that the decision maker's preferences over all monetary gambles can be represented by a utility function $U(\bullet)$. Suppose that another function

Myerson_C03.qxd 09-09-2004 10:32 Page 101

$V(\bullet)$ is an increasing linear transformation of $U(\bullet)$. That is, suppose that there exist two numbers $A$ and $B$ such that $A > 0$ and

$$V(x) = A*U(x) + B \quad \text{for all } x$$

Then $V(\bullet)$ is also a valid utility function for representing this decision maker's preferences over all monetary gambles.

**Proof of Fact 1** For any two gambles $\mathbf{X}$ and $\mathbf{Y}$, if $E(U(\mathbf{X})) > E(U(\mathbf{Y}))$, then

$$E(V(\mathbf{X})) = E(A*U(\mathbf{X}) + B) = A*U(\mathbf{X}) + B$$
$$> A*U(\mathbf{Y}) + B = E(A*U(\mathbf{Y}) + B) = E(V(\mathbf{Y}))$$

So comparing expectations of $V$ and $U$ both yield the same preference ordering over gambles.

## 3.5 Advanced Technical Note on Constant Risk Tolerance

The certainty equivalent of a gamble is the amount of money such that the decision maker would be indifferent between the gamble and this amount of money for sure. So an amount of money $M$ is the certainty equivalent of a gamble $\mathbf{X}$ when

$$M \sim \mathbf{X} \quad \text{and} \quad U(M) = E(U(\mathbf{X}))$$

We may let $CE(\mathbf{X})$ denote the decision maker's certainty equivalent for the gamble $\mathbf{X}$, so

$$U(CE(\mathbf{X})) = E(U(\mathbf{X}))$$

We say that a decision maker is *risk averse* if the decision maker's certainty equivalent for any risky gamble is always less than its expected monetary value, that is,

$$CE(\mathbf{X}) < E(\mathbf{X})$$

whenever the standard deviation of $\mathbf{X}$ is bigger than zero. The *risk premium* is the difference

$$RP(\mathbf{X}) = E(\mathbf{X}) - CE(\mathbf{X})$$

The assumption of *constant risk tolerance* says that a decision maker's preferences over two gambles would not be affected by adding a constant monetary amount to both of them. That is, for any two gambles $\mathbf{X}$ and $\mathbf{Y}$ and for any nonrandom monetary amount $k$,

$$\text{if } \mathbf{X} \succ \mathbf{Y} \text{ then } \mathbf{X} + k \succ \mathbf{Y} + k$$

Equivalently, constant risk tolerance means that adding a constant monetary amount to any gamble would increase its certainty equivalent by this

Myerson_C03.qxd 09-09-2004 10:32 Page 102

amount. That is, if the decision maker has constant risk tolerance, then for any gamble **X** and for any nonrandom monetary amount $k$,

$$CE(\mathbf{X} + k) = CE(\mathbf{X}) + k$$

Notice that constant risk tolerance does not imply that $CE(k*\mathbf{X}) = k*CE(\mathbf{X})$. This condition is called *proportional risk tolerance*. A decision maker cannot have both constant risk tolerance and proportional risk tolerance unless he is risk neutral (in which case he has the simple utility function $U(x) = x$ and always prefers the gamble with the highest expected monetary value).

Our two following main results tell us that constant risk tolerance holds true for a risk-averse decision maker if and only if the decision maker has a utility function satisfying a simple exponential formula.

**Fact 2** Suppose that a decision maker's risk preferences can be represented by a utility function of the form

$$U(x) = -\text{EXP}(-x/r)$$

for some constant $r$. Then the decision maker has constant risk tolerence. Furthermore, for any two gambles **X** and **Y** with monetary payoffs that are independent random variables, we have

$$CE(\mathbf{X} + \mathbf{Y}) = CE(\mathbf{X}) + CE(\mathbf{Y})$$

That is, this exponential utility formula implies that the decision maker's certainty equivalent for the sum of two independent gambles is just the sum of his certainty equivalents for these two gambles separately.

**Fact 3** If a risk-averse decision maker has risk preferences that satisfy utility theory and constant risk tolerance, then the decision maker's preferences can be represented by a utility function of the form

$$U(x) = -\text{EXP}(-x/r)$$

for some constant $r$.

**Proof of Fact 2** We use here some well-known mathematical properties of the EXP function. (See also Section 4.2.) For any two numbers $x$ and $y$,

$$\text{EXP}(x + y) = \text{EXP}(x)*\text{EXP}(y) \quad \text{and}$$
$$\text{EXP}(x*y) = (\text{EXP}(x))\text{^}y$$

Furthermore, if $x > y$, then $\text{EXP}(x) > \text{EXP}(y) > 0$.

Let $M = CE(\mathbf{X})$. With the exponential utility formula, we have

$$-\text{EXP}(-M/r) = U(M) = E(U(\mathbf{X})) = \sum_x P(\mathbf{X} = x)*(-\text{EXP}(-x/r))$$

To prove that this exponential formula implies constant risk tolerance, notice that for any constant $k$,

$$\begin{aligned} U(M+k) &= -\text{EXP}(-(M+k)/r) \\ &= -\text{EXP}(-M/r)*\text{EXP}(-k/r) \\ &= -(\sum_x P(\mathbf{X}=x)*\text{EXP}(-x/r))*\text{EXP}(-k/r) \\ &= -(\sum_x P(\mathbf{X}=x)*\text{EXP}(-(x+k)/r)) \\ &= E(U(\mathbf{X}+k)) \end{aligned}$$

and so $M + k = CE(\mathbf{X} + k)$.

To prove the second property about CEs of independent sums, suppose that $\mathbf{X}$ and $\mathbf{Y}$ are independent, let $M = CE(\mathbf{X})$, and let $N = CE(\mathbf{Y})$. Then we have

$$\begin{aligned} U(M+N) &= -\text{EXP}(-(M+N)/r) \\ &= -\text{EXP}(-M/r)*\text{EXP}(-N/r) \\ &= -(\sum_x P(\mathbf{X}=x)*\text{EXP}(-x/r))*(\sum_y P(\mathbf{Y}=y)*\text{EXP}(-y/r)) \\ &= -\sum_x \sum_y P(\mathbf{X}=x)*P(\mathbf{Y}=y)*\text{EXP}(-(x+y)/r) \\ &= E(U(\mathbf{X}+\mathbf{Y})) \end{aligned}$$

and so $M + N = CE(\mathbf{X} + \mathbf{Y})$.

**Proof of Fact 3** We will prove the result for all integer payoffs.

First, ask the decision maker "for what number $p$ would you be indifferent between getting \$1 for sure or getting a gamble that pays either \$2 or \$0 with probabilities $p$ and $1 - p$, respectively?" Assuming risk aversion, he will specify a number $p$ that is between 1/2 and 1. (For example, he might say $p = 0.51$.)

The certainty equivalent that he has expressed tells us that the decision maker's utility function $U$ must satisfy the equation

$$p*U(2) + (1-p)*U(0) = U(1)$$

Constant risk tolerance tells us that adding any number $x$ to all the prizes, we must have

(1) $p*U(x+2) + (1-\text{p})*U(x) = U(x+1)$, for every $x$

Now let $r = -1/\text{LN}((1-p)/p)$, which is the number such that

$$\text{EXP}(-1/r) = (1-p)/p$$

Notice that $(1 - p)/p$ is between 0 and 1, because $p$ is between 0.5 and 1. With this value of $r$, the basic mathematical properties of the EXP function imply that

$$\text{EXP}(-x/r) = ((1-p)/p)\text{^}x$$

We can now show that equation (1) would also be satisfied when we replace the function $U$ by another function $V$ that is defined by the formula

$$V(x) = -\text{EXP}(-x/r) = 0 - ((1-p)/p)\text{^}x$$

(The 0 at the front of the formula is needed in Excel because, in its interpretation of formulas, Excel applies a negative sign before exponentiation,

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | DERIVATION OF EXPONENTIAL UTILITY FROM CONSTANT RISK TOLERANCE | | | | | | | |
| 2 | Suppose $1 is as good as the lottery (p[$2],(1-p)[$0]). | | | | | | | |
| 3 | So p*U(2)+(1-p)*U(0)=U(1). | | | | | | COMPARE: | |
| 4 | Constant risk tolerance implies that, for any k | | | | | | RiskTol | A |
| 5 | we must also have p*U(k+2)+(1-p)*U(k)=U(k+1). | | | | | | 24.99667 | 25.5 |
| 6 | So for each k: U(k+2) = (U(k+1)-(1-p)*U(k))/p. | | | | | | | B |
| 7 | | | | | | | | 25.5 |
| 8 | Initial parameters | | | Resulting utilities | | | | |
| 9 | p | 0.51 | | k | U(k) | | UTIL | A*UTIL+B |
| 10 | U(0) | 0 | | 0 | 0 | | -1 | 0 |
| 11 | U(1) | 1 | | 1 | 1 | | -0.96078 | 1 |
| 12 | | | | 2 | 1.960784 | | -0.92311 | 1.960784 |
| 13 | | | | 3 | 2.883891 | | -0.88691 | 2.883891 |
| 14 | FORMULAS | | | 4 | 3.770797 | | -0.85213 | 3.770797 |
| 15 | G7. =1/LN(B9/(1-B9)) | | | 5 | 4.622923 | | -0.81871 | 4.622923 |
| 16 | E10. =B10 | | | 6 | 5.441632 | | -0.7866 | 5.441632 |
| 17 | E11. =B11 | | | 7 | 6.228234 | | -0.75576 | 6.228234 |
| 18 | E12. =(E11-(1-$B$9)*E10)/$B$9 | | | 8 | 6.98399 | | -0.72612 | 6.98399 |
| 19 | E12 copied down | | | 9 | 7.710108 | | -0.69764 | 7.710108 |
| 20 | G10. =-EXP(-D10/$H$9) | | | 10 | 8.407751 | | -0.67028 | 8.407751 |
| 21 | G10 copied down | | | 11 | 9.078035 | | -0.644 | 9.078035 |
| 22 | H5. =(B11-B10)/(G11-G10) | | | 12 | 9.722034 | | -0.61874 | 9.722034 |
| 23 | H7. =E10-H5*G10 | | | 13 | 10.34078 | | -0.59448 | 10.34078 |
| 24 | H10. =$H$5*G10+$H$7 | | | 14 | 10.93526 | | -0.57117 | 10.93526 |
| 25 | H10 copied down | | | 15 | 11.50642 | | -0.54877 | 11.50642 |
| 26 | | | | 16 | 12.05519 | | -0.52725 | 12.05519 |
| 27 | | | | 17 | 12.58244 | | -0.50657 | 12.58244 |
| 28 | | | | 18 | 13.08901 | | -0.48671 | 13.08901 |
| 29 | | | | 19 | 13.57571 | | -0.46762 | 13.57571 |
| 30 | | | | 20 | 14.04333 | | -0.44928 | 14.04333 |

FIGURE **3.5**
**Derivation of exponential utility from constant risk tolerance**

but Excel does exponentiation before subtraction.) To verify that this formula yields a solution to (1), notice that it gives us

$$p*V(x+2) + (1-p)*V(x) = 0 - p*((1-p)/p)\text{^}(x+2) - (1-p)*((1-p)/p)\text{^}x$$
$$= 0 - (p*((1-p)/\text{p})\text{^}2 + (1-\text{p}))*((1-p)/p)\text{^}x$$
$$= 0 - ((1-p)/p)*((1-p)/p)\text{^}x$$
$$= 0 - ((1-p)/p)\text{^}(x+1) = V(x+1)$$

But equation (1) is a linear difference equation, and any two solutions to such equations must differ by a linear transformation. Thus, the utility function $U(x)$ must be a linear transformation of the exponential function $V(x) = -\text{EXP}(-x/r)$. That is, there must exist constants $A$ and $B$ such that, for every $x$,

$$U(x) = A*V(x) + B$$

In fact, the constants are

$$A = (U(1) - U(0))/(V(1) - V(0))$$
$$B = U(0) - A*V(0)$$

Because $U(x)$ and $V(x)$ are both increasing in $x$, we know that $A$ must be a positive number. So by Fact 1, the decision maker's preferences over gambles can be equally well represented by the utility function $V(x) = -\text{EXP}(-x/r)$.

Myerson_C03.qxd 09-09-2004 10:32 Page 105

## 3.6 Summary

A decision maker's certainty equivalent of a gamble can be defined as the amount of money such that the decision maker would be just indifferent between the gamble and this amount of money for sure. Under the assumptions of risk neutrality from Chapter 2, the certainty equivalent of any gamble was identified with the gamble's expected monetary value. We saw in this chapter that when the expected monetary value seems seriously questionable as a criterion for optimal decision making, utility theory can provide a more sophisticated decision criterion that takes account of the decision maker's personal attitudes toward risks. According to utility theory, any decision maker has a personal utility function $U$ that summarizes the willingness to take monetary risks. When we know the decision maker's personal utility function $U$, the decision maker's certainty equivalent for any gamble that pays some random amount of money $\mathbf{X}$ can be computed by the formula $U(CE) = EU(\mathbf{X})$.

To simplify the problem of assessing a decision maker's utility function, we introduced the assumption of constant risk tolerance, which is satisfied by a one-parameter family of utility functions. Under this assumption, we studied ways to measure a decision maker's personal risk-tolerance parameter by asking his certainty equivalent for a simple gamble with two equally likely prizes. Then, given such a measurement of the decision maker's risk tolerance, we saw how to compute the decision maker's certainty equivalent for any gamble. To illustrate the basic ideas in these certainty-equivalent calculations, we first considered simple gambles that have a known discrete probability distribution. Then, to compute certainty equivalents for more complex gambles, we saw how to estimate certainty equivalents from simulation data.

Simtools functions introduced in this chapter include UTIL, UINV, RISKTOL, CEPR, and CE. Excel's GoalSeek tool was also mentioned as an alternative to using RISKTOL.

## Exercises

**1.** A decision maker subjectively assessed the following certainty equivalents:

**(i)** "A lottery that pays either \$10,000 or \$0, each with probability 1/2, would be worth \$3800 (for sure) to me."

**(ii)** "A lottery that pays either \$5000 or \$0, each with probability 1/2, would be worth \$2200 to me."

**(iii)** "A lottery that pays either \$1000 or \$0, each with probability 1/2, would be worth \$440 to me."

**(a)** If the decision maker has constant risk tolerance, then what risk tolerance would be implied by each of these three statements?

**(b)** One of these statements implies a risk tolerance that is substantially different from the other

Myerson_C03.qxd 09-09-2004 10:32 Page 106

two. If the decision maker reassessed the certainty equivalent for the lottery described in this statement, what certainty equivalent would correspond to the same level of risk tolerance expressed in the other two statements?

Now, assume that the decision maker has constant risk tolerance equal to the average of the two closer values that you found in part (a). Then what should be the decision maker's certainty equivalent for each of the following lotteries?

**(c)** A lottery that pays either \$2000 or \$6000 or \$10,000, each with probability 1/3

**(d)** A lottery that pays either \$2000 or \$6000 or \$10,000, with probabilities 0.4, 0.2, and 0.4, respectively

**(e)** A lottery that pays either \$0 or \$4000 or \$8000, with probabilities 0.4, 0.2, and 0.4, respectively

**(f)** A lottery that pays either \$0 or \$8000 or \$16,000, with probabilities 0.4, 0.2, and 0.4, respectively

**2.** A decision maker has to choose among the following three gambles:

Lottery 1 pays either \$20,000 or \$30,000 or \$70,000, with probabilities 0.260, 0.675, and 0.065, respectively.

Lottery 2 pays either –\$10,000 or \$30,000 or \$40,000, with probabilities 0.065, 0.675, and 0.260, respectively.

Lottery 3 pays a dollar amount drawn from a Normal distribution with the same expected monetary value and the same standard deviation as Lottery 1.

Suppose that the decision maker has a risk-tolerance constant equal to \$20,000.

**(a)** Compute the certainty equivalent of Lotteries 1 and 2 from their discrete probability distributions.

**(b)** Estimate the certainty equivalent of each of these three lotteries from simulation data, using data from at least 500 simulations.

**(c)** Based on the simulation data that you used in part (b), compute a 95% confidence interval for the certainty equivalent of each of these three lotteries. Then check to see whether the intervals for Lotteries 1 and 2 include the actual values that you computed in part (a).

**3.** Consider again Exercise 4 at the end of Chapter 2, where the Acme Widget Company's profit from selling widgets next year is uncertain. Recall that the profit uncertainty is derived from uncertainty about several different factors: the total demand for widgets next year, Acme's share of this demand, the selling price of widgets, Acme's variable cost per widget sold, and Acme's fixed cost of selling any widgets next year.

Let us assume here that Acme's chief executive officer wants Acme to evaluate risks according to a utility function with constant risk tolerance. To assess this risk tolerance, the chief executive officer was asked to consider a simple hypothetical lottery that would pay either \$500,000 or –\$100,000, each with probability 1/2. The CEO said that if Acme owned a "ticket" to such a lottery, then it should certainly not give it away for \$0, but it should be glad to give it up in exchange for a sure \$100,000.

**(a)** Based on the chief executive officer's attitudes toward the hypothetical simple lottery, find a lower bound and an upper bound on Acme's risk tolerance.

**(b)** Make a table of at least 1000 simulations of Acme's profit from selling widgets next year, using the discrete probability distributions for each factor that are described in Exercise 4 of Chapter 2. Use this simulation data to answer parts (c) and (d) of this exercise.

**(c)** Make a table showing, for each number $X$ that is a multiple of \$10,000 between \$0 and \$100,000, what would be Acme's *risk tolerance* and Acme's *certainty equivalent* of its profit from selling widgets next year if Acme's certainty equivalent of the simple hypothetical lottery were $X$.

**(d)** Estimate the risk tolerance that would make Acme's certainty equivalent of its profit from selling widgets next year equal to \$100,000. With this risk tolerance, what would be Acme's certainty equivalent of the simple hypothetical lottery? (*Hint:* You may use Excel's GoalSeek tool here.)

Myerson_C04.qxd 09-09-2004 11:55 Page 107

# 4

# Continuous Random Variables

In this chapter, we study ways to model an unknown quantity that has infinitely many possible values. To describe our beliefs about such an unknown quantity, a probability distribution must specify probabilities for infinitely many possible events, which may seem a rather daunting task. But with some technical assumptions we can make this task quite tractable.

The kind of situation that we want to model by these techniques is illustrated by Part C of the Superior Semiconductor case, where the decision analyst is confronted by a broader scope of uncertainty. In this example, we learn that the cost of developing a new product, which had been roughly estimated (in Part A) at $26 million, might be significantly more or less than that amount. When the full scope of uncertainty about this dollar cost is recognized, the decision analyst must admit that the actual development cost could be virtually any positive number.

The way to make such a probability assessment task tractable is to assume that our beliefs about the unknown quantity can be reasonably approximated by some probability distribution in one of the mathematical families of probability distributions that statisticians have studied. Each of these families includes a variety of different probability distributions that are characterized by a few parameters. So our task is first to pick an appropriate family of probability distributions, and then to assess values for the corresponding parameters that can fit our actual beliefs about the unknown quantity. The most commonly used families of probability distributions typically have only two or three parameters to assess. So this parameter assessment problem may actually be easier than the problem of assessing a discrete probability distribution that we saw in Part A of the Superior Semiconductor case, where the decision analyst had to generate five probability numbers that sum to 1.

This chapter considers some of the best-known families of probability distributions, each of which is considered particularly appropriate for certain special kinds of unknown quantities. The Normal and Lognormal

Myerson_C04.qxd 09-09-2004 11:55 Page 108

## CASE ■ Superior Semiconductor (Part C)

When Eastmann presented to the executive committee her analysis of the proposed new T-regulator product, Suttcliff expressed his appreciation of the frank and clear way that she had acknowledged the difficulty of predicting competitors' reactions and Superior's resulting share of the new market. But other people on the committee questioned whether there might be other important risk factors that Eastmann had ignored. In the ensuing discussion, Eastmann discovered that she might have misinterpreted the confidence with which the business marketing director and the chief production engineer had expressed their estimates of market size and development costs. They felt confident that their numbers were the best working estimates, but they did not mean to imply that the actual outcomes might not be quite different. If uncertainty was to be expressed, then they had a lot more to say.

The chief production engineer said that total development cost for the new product was equally likely to be above or below his original cost estimate of \$26 million. After further questioning by Eastmann, he also suggested that there was a 25% probability of costs being below \$23 million, and a 25% probability of costs being above \$29 million. He then added that, even after all these development expenses, there was some chance that Superior might find that the new device was unmarketable because of some unforeseen technical problem. When Eastmann asked him to quantify the probability of such a technical failure, he said it had a small 5% probability.

Then the business marketing manager was similarly asked to assess quartiles for the potential value of the T-regulator market. She said that if the product was successfully developed, then the total discounted value of the market was equally likely to be above or below \$100 million, with a 25% probability of being below \$80 million, and a 25% probability of being above \$125 million.

Eastmann promised to do some further analysis, to re-estimate the expected net profit and the cumulative risk profile of the proposed T-regulator product, taking account of the uncertainties that had been expressed in this meeting.

distributions are widely recognized as applicable in the broadest range of situations, and so we will give most careful consideration here to these two families. In addition, we will introduce the Generalized-Lognormal family, which includes the Normal and Lognormal families as special cases and which is particularly convenient for subjective assessments in cases where hard data are scarce. Other families of probability distributions that are used in various specialized areas of application are discussed briefly at the end of the chapter.

## 4.1 Normal Distributions

The Central Limit Theorem (which was introduced in Section 2.6) tells us that if we make a large number of independent samples from virtually any probability distribution, then our sample average will have a probability distribution that is approximately Normal. This remarkable theorem

Myerson_C04.qxd 09-09-2004 11:55 Page 109

assures us that there will be many unknowns in the real world that are well described by Normal distributions. In fact, a Normal probability distribution is generally a good fit for any unknown quantity that can be expressed as the sum of many small contributions that are independently drawn from some fixed probability distribution. For example, if our uncertainty about each day's revenue is the same as any other day and is independent of all other days' revenues, and each day's revenue will be small relative to the whole year, then total annual revenue should be a Normal random variable. (Here the phrase "Normal random variable" is used to denote any unknown quantity that has a Normal probability distribution.)

The Normal probability distributions are a two-parameter family that are indexed by the mean and standard deviation. That is, for any number $\mu$ and any positive number $\sigma$, there is a Normal probability distribution that has mean $\mu$ and standard deviation $\sigma$.

A linear transformation of any Normal random variable is also Normal. That is, if a random variable $\mathbf{X}$ has a Normal probability distribution with mean $\mu$ and standard deviation $\sigma$, and if $c$ and $d$ are any two nonrandom numbers, then the random variable $c*\mathbf{X} + d$ also has a Normal probability distribution, with mean $c*\mu + d$ and standard deviation $|c|*\sigma$.

The range of possible values of a Normal random variable includes all numbers, positive or negative. Of course, there are infinitely many numbers, and so no single number gets positive probability under a Normal distribution. Instead, we can only talk meaningfully about intervals of numbers having positive probability. In the language of probability theory, such distributions that do not put positive probability on any single value are called *continuous*.

Excel provides several functions for dealing with Normal distributions. Here we will use two: NORMINV and NORMSDIST. Of these, the function NORMINV will be much more important in this book, but both will be explained now.

The cumulative probability below any value in a Normal distribution depends on how many standard deviations this value is above or below the mean, according to the Excel function NORMSDIST. That is, if $\mathbf{X}$ is any Normal random variable with mean $\mu$ and standard deviation $\sigma$, and $w$ is any number, then

$$P(\mathbf{X} < w) = \text{NORMSDIST}((w - \mu)/\sigma)$$

If we want to know the probability that $\mathbf{X}$, a Normal random variable with mean $\mu$ and standard deviation $\sigma$, is between some pair of numbers $w$ and $y$, where $w < y$, then we can use the formula

$$\begin{aligned} P(w < \mathbf{X} < y) &= P(\mathbf{X} < y) - P(\mathbf{X} < w) \\ &= \text{NORMSDIST}((y - \mu)/\sigma) - \text{NORMSDIST}((w - \mu)/\sigma) \end{aligned}$$

as illustrated by cell F3 in Figure 4.1.

Notice that there are two S's in the name of the function NORMSDIST. The first S actually stands for "standard," because a Normal random

Myerson_C04.qxd 09-09-2004 11:55 Page 110

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | Mean | Stdev | Normal random variable | | | 19.9 | 20.1 | |
| 2 | 26 | 4.5 | | 29.6312 | | P(F1 < D2 < G1) | | |
| 3 | | | | | | 0.00729 | | |
| 4 | NORMAL DISTRIBUTIO | | | | FORMULAS | | | |
| 5 | CumvProby | Value | ProbyDensity | | D2. =NORMINV(RAND(),A2,B2) | | | |
| 6 | 0.001 | 12.0939 | | F3. =NORMSDIST((G1-A2)/B2)-NORMSDIST((F1-A2)/B2) | | | | |
| 7 | 0.01 | 15.5315 | 0.00407 | | B6. =NORMINV(A6,$A$2,$B$2) | | | |
| 8 | 0.02 | 16.7581 | 0.00998 | | B6 copied to B6:B106 | | | |
| 9 | 0.03 | 17.5364 | 0.01467 | | C7. =(A8-A6)/(B8-B6) | | | |
| 10 | 0.04 | 18.1219 | 0.01884 | | C7 copied to C7:C105 | | | |
| 11 | 0.05 | 18.5982 | 0.02269 | | A6 contains 0.001 | | | |
| 12 | 0.06 | 19.0035 | 0.02629 | | A7:A105 is filled by a series from .01 to .99, step .01 | | | |
| 13 | 0.07 | 19.3589 | 0.02969 | | A106 contains 0.999 | | | |
| 14 | 0.08 | 19.6772 | 0.03291 | | | | | |
| 15 | 0.09 | 19.9666 | 0.03598 | **Inverse Cumulative Chart plots (A6:A106,B6:B106)** | | | | |
| 16 | 0.1 | 20.233 | 0.03891 | | | | | |
| 17 | 0.11 | 20.4806 | 0.04171 | | | | | |
| 18 | 0.12 | 20.7126 | 0.04438 | | | | | |
| 19 | 0.13 | 20.9312 | 0.04695 | | | | | |
| 20 | 0.14 | 21.1386 | 0.04941 | | | | | |
| 21 | 0.15 | 21.3361 | 0.05176 | | | | | |
| 22 | 0.16 | 21.5249 | 0.05402 | | | | | |
| 23 | 0.17 | 21.7063 | 0.05619 | | | | | |
| 24 | 0.18 | 21.8809 | 0.05827 | | | | | |
| 25 | 0.19 | 22.0495 | 0.06027 | | | | | |
| 26 | 0.2 | 22.2127 | 0.06218 | | | | | |
| 27 | 0.21 | 22.3711 | 0.06401 | | | | | |
| 28 | 0.22 | 22.5251 | 0.06577 | | | | | |
| 29 | 0.23 | 22.6752 | 0.06745 | | | | | |
| 30 | 0.24 | 22.8216 | 0.06906 | | | | | |
| 31 | 0.25 | 22.9648 | 0.07059 | | | | | |
| 32 | 0.26 | 23.1049 | 0.07206 | **Probability Density Chart plots (B7:B105,C7:C105)** | | | | |
| 33 | 0.27 | 23.2423 | 0.07346 | | | | | |
| 34 | 0.28 | 23.3772 | 0.07479 | | | | | |
| 35 | 0.29 | 23.5098 | 0.07605 | | | | | |
| 36 | 0.3 | 23.6402 | 0.07725 | | | | | |
| 37 | 0.31 | 23.7687 | 0.07838 | | | | | |
| 38 | 0.32 | 23.8954 | 0.07946 | | | | | |
| 39 | 0.33 | 24.0204 | 0.08046 | | | | | |
| 40 | 0.34 | 24.1439 | 0.08141 | | | | | |
| 41 | 0.35 | 24.2661 | 0.0823 | | | | | |
| 42 | 0.36 | 24.3869 | 0.08313 | | | | | |
| 43 | 0.37 | 24.5067 | 0.08389 | | | | | |
| 44 | 0.38 | 24.6253 | 0.0846 | | | | | |
| 45 | 0.39 | 24.7431 | 0.08525 | | | | | |
| 46 | 0.4 | 24.8599 | 0.08584 | | | | | |
| 47 | 0.41 | 24.976 | 0.08638 | | | | | |

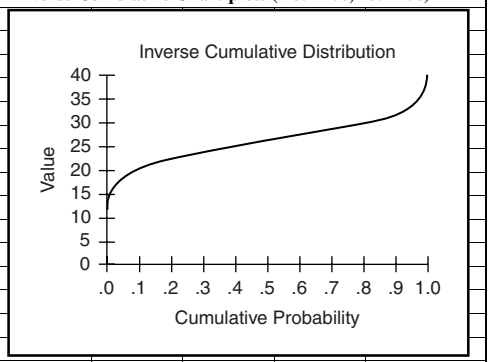


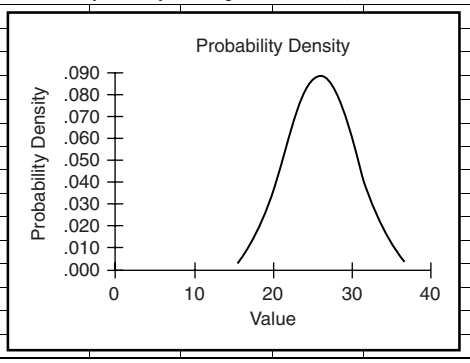


FIGURE **4.1**
**The Normal distribution**

variable with mean 0 and standard deviation 1 is sometimes called a *standard Normal* random variable. For any number $w$, the Excel function NORMSDIST($w$) returns the probability of such a standard Normal random variable being less than $w$.

The Normal distribution has a symmetry around its mean so that, for any number $x$,

$$\text{NORMSDIST}(-w) = 1 - \text{NORMSDIST}(w)$$

Myerson_C04.qxd 09-09-2004 11:55 Page 111

That is, a Normal random variable with mean μ and standard deviation σ is as likely to be less than $\mu - w*\sigma$ as it is to be greater than $\mu + w*\sigma$.

Some special values of the NORMSDIST function worth noting are NORMSDIST(0) = 0.5, NORMSDIST(0.675) = 0.75, NORMSDIST(1) = 0.841, NORMSDIST(1.96) = 0.975, NORMSDIST(3) = 0.99865. These cumulative probabilities imply that when **X** is any Normal random variable with mean μ and standard deviation σ,

$$P(\mathbf{X} < \mu) = P(\mathbf{X} > \mu) = 0.5$$
$$P(\mathbf{X} < \mu - 0.675*\sigma) = 0.25$$
$$P(\mathbf{X} < \mu + 0.675*\sigma) = 0.75$$
$$P(\mu - \sigma < \mathbf{X} < \mu + \sigma) = 0.683$$
$$P(\mu - 1.96*\sigma < \mathbf{X} < \mu + 1.96*\sigma) = 0.95$$
$$P(\mu - 3*\sigma < \mathbf{X} < \mu + 3*\sigma) = 0.9973$$

The Excel function NORMINV is the inverse cumulative probability function for Normal probability distributions. That is, if **X** is any Normal random variable with mean μ and standard deviation σ, then NORMINV($q$, μ, σ) is the value that has probability $q$ of being greater than **X**,

$$P(\mathbf{X} < \text{NORMINV}(q, \mu, \sigma)) = q$$

Here NORMINV's parameters $q$, μ, and σ can be any three numbers such that $\sigma > 0$ and $0 < q < 1$. Using NORMINV, we can make an inverse cumulative distribution chart by plotting probability numbers from 0 to 1 against the corresponding NORMINV(probability, μ, σ) values, as shown in Figure 4.1 (in rows 16–30) for the parameters $\mu = 26$ and $\sigma = 4.5$. (We will see that such a Normal distribution, with mean 26 and standard deviation 4.5, may be a reasonable model of the beliefs expressed about the development cost in Part C of the Superior Semiconductor case.)

Recall from Chapter 2 that we can simulate a random variable with any probability distribution by applying a RAND( ) value as input into the inverse cumulative probability function. Thus, we can simulate a Normal random variable with mean μ and standard deviation σ by the formula

NORMINV(RAND( ), μ, σ)

as shown in cell D2 of Figure 4.1 (where the mean μ is in cell A2 and the standard deviation σ is in cell B2).

We have tried to argue that the inverse cumulative probability curve shows us everything that we could want to know about the probability distribution of a random variable. But people often prefer to look at a different way of representing probability distributions graphically, called the *probability density* curve. We will not need to use probability density functions for any calculations in this book, but, admittedly, they are handy for drawing intuitive pictures of probability distributions. The probability density function provides a useful way to visualize a random variable's probability

Myerson_C04.qxd 09-09-2004 11:55 Page 112

distribution because for any interval of possible values, the probability of the interval is the area under the probability density curve over this interval. So in case you want to look at these pictures too, how to make them and what they mean will be explained. Excel provides a rather complicated formula for computing Normal probability density functions [NORMDIST(•, μ, σ, FALSE)]. But to emphasize the importance of the inverse cumulative function NORMINV, and to help you better understand what probability densities are, we will instead show you here how to estimate such probability densities with NORMINV.

A probability density for a random variable at any possible value is actually defined mathematically to be the slope of the cumulative probability function at that value. So when a mathematician says that a random variable **X** has a probability density $f(y)$ at some value $y$, this statement means that for any small interval around $y$, the probability of **X** being in that interval is approximately equal to the length of the interval multiplied by the density $f(y)$. According to Figure 4.1, for example, a Normal random variable with mean 26 and standard deviation 4.5 has, at the value 20, a probability density of approximately 0.036. So we may infer that the probability of such a random variable being between 19.9 and 20.1 (an interval of length 0.2 that includes the value 20) is approximately 0.2∗0.036 = 0.0072. Notice that this result is very close to the probability that is computed with NORMSDIST in cell F3 of Figure 4.1.

To make a chart of probability densities for a given distribution, let us begin by filling a range of cells in column A with probability values that go from 0 to 1. For the spreadsheet of Figure 4.1, such probability values from 0 to 1 with step size 0.01 are filled into the range A6:A106, except that the values 0 and 1 at the top and bottom of this range are changed to 0.001 and 0.999 because NORMINV and many other inverse cumulative functions are not defined at the probabilities 0 and 1. Next, in column B, let us compute the corresponding inverse cumulative values. So for the spreadsheet in Figure 4.1, with the mean in cell B2 and the standard deviation in cell C2, we enter in cell B6 the formula

```
=NORMINV(A6,$A$2,$B$2)
```

and we copy B6 to cells B6:B106. Then the probability density can be estimated in column C by entering the formula

```
=(A8–A6)/(B8–B6)
```

into cell C7, and then copying C7 to cells C7:C105. (This formula in the C column is not applied in the top and bottom rows 6 and 106 because it needs to refer to A and B values in the rows above and below the current row.) In this C7 formula, the denominator B8 – B6 is the length of the interval from the value of cell B6 (12.09), which is slightly less than B7 (15.53), to the value of cell B8 (16.76), which is slightly greater than B7. The numerator A8 – A6 is the probability of the random variable being in

Myerson_C04.qxd 09-09-2004 11:55 Page 113

this interval, because A8 (0.02) is the probability of the random variable being less than B8, whereas A6 (0.001) is the probability of the random variable being less than B6. So (A8 – A6)/(B8 – B6) is the probability in a short interval around B7 divided by the length of this interval, and thus it can be used to approximate the probability density of this random variable at B7.

The probability density chart shown in Figure 4.1 (in rows 33 to 47) is plotted from the values in B7:B105 and the estimated densities in C7:C105. The characteristic shape of this Normal density chart is often called a *bell-shaped curve*. The probability density function provides a useful way to visualize a random variable's probability distribution because, for any interval of possible values, the probability of the interval is the area under the probability density curve over this interval.

Normal random variables can take positive or negative values, and their probability of going negative can be nonnegligible when the mean is less than twice the standard deviation. But there are many unknown quantities, such as prices, for which a negative value would be nonsensical. For such unknowns, it is more appropriate to use a probability distribution for which the range of possible values includes only the positive numbers, not the negative numbers. The statistics literature features several families of such distributions, and among these the one that is most commonly used in finance and decision analysis is the family of Lognormal probability distributions. But before introducing Lognormal random variables, we briefly turn aside to introduce two important mathematical functions: EXP and LN.

## 4.2 EXP and LN

The function EXP is called the *exponential function*. If you enter the formula =EXP(1) into any cell, you will get the magic number 2.71828 . . . , which mathematicians often denote by the symbol "e." For any other number *r*, the formula EXP(*r*) returns the value of e raised to the power *r*, that is EXP(*r*) equals 2.71828^*r*. (In Figure 4.2, see cells B6, D6, and E6.)

The function LN is called the *natural logarithm function*, and it is the inverse of EXP. That is, for any positive number *x*, the formula LN(*x*) returns the number *r* such that EXP(*r*) = *x*. (In Figure 4.2, look at cells B3, B6, and B7.)

You may have studied these functions EXP and LN in mathematics classes, but they are introduced here with a basic financial application that may help you to see what is so special about the powers of this magic number e. Suppose that you put money into a bank account that offers an annual interest rate of some nominal interest rate *r* that will be compounded

Myerson_C04.qxd 09-09-2004 11:55 Page 114

| | A | B | C | D | E | F |
|---|---|---|---|---|---|---|
| 1 | THE EXP FUNCTION FROM COMPOUND INTEREST: | | | | | |
| 2 | Suppose we start with $1 in a bank account. | | | | | |
| 3 | Nominal interest rate | 0.12 | | | | |
| 4 | Times compounded/year | 4 | *(try 1,2,4,12,365)* | | | |
| 5 | $Value at end of year | 1.1255088 | | The number e: | e^r | |
| 6 | EXPonential formula | 1.1274969 | | 2.7182818 | 1.1274969 | |
| 7 | LN, the inverse of EXP | 0.12 | | | | |
| 8 | | | | | | |
| 9 | So EXP(r) is the actual yield, per $1 initial investment, after a year | | | | | |
| 10 | when nominal interest rate r is compounded continuously. | | | | | |
| 11 | LN(x) is the nominal interest rate that would yield $x after a year, | | | | | |
| 12 | per $1 initial investment, with continuous compounding. | | | | | |
| 13 | | | | | | |
| 14 | AN EXAMPLE TO SHOW THE ADVANTAGE OF LOGARITHMIC GROWTH RATES | | | | | |
| 15 | | Values: | Growth rates: | | | |
| 16 | Year 0 | 100 | Percent | Logarithmic | | |
| 17 | Year 1 | 80 | -20.00% | -22.31% | | |
| 18 | Year 2 | 100 | 25.00% | 22.31% | | |
| 19 | Average growth rate: | | 2.50% | 0.00% | | |
| 20 | | | | | | |
| 21 | FORMULAS FROM RANGE A1:D19 | | | | | |
| 22 | B5. =(1+B3/B4)^B4 | C17. =(B17-B16)/B16 | | | | |
| 23 | B6. =EXP(B3) | C18. =(B18-B17)/B17 | | | | |
| 24 | B7. =LN(B6) | D17. =LN(B17/B16) | | | | |
| 25 | D6. =EXP(1) | D18. =LN(B18/B17) | | | | |
| 26 | E6. =D6^B3 | C19. =AVERAGE(C17:C18) | | | | |
| 27 | | D19. =AVERAGE(D17:D18) | | | | |

FIGURE **4.2**
**Compound interest and growth rate calculations with EXP and LN**

*k* times during the year, where *k* is some positive integer. Then *k* times per year, the bank will pay you $r/k$ dollars for every dollar in your account (including past interest payments as well as your original investment). That is, *k* times during the year, the value of your account will be multiplied by $1 + r/k$, and so every dollar that you initially invested will yield $(1 + r/k)^k$ dollars at the end of the year. For example, Figure 4.2 shows that when the nominal interest rate $r$ is 0.12 (in cell B3) and the times compounded is $k = 4$ (in cell B4), then each dollar initially invested will grow to $(1 + r/k)^k = 1.1255$ dollars (in cell B5) at the end of the year. This 1.1255 ratio of the end-of-year value divided by the beginning-of-year value for the account may be called its *annual growth ratio*.

Now consider what happens when the same nominal interest rate $r$ is compounded more times per year. (Try $k = 12$ or $k = 365$ in cell B4 of Figure 4.2 to represent monthly or daily compounding.) As *k* becomes large, the annual growth ratio $(1 + r/k)^k$ approaches the limiting value of EXP($r$). That is, EXP($r$) is the annual growth ratio (the end-of-year value divided by the beginning-of-year value) for an account that pays the nominal interest rate $r$ but is compounded a very large number of times during the year.

Conversely, for any annual growth ratio $x$, the formula LN($x$) is the nominal annual interest rate that, if compounded many times during the year, would yield $x$ dollars at the end of the year for every dollar invested at the beginning of the year. We may call LN($x$) the *logarithmic growth rate* that corresponds to the growth ratio $x$.

Myerson_C04.qxd 09-09-2004 11:55 Page 115

| | H | I | J | K | L |
|---|---|---|---|---|---|
| 16 | Other facts about EXP and LN: | | | | |
| 17 | EXP(r1+r2) = EXP(r1)*EXP(r2) | | | | |
| 18 | LN(x1*x2) = LN(x1)+LN(x2) | | | | |
| 19 | | | | | |
| 20 | Example: | | | | |
| 21 | r (random) | s (random) | | | |
| 22 | -0.8441 | 1.1442 | | | |
| 23 | x = EXP(r) | y = EXP(s) | | EXP(r+s) | EXP(r)*EXP(s) |
| 24 | 0.4299 | 3.1401 | | 1.3500 | 1.3500 |
| 25 | LN(x) | LN(y) | | LN(x*y) | LN(x)+LN(y) |
| 26 | -0.8441 | 1.1442 | | 0.3001 | 0.3001 |
| 27 | | | | | |
| 28 | FORMULAS FROM RANGE H21:L26 | | | | |
| 29 | H22. =NORMINV(RAND(),0,1) | | | | |
| 30 | I22. =NORMINV(RAND(),0,1) | | | | |
| 31 | H24. =EXP(H22) | | H26. =LN(H24) | | |
| 32 | I24. =EXP(I22) | | I26. =LN(I24) | | |
| 33 | K24. =EXP(H22+I22) | | K26. =LN(H24*I24) | | |
| 34 | L24. =H24*I24 | | L26. =H26+I26 | | |

FIGURE **4.3**
**The multiplication and addition properties of EXP and LN**

If an investment that was worth $w$ dollars at the beginning of the year becomes worth $y$ dollars at the end of the year, then people would traditionally say that its growth rate during the year has been $(y - w)/w$. But its logarithmic growth rate is the natural logarithm of the growth ratio, that is, LN($y/w$). When the growth ratio y/$w$ is close to 1, these logarithmic growth rates are close to the traditional growth rates.

Here is an example to show the advantage of using logarithmic growth rates instead of traditional growth rates in sophisticated quantitative analysis. Consider a two-year investment that was worth $100 initially, was worth $80 at the end of the first year, and was worth $100 again at the end of the second year. In traditional terms, the growth rate in the first year was (80 – 100)/100 = –0.20, and the growth rate in the second year was (100 – 80)/80 = 0.25. So the average growth rate for these two years seems to be (–0.20 + 0.25)/2 = 0.025, which gives us a positive average annual growth rate during a two-year period when the value did not grow at all! But the logarithmic growth rates in this example are LN(80/100) = –0.2231 in the first year and LN(100/80) = 0.2231 in the second year, which correctly yields an average annual growth rate of (–0.2231 + 0.2231)/2 = 0 over the two-year period.

There are two other formulas that you should know about EXP and LN:

$$\text{EXP}(x + y) = \text{EXP}(x)*\text{EXP}(y)$$

$$\text{LN}(r*s) = \text{LN}(r) + \text{LN}(s)$$

(Here $x$ and $y$ can be any two numbers, and $r$ and $s$ can be any two positive numbers.) The first of these formulas tells us that EXP converts addition into multiplication, while the second tells us that LN converts multiplication into addition. These formulas are illustrated in Figure 4.3 by the equality between cells K24 and L24 and by the equality between cells K26 and L26.

Myerson_C04.qxd 09-09-2004 11:55 Page 116

## 4.3 Lognormal Distributions

By definition, a random variable **Y** is a *Lognormal random variable* if its natural logarithm LN(**Y**) is a Normal random variable. Because the EXP function is the inverse of LN, a Lognormal random variable can be equivalently defined as any random variable that can be computed by applying the EXP function to a Normal random variable. Thus, for any number *m* and any positive number *s*,

EXP(NORMINV(RAND( ), *m, s*))

is a Lognormal random variable. These parameters *m* and *s* are called the *log-mean* and *log-standard-deviation* of the random variable. The log-mean and the log-standard-deviation are not the mean and standard deviation of this Lognormal random variable, instead they are the mean and standard deviation of its natural logarithm.

So the Lognormal random variables have a family of probability distributions that can be parameterized by their log-mean and log-standard-deviation. If you multiply a Lognormal random variable with log-mean *m* and log-standard-deviation *s* by a (nonrandom) positive constant *c*, the result is another Lognormal random variable that has log-mean $m + \text{LN}(c)$ and the same log-standard-deviation *s*. The EXP function always yields positive values, and so a Lognormal random variable is always positive.

Lognormal random variables are often used in finance to describe the growth ratio of an asset's value over some period of time. To see why, let **Y** denote the growth ratio for the value of some financial asset over the coming year. That is, **Y** is the unknown value of this asset a year from today divided by the asset's value today. For each integer *i* from 1 to 52, let $\mathbf{X}_i$ denote the growth ratio of this asset during the *i*th week of the coming year. So

$$\mathbf{Y} = \mathbf{X}_1 * \mathbf{X}_2 * \cdots * \mathbf{X}_{52}$$

Taking the natural logarithm of both sides, we get

$$\text{LN}(\mathbf{Y}) = \text{LN}(\mathbf{X}_1 * \mathbf{X}_2 * \cdots * \mathbf{X}_{52}) = \text{LN}(\mathbf{X}_1) + \text{LN}(\mathbf{X}_2) + \cdots \text{LN}(\mathbf{X}_{52})$$

because LN converts multiplication to addition. Now suppose that the growth ratio each week is drawn independently from the same probability distribution. Suppose also that there cannot be a large change in any one week, that is, the growth ratio $\mathbf{X}_i$ in any one week must be close to 1. Then $\text{LN}(\mathbf{X}_1), \ldots, \text{LN}(\mathbf{X}_{52})$ will be independent random variables, each drawn from the same distribution, and their values must be close to LN(1) = 0. Now recall (from Section 4.1) that a Normal probability distribution will be a good fit for any unknown quantity that can be expressed as the sum of many small (near 0) contributions that are independently drawn from a fixed probability distribution. So as the sum of all these small $\text{LN}(\mathbf{X}_i)$ random variables, LN(**Y**) must be an approximately Normal random variable.

Myerson_C04.qxd 09-09-2004 11:55 Page 117

This approximate Normality of LN(**Y**) implies that the growth ratio **Y** itself must be an approximately Lognormal random variable. Furthermore, when we multiply this growth ratio **Y** by the known current value of this asset, we find that this asset's value a year from today should also be an approximately Lognormal random variable.

When we suggest that a growth ratio should be a Lognormal random variable, we are equivalently suggesting that the corresponding logarithmic growth rate should be a Normal random variable, because **Y** is a Lognormal random variable when its logarithm LN(**Y**) is a Normal random variable. The values of assets that are traded in the stock market can never become negative, and so their growth ratios over any period of time cannot be negative. Thus it may seem intuitive to assume that such growth ratios are Lognormal random variables, because Lognormal random variables also cannot be negative. On the other hand, a logarithmic growth rate can be positive or negative, just as a Normal random variable can be positive or negative.

To see how the Normal and Lognormal distributions can arise in practice, consider a simple repeated-investment process, which is illustrated in Figure 4.4. In each round of this process, the decision maker decides how much money to invest. Then a fair coin is tossed and if it is Heads, then the return is 2 times the amount invested, but if it is Tails, then the return is 1/4 of the amount invested. So the round starts with $W$ dollars and $x$ dollars are invested, then the wealth at the end of the round is either $W - x + 2*x$ or $W - x + x/4$, each with probability 1/2.

Consider now a decision maker who starts with $1000 and gets to play 100 rounds of this investment process. Figure 4.4 analyzes the results of two different strategies that the decision maker could use. The first strategy always invests a *fixed amount*, $100 at each round. The second strategy always invests a *fixed fraction*, 10% of the wealth, at each round. Notice that in the first round, both strategies invest $100, but thereafter their results may differ.

The results of these two strategies are simulated in the spreadsheet of Figure 4.4. The returns on amounts invested in 100 independent rounds are simulated in cells A8:A107 by entering

```
=IF(RAND( )<0.5,2,0.25)
```

into cell A8 and copying it to A8:A107.

For the first strategy, the fixed $100 amount to be invested at each round is entered into cell B5 of Figure 4.4. Cell B7 contains the decision maker's initial wealth, $1000. Then the wealth at the end of the first round is computed in cell B8 by the formula

```
=B7+$B$8*(A8–1)
```

Copying cell B8 to cells B8:B107 then simulates the results of this fixed-amount strategy in the 100 repetitions of this investment process. (Notice that cell B8's reference to B7 is adjusted in each copy to become the decision

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | EXAMPLE: Repeated independent investments returning $2.00 or $0.25 per $1 invested. | | | | | | | | | | | | | |
| 2 | | | | | | Fixed-amount model | | | Fixed-fraction model | | | | | |
| 3 | Investment strategy: | | | | | Realizd | Approx distns | | Realizd | Approximating distns | | | | |
| 4 | FixedAmount | | FixedFractn | | %Rank | final$ | NORM | GENL | final$ | ExpNORM | LNORM | GENL | | |
| 5 | | 100 | 0.1 | | 0.01 | 325 | 246 | 439 | 354 | 327 | 331 | 396 | | |
| 6 | | Wealth | Wealth | | 0.1 | 1200 | 1141 | 1252 | 842 | 794 | 800 | 887 | | |
| 7 | retn | 1000 | 1000 | | 0.2 | 1550 | 1518 | 1595 | 1190 | 1153 | 1160 | 1245 | | |
| 8 | 2.00 | 1100 | 1100 | | 0.3 | 1725 | 1790 | 1842 | 1416 | 1509 | 1516 | 1589 | | |
| 9 | 2.00 | 1200 | 1210 | | 0.4 | 2075 | 2022 | 2053 | 2002 | 1899 | 1906 | 1959 | | |
| 10 | 2.00 | 1300 | 1331 | | 0.5 | 2250 | 2239 | 2250 | 2381 | 2355 | 2360 | 2381 | | |
| 11 | 0.25 | 1225 | 1231 | | 0.6 | 2425 | 2456 | 2447 | 2831 | 2919 | 2923 | 2894 | | |
| 12 | 2.00 | 1325 | 1354 | | 0.7 | 2600 | 2688 | 2658 | 3367 | 3673 | 3674 | 3566 | | |
| 13 | 0.25 | 1250 | 1253 | | 0.8 | 2950 | 2960 | 2905 | 4761 | 4808 | 4802 | 4554 | | |
| 14 | 0.25 | 1175 | 1159 | | 0.9 | 3300 | 3337 | 3248 | 6733 | 6982 | 6961 | 6392 | | |
| 15 | 0.25 | 1100 | 1072 | | 0.99 | 4175 | 4232 | 4061 | 16013 | 16936 | 16815 | 14301 | | |
| 16 | 2.00 | 1200 | 1179 | | | | | | FORMULAS | | | | | |
| 17 | 0.25 | 1125 | 1091 | | Mean | 2239 | 3370 | 7.76412 | A8. =IF(RAND()<0.5,2,0.25) | | | | | |
| 18 | 2.00 | 1225 | 1200 | | Stdev | 857 | 3435 | 0.84815 | B8. =B7+$B$5*(A8-1) | | | C8. =C7*(1+$C$5*(A8-1)) | | |
| 19 | 0.25 | 1150 | 1110 | | Q1 | 1725 | 1416 | | A8:C8 copied to A8:C107 | | | | | |
| 20 | 2.00 | 1250 | 1221 | | Q2 | 2250 | 2381 | | F24. =B107 | | G24. =C107 | | | |
| 21 | 2.00 | 1350 | 1343 | | Q3 | 2775 | 4004 | | H25. =LN(G25) | | H25 copied to H25:H525 | | | |
| 22 | 2.00 | 1450 | 1477 | | | Final Wealth | | | F17. =AVERAGE(F25:F525) | | | F18. =STDEV(F25:F525) | | |
| 23 | 0.25 | 1375 | 1366 | | | FixAmt | FixFrac | | F19. =PERCENTILE(F25:F525,0.25) | | | | | |
| 24 | 2.00 | 1475 | 1503 | | SimTable | 4525 | 22646 | LN | F20. =PERCENTILE(F25:F525,0.5) | | | | | |
| 25 | 2.00 | 1575 | 1653 | | 0 | 3300 | 6733 | 8.81482 | F21. =PERCENTILE(F25:F525,0.75) | | | | | |
| 26 | 2.00 | 1675 | 1818 | | 0.002 | 2950 | 4761 | 8.46827 | F17:F21 copied to G17:G21, F17:F18 to H17:H18 | | | | | |
| 27 | 2.00 | 1775 | 2000 | | 0.004 | 3825 | 11323 | 9.33463 | F5. =PERCENTILE($F$25:$F$525,E5) | | | | | |
| 28 | 0.25 | 1700 | 1850 | | 0.006 | 2950 | 4761 | 8.46827 | G5. =NORMINV(E5,$F$17,$F$18) | | | | | |
| 29 | 2.00 | 1800 | 2035 | | 0.008 | 1375 | 1001 | 6.90883 | H5. =GENLINV(E5,$F$19,$F$20,$F$21) | | | | | |
| 30 | 0.25 | 1725 | 1883 | | 0.01 | 4875 | 32025 | 10.3743 | I5. =PERCENTILE($G$25:$G$525,E5) | | | | | |
| 31 | 2.00 | 1825 | 2071 | | 0.012 | 1375 | 1001 | 6.90883 | J5. =EXP(NORMINV(E5,$H$17,$H$18)) | | | | | |
| 32 | 2.00 | 1925 | 2278 | | 0.014 | 1900 | 1684 | 7.42864 | K5. =LNORMINV(E5,$G$17,$G$18) | | | | | |
| 33 | 0.25 | 1850 | 2107 | | 0.016 | 3125 | 5662 | 8.64155 | L5. =GENLINV(E5,$G$19,$G$20,$G$21) | | | | | |
| 34 | 0.25 | 1775 | 1949 | | 0.018 | 2250 | 2381 | 7.77519 | F5:L5 copied to F5:L15 | | | | | |

FIGURE **4.4** **Repeated investment strategies yielding payoffs that are approximately Normal or Lognormal**

Myerson_C04.qxd 09-09-2004 11:55 Page 118

Myerson_C04.qxd 09-09-2004 11:55 Page 119

maker's wealth from the end of the previous round, to which this round then adds either 100*(2 – 1) = $100 or 100*(1/4 – 1) = –$75.) The final wealth after 100 rounds of this strategy is contained in cell B107, which is echoed in cell F24 (by the formula =B107).

For the second strategy, the fixed fraction 0.1 to be invested at each round is entered into cell C5 of Figure 4.4. Cell C7 again contains the decision maker's initial wealth, $1000. Then the wealth at the end of the first round is computed in cell C8 by the formula

```
=C7*(1+$C$8*(A8-1))
```

Copying cell C8 to cells C8:C107 then simulates the results of this fixed-amount strategy in the 100 repetitions of this investment process. (Again, notice that cell C8's reference to C7 is adjusted in each copy to become the decision maker's wealth from the end of the previous round, which this round then multiplies either by 1 + 0.1*(2 – 1) = 1.1 or by 1 + 0.10(1/4 – 1) = 0.925.) The final wealth after 100 rounds of this strategy is contained in cell C107, which is echoed in cell G24.

The spreadsheet in Figure 4.4 has a simulation table containing data from 501 independent simulations of this 100-round investment process. Cells F25:F525 contain simulated 501 values of the final wealth from the first (fixed-amount) strategy. Cells G25:G525 contain 501 values of the final wealth from the second (fixed-fraction) strategy. Percentile statistics from this simulation data are summarized in cells F6:F15 and cells I6:I15. Notice that cells E6:E15 contain selected cumulative probability fractions from 0.01 to 0.99. For each of these fractions, the value that is greater than this fraction of the simulated payoffs from the first strategy is computed in the corresponding cell in F5:F15 by entering the formula

```
=PERCENTILE($F$25:$F$525,E5)
```

into cell F5, and then copying F5 to cells F5:F15. The corresponding percentile values in the simulated payoff data for the second strategy are listed in cells I5:I15 by entering the formula

```
=PERCENTILE($G$25:$G$525,E5)
```

into cell I5, and then copying I5 to I5:I15.

Under the first strategy, the decision maker's wealth changes by adding an independent draw from the same distribution (+100 or –75, each with probability 1/2) in each of the 100 rounds. This repeated independent addition implies that the final wealth after 100 rounds with this strategy should have an approximately Normal distribution. One characteristic of a Normal distribution is a symmetry around the mean of the distribution. Notice in this data, for example, that the 0.10 percentile level (1200, in cell F6) is $1050 below the median 0.50 percentile level (2250, in cell F10), whereas the 0.90 percentile level (3300, in cell F14) is $1050 above the median.

Myerson_C04.qxd 09-09-2004 11:55 Page 120

To verify the approximate Normality of the first strategy's payoff distribution, the mean and standard deviation from this payoff data are calculated in cells F17 and F18 by the formulas =AVERAGE(F25:F525) and =STDEV(F25:F525), respectively. Then the corresponding percentile levels of a Normal distribution with this mean and standard deviation are computed in cells G5:G15 by entering the formula

=NORMINV(E5,$F$17,$F$18)

into cell G5 and then copying G5 to G5:G15. The close similarity of these Normal percentiles in G5:G15 to the actual percentiles of the simulation data in cells F5:F15 is strong evidence that these payoffs from the fixed-amount strategy are indeed approximately Normal.

Under the second strategy, the decision maker's wealth is multiplied by a growth ratio drawn independently from the same distribution (1.10 or 0.925, each with probability 1/2) in each of the 100 rounds. This repeated independent multiplication implies that the final wealth after 100 rounds with this strategy should have an approximately Lognormal distribution. The symmetry of the Normal distribution is conspicuously absent from this fixed-fraction payoff data. Notice, for example, that the 0.10 percentile level (842, in cell I6) is only $1539 below the median 0.50 percentile level (2381, in cell I10), whereas the 0.90 percentile level (6733, in cell I14) is $4352 above the median. But if we compute ratios, then we may find a different kind of symmetry (842/2381 = 0.354 = 2381/6733) that is characteristic of Lognormal distributions.

To verify the approximate Lognormality of the second strategy's payoff distribution is a bit more complicated. The payoffs are Lognormal if their logarithms (LN) are Normal. So the logarithms of the second strategy's final payoffs are computed in Figure 4.4 by entering =LN(G25) into cell H25 and then copying H25 to H25:H525. Then the mean and standard deviation of these logarithms are computed in cells H17 and H18 by the formulas =AVERAGE(H25:H525) and =STDEV(H25:H525), respectively. This mean and standard deviation of the payoffs' logarithms are our estimates of the log-mean and log-standard-deviation of the payoffs' Lognormal distribution. The corresponding percentile levels of a Lognormal distribution with this log-mean and log-standard-deviation are computed in cells J5:J15 by entering the formula

=EXP(NORMINV(E5,$H$17,$H$18))

into cell J5 and then copying J5 to J5:J15. The close similarity of these Lognormal percentiles in J5:J15 to the actual percentiles of the simulation data in cells I5:I15 is strong evidence that these payoffs from the fixed-fraction strategy are indeed approximately Lognormal.

In some situations, you may prefer to parameterize a Lognormal random variable by its own expected value and standard deviation instead of by the expected value and standard deviation of its natural logarithm. The formulas for computing the mean $\mu$ and the standard deviation $\sigma$ of a Lognormal

Myerson_C04.qxd 09-09-2004 11:55 Page 121

random variable from its log-mean $m$ and log-standard-deviation $s$ are rather complicated:

$$\mu = \text{EXP}(m + 0.5*s\text{^}2), \qquad \sigma = ((\text{EXP}(s\text{^}2) - 1)\text{^}0.5)*\mu$$

To let you avoid these formulas, Simtools provides a function called LNORMINV, which is the inverse cumulative probability function for Lognormal random variables parameterized by the mean and standard deviation. That is, when the mean $\mu$ and the standard deviation $\sigma$ are derived from the log-mean $m$ and the log-standard-deviation $s$ according to the (complicated) two formulas just shown, then for any cumulative probability $q$ between 0 and 1, the formula LNORMINV($q$, $\mu$, $\sigma$) will yield the same value as EXP(NORMINV($q$, $m$, $s$)). Thus, a Lognormal random variable that has expected value $\mu$ and standard deviation $\sigma$ can be simulated in Excel with Simtools by the formula

LNORMINV(RAND( ), $\mu$, $\sigma$)

This second parameterization of the Lognormal distribution is applied to our repeated-investment example in cells K5:K15 of Figure 4.4. Notice that the mean and standard deviation of the payoffs from the second (fixed-fraction) strategy are computed in cells G17 and G18 by the formulas =AVERAGE(G25:G525) and =STDEV(G25:G525), respectively. For the selected cumulative probabilities in E5:E15, the corresponding percentile levels of a Lognormal distribution with this (true) mean and standard deviation are then computed in cells K5:K15 by entering the formula

=LNORMINV(E5,$G$17,$G$18)

into cell K5 and then copying K5 to K5:K15. The similarity of these Lognormal percentiles in K5:K15 to the actual percentiles of the simulation data in cells I5:I15 is, again, evidence that these payoffs from the fixed-fraction strategy are indeed approximately Lognormal.

Cells J5:J15 and K5:K15 in Figure 4.4 contain Lognormal distributions that were statistically estimated from the same data by two different methods. Statistical theory suggests that the logarithmic estimation method used in J5:J15 may be more accurate in general, but the differences between these two estimates should become small as the size of our sample becomes large.

Recall that a Normal random variable can be either positive or negative, but a Lognormal random variable must always be positive. Under the decision-maker's strategy of investing a fixed 0.1 fraction of wealth, the wealth will always be positive in this example, because the gambles would become proportionally smaller if the wealth became small. But a strategy of investing a fixed $100 amount in every round could indeed cause the wealth to become negative. The probability of a negative final wealth under this fixed-amount strategy is only about 0.005, as may be estimated in the spreadsheet of Figure 4.4 by the formula NORMSDIST((0 – F17)/F18). (Of course, such negative outcomes would not occur in this example if the

Myerson_C04.qxd 09-09-2004 11:55 Page 122

decision maker is not allowed to invest more than the current wealth, but then the decision maker might also be unable to implement the strategy of investing $100 in every round.)

## 4.4 Generalized Lognormal Distributions

Now we introduce a three-parameter family of continuous probability distributions that may be called the *Generalized-Lognormal* family. These Generalized-Lognormal distributions are defined to include the probability distributions of all random variables of the form $c*\mathbf{X} + d$, where $\mathbf{X}$ is a Normal or Lognormal random variable and $c$ and $d$ are nonrandom constants. So these Generalized-Lognormal distributions include the Normals and Lognormals as special cases.

Any Generalized-Lognormal probability distribution can be characterized by specifying its value at the 0.25, 0.50, and 0.75 cumulative probability levels, which we may call the *quartile boundary points* for the distribution, denoted by $Q_1$, $Q_2$, and $Q_3$, respectively. That is, we may say that $Q_1$, $Q_2$, and $Q_3$ are the quartile boundary points for an unknown quantity $\mathbf{X}$ if

$$P(\mathbf{X} < Q_1) = 0.25, \quad P(\mathbf{X} < Q_2) = 0.50, \quad \text{and} \quad P(\mathbf{X} < Q_3) = 0.75$$

These three boundary points divide the number line into four separate *quartiles,* which are equally likely to include the actual value of $\mathbf{X}$. The first quartile is the set of all numbers less than $Q_1$, the second quartile is the interval between $Q_1$ and $Q_2$, the third quartile is the interval between $Q_2$ and $Q_3$, and the fourth quartile is the set of numbers greater than $Q_3$. When $\mathbf{X}$ is a continuous random variable, each of these quartiles has the same probability, 1/4, of including the actual value of $\mathbf{X}$. The second quartile boundary point, $Q_2$, is also called the *median* of the distribution of $\mathbf{X}$.

Simtools provides a function called GENLINV, which is the inverse cumulative probability function for a Generalized-Lognormal probability distribution parameterized by its quartile boundary points. That is, if $Q_1$, $Q_2$, and $Q_3$ are any three numbers such that $Q_1 < Q_2 < Q_3$, then the formula

$$\text{GENLINV(RAND( ),}Q_1, Q_2, Q_3)$$

returns a Generalized-Lognormal random variable that has quartile boundary points $Q_1$, $Q_2$, and $Q_3$.

In the special case where the quartile boundary points satisfy the equal-difference condition

$$Q_3 - Q_2 = Q_2 - Q_1$$

this Generalized-Lognormal random variable is a Normal random variable with mean

$$\mu = Q_2$$

Myerson_C04.qxd 09-09-2004 11:55 Page 123

and standard deviation

$$\sigma = (Q_3 - Q_1)/1.349$$

On the other hand, in the special case where the quartile boundary points satisfy the equal-ratio condition

$$Q_3/Q_2 = Q_2/Q_1$$

the Generalized-Lognormal random variable is a Lognormal random variable with log-mean

$$m = \text{LN}(Q_2)$$

and log-standard-deviation

$$s = (\text{LN}(Q_3) - \text{LN}(Q_1))/1.349$$

Thus, the GENLINV function gives us another way of constructing Normal or Lognormal random variables, based on their quartiles.

For the repeated-investment process, the quartile boundary points for the payoffs from the first (fixed-amount) strategy are estimated in cells F19:F21 of Figure 4.4 based on the simulation data in F25:F525. For example, cell F19 contains the formula

```
=PERCENTILE(F25:F525,0.25)
```

Then, for the selected cumulative probabilities in cells E5:E16, the corresponding percentile levels of a Generalized-Lognormal distribution with these three quartile boundary points are computed in cells H5:H15 by entering the formula

```
=GENLINV(E5,$F$19,$F$20,$F$21)
```

into cell H5 and then copying H5 to H5:H15. Similarly, the quartile boundary points for the payoffs from the second (fixed-fraction) strategy are estimated in cells G19:G21 from the simulation data below. The corresponding percentile levels of a Generalized-Lognormal distribution with these quartile boundary points are computed in cells L5:L15 by entering the formula

```
=GENLINV(E5,$G$19,$G$20,$G$21)
```

into cell L5 and then copying L5 to L5:L15. Notice that these computed Generalized-Lognormal percentiles in L5:L15 are reasonably close to the second strategy's actual data percentiles in I5:I15. Notice also that the computed Generalized-Lognormal percentiles in H5:H15 are similarly close to the actual data percentiles in F5:F15. Thus, the Generalized-Lognormal offers a reasonable approximation to the payoff distributions from both the fixed-amount and the fixed-fraction investment strategies.

Statistical theory suggests that the estimation methods used in cells G5:G15 and in cells J5:J15 will generally tend to be more accurate than the Generalized-Lognormal quartile-boundary estimation method used in cells H5:H15 and L5:L15 of Figure 4.4. But the differences between these various

Myerson_C04.qxd 09-09-2004 11:55 Page 124

estimation techniques should become small when the size of the data sample becomes large.

**TECHNICAL NOTE** (For people who want more technical details about GENLINV, the following explains how one could equivalently compute these Generalized-Lognormal random variables without Simtools.xla. First, given the quartile boundary points $Q_1$, $Q_2$, and $Q_3$, compute the ratio of quartile widths

$$\beta = (Q_3 - Q_2)/(Q_2 - Q_1)$$

If β is different from 1, then GENLINV(RAND( ), $Q_1$, $Q_2$, $Q_3$) is equivalent to the formula

$$Q_2 + (Q_3 - Q_2)*(\beta\text{^}(\text{NORMINV}(\text{RAND}(\,), 0, 1/0.67449) - 1)/(\beta - 1)$$

If β is equal to 1, then GENLINV(RAND( ), $Q_1$, $Q_2$, $Q_3$) is instead equivalent to the formula

$$\text{NORMINV}(\text{RAND}(\,), Q_2, (Q_3 - Q_2)/0.67449))$$

Beyond the cumulative probability parameter and the three quartile boundary points, the GENLINV function can also be given two other optional parameters that impose bounds on the possible values of GENLINV. If you want to make sure that the value of GENLINV never goes below some lowest value *L*, then you can include this number *L* as the optional fifth parameter. The value of the formula GENLINV(*p*, $Q_1$, $Q_2$, $Q_3$, *L*) is the same as GENLINV(*p*, $Q_1$, $Q_2$, $Q_3$) unless it would be less than *L*, in which case the value is rounded up to *L*. For example, GENLINV(RAND( ), 2, 5, 9) is negative with probability 0.103, whereas GENLINV(RAND( ), 2, 5, 9, 0) is never negative but instead equals 0 with probability 0.103. Similarly, if you want to make sure that the value of GENLINV never goes above some highest value *H*, then you can include this number *H* as the optional sixth parameter. The value of the formula GENLINV(*p*, $Q_1$, $Q_2$, $Q_3$, *L*, *H*) is the same as GENLINV(*p*, $Q_1$, $Q_2$, $Q_3$, *L*) unless it would be greater than *H*, in which case the value is rounded down to *H*.

## 4.5 Subjective Probability Assessment

There are some unknown quantities for which finding comparable data may be relatively straightforward. For example, if we want to predict the rate of return next year for a share of stock in an old established company, we might naturally assume that next year's annual rate of return will be another independent draw out of the same probability distribution that generated the past 30 years' annual rates of return for this stock. In such cases, we can collect data about past realizations of this quantity, and we can use our sample's mean and standard deviation to estimate the mean

Myerson_C04.qxd 09-09-2004 11:55 Page 125

and standard deviation (or the log-mean and log-standard-deviation) of a Normal or Lognormal distribution that is appropriate for this unknown future rate of return.

But we may also be confronted by unknown quantities that have a more unique quality, for which there may not be any obviously comparable data of past realizations. For example, if we want to estimate the potential demand for a particular new product like the one in the Superior Semiconductor case, we might be very uncomfortable with a naive assumption that this new product will look like a random draw from the previous 30 new products that we have introduced. Such an assumption would ignore all the unique technical characteristics of this new product and its potential users, which should be shaping our beliefs in this area.

So let us now consider the problem of assessing a probability distribution for an unknown quantity where analysis of data about similar quantities seems to be out of the question. Even where there is no data, if we have uncertainty about an unknown quantity, then we should be able to describe our uncertainty by a probability distribution. To make this point forcefully, let us take up a specific example to which you might have thought that probability modeling could not even be applied: the question of when Napoleon Bonaparte was born.

Let **N** denote the year when Napoleon Bonaparte was born. We can make **N** a continuous random variable by allowing fractions, so that **N** = 1066.5 would mean that Napoleon was born in the middle of the year A.D. 1066. If we do not know when Napoleon was born, then from our perspective, **N** is an unknown quantity. But how can we define its probability distribution?

A probability distribution is a description of uncertainty and beliefs, and different people may have different beliefs about an unknown quantity such as **N**. A French professor of history might feel sure about the year in which Napoleon was born, whereas an American professor of economics might be unsure about the century in which he was born. To describe the different beliefs of different people, we must admit that they may subjectively assess different probability distributions for the same unknown. So we should not simply speak about "the" probability distribution of **N**. Instead, we should speak of a particular individual's *subjective probability* distribution for the unknown quantity **N**, as assessed at a particular point in time with reference to the information available to this individual at this point in time. This subjective probability distribution is thus defined to be a quantitative description of what this individual believed about this unknown quantity at this point in time.

When we set out to measure your subjective probability distribution for **N**, given your current information, you might feel a sense of reluctance if you feel disqualified by relative ignorance of French history. In some circumstances, an individual should be reluctant to express beliefs where others have better information. In the Superior Semiconductor case, the decision analyst and the executive decision maker rightly defer to the marketing manager's beliefs about the potential demand for the new T-regulator device.

Myerson_C04.qxd 09-09-2004 11:55 Page 126

So you might say that your beliefs about this unknown quantity should be the beliefs of the best expert that you can consult. And if the best expert that you can consult is an encyclopedia that has the exact date of Napoleon's birth, then of course you should check this encyclopedia and assign subjective probability 1 to the date that it gives you. But, to prepare ourselves for practical problems like demand forecasting where the encyclopedia cannot help us, let us suppose that you have no access to an encyclopedia or a French history professor. Instead, let us suppose that I am your best available expert, with the information and beliefs that I had at the time that I first wrote this essay on subjective probability assessment. In fact, I have taken this exercise seriously (resisting the temptation to open that forbidden encyclopedia) and have assessed my subjective probability distribution for this unknown quantity. As an illustration of subjective probability in practice, let me now describe how I assessed this subjective probability distribution.

The first thing that I began to do was to recall some of the things that I knew about Napoleon from high school. I knew that the French Revolution began in 1789, but I did not know what Napoleon was doing at that time. I knew that Napoleon invaded Russia in 1812, which I could remember because of the name of the famous 1812 Overture (by some famous composer). I believed that Napoleon died in the 1820s or 1830s, but it might have been a relatively young death from some kind of cancer. I felt quite sure that he had risen to be the most powerful man in France by 1799 (that date stuck in my mind for some reason), and I felt pretty sure that he had been remarkably young for such a position of power at that time. I recalled someone saying that revolutions have often been times of great opportunity for ambitious young men like Napoleon. So I knew that he had been born sometime in the 1700s, and I drew a time line for this century with intervals marked to represent the decade years 1710, 1720, through 1790. Then I started thinking about each of these years as possible birth years for Napoleon, starting in the middle, at 1750. I soon found myself calculating, for each of these years, how old Napoleon would have been in late 1799. It just seemed easier to think about these ages-in-1799 than to think about the birth years themselves.

You might have supposed that the first questions that I would undertake would be to assess an expected value and a standard deviation for my probability distribution, where the mean is supposed to be a measure of the center of my subjective probability distribution, and the standard deviation is supposed to be a measure of the amount of uncertainty around this center. But we have seen that expected values and standard deviations are rather sophisticated mathematical concepts that are not so easy to understand. Even I, a professional statistics teacher, would have trouble judging whether my subjective standard deviation of **N** should be (say) 4 years or 7 years. It is simpler (and thus better) to think first about assessing the probabilities of events.

Myerson_C04.qxd 09-09-2004 11:55 Page 127

I might try, for example, to assess my subjective probability of **N** being less than 1750 (that is, of Napoleon being older than 49 at the end of 1799). My subjective probability of the event "**N** < 1750" could be defined as the number $p$ such that I would be indifferent between the following two alternatives, if they were offered to me uninformatively:

1. a lottery ticket that pays \$1000 if **N** < 1750, and pays \$0 otherwise
2. a lottery ticket that pays \$1000 if a particular RAND( ) in a spreadsheet is less than $p$ after its next recalculation, and pays \$0 otherwise

By being "offered uninformatively," I mean that these lotteries are not supposed to be available only because someone who knows more than I do wants to see me lose (or win), so that getting such a choice should not change my beliefs about **N**. I could say that I would prefer alternative (1) above when $p$ is 1/2 or 1/10, but I would prefer alternative (2) when $p$ is 1/100, and so my subjectively assessed $P(\mathbf{N} < 1750)$ must be somewhere between 1/100 and 1/10. As I think about $p$ decreasing below 1/10, alternative (2) becomes worse until somewhere it ceases to be better than alternative (1) for me. But to assess that point of indifference that defines $P(\mathbf{N} < 1750)$, I may need to consider some probability numbers for which I have relatively little intuitive feeling. Is $p = 0.05437$ a number that makes lottery (2) just as good as lottery (1), or might it be $p = 0.06819$? Even a trained statistician might have little or no confidence in his ability to answer such questions subjectively.

So it is better to think instead about events that have simple probabilities like 1/2 or 1/4. Thus, subjective probability assessment in practice often begins with an assessment of the three quartile boundary points that (for an unknown quantity with a continuous distribution) divide the number line into four equally likely intervals.

Thus, the first question that I actually asked myself was: "For what number $Q_2$ would I be indifferent between (1a) a bet that pays me \$1000 if $\mathbf{N} < Q_2$ and pays me \$0 otherwise, and (2a) a bet that pays me \$1000 if $\mathbf{N} \geq Q_2$ and pays me \$0 otherwise?" After trying several numbers, I settled on $Q_2 = 1764$, because I felt that Napoleon was equally likely to have been more or less than 36 years old in late 1799. (Actually, with $Q_2 = 1764$, if I imagined I already had either of these two bets and that I was approached by someone with an offer to exchange it for the other bet, then I found that I always wanted to stay with the bet that I already had. But the intensity of my desire to reject the "new" offer was equally intense, whichever side I imagined myself having first, and so I felt that I really was indifferent.) Thus my subjective probabilities should satisfy the equations

$$P(\mathbf{N} < 1764) = P(\mathbf{N} \geq 1764) = 1/2$$

Next, I asked myself, for what number $Q_1$ would I be indifferent between (1b), a bet that pays me \$1000 if $\mathbf{N} < Q_1$, and (2b), a bet that pays me \$1000

Myerson_C04.qxd 09-09-2004 11:55 Page 128

if $Q_1 \le \mathbf{N} < 1764$? After going back and forth for a while, I finally decided that I would be indifferent for $x = 1758$. That is, I figured that the events "$\mathbf{N} < 1758$" (aged at least 42 in late 1799) and "$1758 \le \mathbf{N} < 1764$" (aged at least 36 but not yet 42 in late 1799) were equally likely, in my opinion. Since the union of these two equally likely and mutually exclusive events was "$\mathbf{N} < 1764$," which I had just assigned subjective probability 1/2, I could conclude

$$P(\mathbf{N} < 1758) = 1/4$$

Next, I asked myself, for what number $Q_3$ would I be indifferent between (1c), betting on the event "$1764 \le \mathbf{N} < Q_3$," and (2c), betting on the event "$\mathbf{N} \ge Q_3$"? After more introspection, I decided that I would be indifferent for $Q_3 = 1767$. That is, I figured that the events "$1764 \le \mathbf{N} < 1767$" (aged at least 33 but not yet 36 in late 1799) and "$\mathbf{N} \ge 1767$" (younger than 33 in late 1799) were equally likely, in my opinion. So I now had two more mutually exclusive and equally likely events whose union ("$\mathbf{N} \ge 1764$") had probability 1/2, and so I could conclude that $P(\mathbf{N} \ge 1767) = P(1764 \le \mathbf{N} < 1767) = 1/4$. Thus, I wrote that

$$P(\mathbf{N} < 1767) = 1 - P(\mathbf{N} \ge 1767) = 3/4$$

To check my judgments, I asked myself whether I would be indifferent between (1d), betting on the event "$1758 \le \mathbf{N} < 1767$" (aged at least 33 but not yet 42 in late 1799), and (2d), betting on the event "$\mathbf{N} > 1767$ or $\mathbf{N} \le 1758$" (younger than 33 or more than 42 years old in late 1799). This was a hard question, but I decided that the answer was probably yes, so I felt comfortable with my assessments. If I had felt that "$1758 \le \mathbf{N} < 1767$" was less likely than "$\mathbf{N} > 1767$ or $\mathbf{N} \le 1758$," then I would have reconsidered my earlier answers, to see whether I might have underestimated $Q_3$ or overestimated $Q_1$ in the preceding two paragraphs.

This process of assessing equally likely intervals had thus yielded three points on my (inverse) cumulative probability curve for this unknown **N**:

$$P(\mathbf{N} < 1758) = 0.25, \qquad P(\mathbf{N} < 1764) = 0.50, \qquad P(\mathbf{N} < 1767) = 0.75$$

I could have continued splitting intervals to get values with cumulative probability 0.125, 0.375, etc., but instead I tried a simpler tactic. I just assumed that my beliefs about this unknown quantity can be represented by a Generalized-Lognormal distribution with my subjective assessed quartile boundary points $Q_1 = 1758$, $Q_2 = 1764$, and $Q_3 = 1767$, and I used the GENLINV function with these parameters to see what the rest of the inverse cumulative distribution would look like. But this assumption needed to be checked by comparing some of the other probabilities that the distribution would predict with my subjective beliefs.

The shape of a Generalized-Lognormal may be symmetric around its median or it may be skewed, depending on the ratio of the widths of the two inner quartiles $Q_3 - Q_2$ and $Q_2 - Q_1$. In general, the possible values of a

Myerson_C04.qxd 09-09-2004 11:55 Page 129

Generalized-Lognormal distribution are bounded on the side of the shorter inner quartile, while the possible values extend indefinitely on the other side where the wider inner quartile expressed more uncertainty. With my assessed quartile boundaries, the ratio of quartile widths is (1767 – 1764)/(1764 – 1758) = 0.5, which expresses much more uncertainty on the low side than on the high side. So this Generalized Lognormal distribution is very skewed, with the region of substantial probability extending much farther on the low side than on the high side. In this case, the extreme values that have a probability 1/100 below or above them are

GENLINV(0.01, 1758, 1764, 1767) = 1704

GENLINV(0.99, 1758, 1764, 1767) = 1769

Notice that the distance from the low value with cumulative probability 0.01 (1704) to the lower quartile boundary point (1758) is much farther than the distance from the higher quartile boundary point (1767) to the high value with cumulative probability 0.99 (1769).

Looking at such low and high values in the probability distribution can be very important for deciding whether the estimated Generalized-Lognormal distribution is a good fit for the actual subjective probability distribution that we are trying to assess. So I asked myself, did I really believe that the probability of Napoleon being born before 1704 (aged 96 or more in late 1799) was 1/100? And did I really believe that the probability of Napoleon being born in 1769 or later (younger than 31 in late 1799) was also 1/100? An event that has probability 1/100 should be considered extremely unlikely, but not completely out of the range of reasonable possibilities. By these standards, I felt strongly that "$\mathbf{N} < 1704$" was not even in the range of reasonable possibilities, and so $P(\mathbf{N} < 1704)$ should be substantially less than 1/100. I also felt that, in my subjective probability distribution, $P(\mathbf{N} \geq 1769)$ was substantially larger than 1/100. So one possible conclusion was that my true beliefs about $\mathbf{N}$ could not be fit into the Generalized-Lognormal family of distributions. But rather than give up on such a nice functional form, it was worthwhile to ask instead whether I might have been somewhat mistaken about my subjective quartile boundary points. After all, it was not so easy to decide whether those intervals were equally likely.

So I thought next about what values with subjective probability 1/100 below or above would have seemed reasonable for GENLINV(0.01, $Q_1$, $Q_2$, $Q_3$) and GENLINV(0.99, $Q_1$, $Q_2$, $Q_3$). After some introspection about the low end of the likely values of $\mathbf{N}$, I estimated that the probability of $\mathbf{N}$ being less than 1745 (55 or older in late 1799) should be about 1/100. At the high end of the distribution, I estimated that the probability of $\mathbf{N}$ being greater than 1776 (not yet 24 in late 1799) should also be about 1/100.

Next I thought about small adjustments to my assessed quartiles. I felt more confident about the median value $Q_2$ that I had assessed, and so I focused on changing $Q_1$ and $Q_3$. The problem seemed to be that my first

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | | Quartiles of Unknown Quantity | | | | Ratio of quartile widths | | |
| 2 | | 1759 | 1764 | 1768 | | 0.8 | | |
| 3 | 1%-point | 25%point | 50%point | 75%point | 99%-point | | Simulated value | |
| 4 | 1740.82 | | | | 1774.74 | | 1770.60 | |
| 5 | | | | | | | | |
| 6 | GENERALIZED-LOGNORMAL DISTRIBUTION | | | | | | | |
| 7 | CumvProby | Value | ProbyDensity | | FORMULAS | | | |
| 8 | 0.001 | 1728.41 | | | A4. =GENLINV(0.01,$B$2,$C$2,$D$2) | | | |
| 9 | 0.01 | 1740.82 | 0.0012 | | E4. =GENLINV(0.99,$B$2,$C$2,$D$2) | | | |
| 10 | 0.02 | 1744.54 | 0.0034 | | G4. =GENLINV(RAND(),$B$2,$C$2,$D$2) | | | |
| 11 | 0.03 | 1746.74 | 0.0053 | | F2. =(D2-C2)/(C2-B2) | | | |
| 12 | 0.04 | 1748.31 | 0.0071 | | B8. =GENLINV(A8,$B$2,$C$2,$D$2) | | | |
| 13 | 0.05 | 1749.54 | 0.0089 | | B8 copied to B8:B108 | | | |
| 14 | 0.06 | 1750.55 | 0.0107 | | C9. =(A10-A8)/(B10-B8) | | | |
| 15 | 0.07 | 1751.41 | 0.0124 | | C9 copied to C9:C107 | | | |
| 16 | 0.08 | 1752.16 | 0.0140 | | | | | |
| 17 | 0.09 | 1752.83 | 0.0157 | Inverse Cumulative Chart plots (A8:A108,B8:B108) | | | | |
| 18 | 0.1 | 1753.44 | 0.0173 | | | | | |
| 19 | 0.11 | 1753.99 | 0.0189 | | | | | |
| 20 | 0.12 | 1754.50 | 0.0205 | | | | | |
| 21 | 0.13 | 1754.97 | 0.0220 | | | | | |
| 22 | 0.14 | 1755.41 | 0.0235 | | | | | |
| 23 | 0.15 | 1755.82 | 0.0250 | | | | | |
| 24 | 0.16 | 1756.21 | 0.0264 | | | | | |
| 25 | 0.17 | 1756.58 | 0.0279 | | | | | |
| 26 | 0.18 | 1756.93 | 0.0293 | | | | | |
| 27 | 0.19 | 1757.26 | 0.0307 | | | | | |
| 28 | 0.2 | 1757.58 | 0.0320 | | | | | |
| 29 | 0.21 | 1757.88 | 0.0333 | | | | | |
| 30 | 0.22 | 1758.18 | 0.0346 | | | | | |
| 31 | 0.23 | 1758.46 | 0.0359 | | | | | |
| 32 | 0.24 | 1758.74 | 0.0372 | | | | | |
| 33 | 0.25 | 1759.00 | 0.0384 | | | | | |
| 34 | 0.26 | 1759.26 | 0.0396 | Probability Density Chart plots (B9:B107,C9:C107) | | | | |
| 35 | 0.27 | 1759.51 | 0.0408 | | | | | |
| 36 | 0.28 | 1759.75 | 0.0419 | | | | | |
| 37 | 0.29 | 1759.98 | 0.0431 | | | | | |
| 38 | 0.3 | 1760.21 | 0.0442 | | | | | |
| 39 | 0.31 | 1760.43 | 0.0452 | | | | | |
| 40 | 0.32 | 1760.65 | 0.0463 | | | | | |
| 41 | 0.33 | 1760.87 | 0.0473 | | | | | |
| 42 | 0.34 | 1761.08 | 0.0483 | | | | | |
| 43 | 0.35 | 1761.28 | 0.0493 | | | | | |
| 44 | 0.36 | 1761.48 | 0.0502 | | | | | |
| 45 | 0.37 | 1761.68 | 0.0511 | | | | | |
| 46 | 0.38 | 1761.87 | 0.0520 | | | | | |
| 47 | 0.39 | 1762.06 | 0.0529 | | | | | |
| 48 | 0.4 | 1762.25 | 0.0537 | | | | | |
| 49 | 0.41 | 1762.44 | 0.0545 | | | | | |

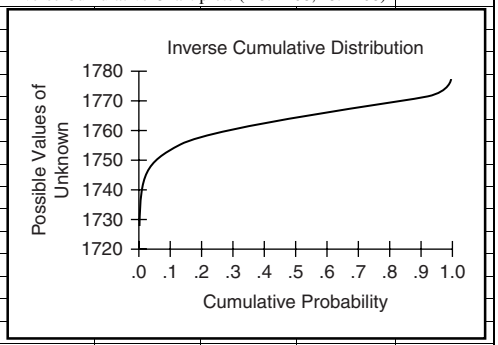


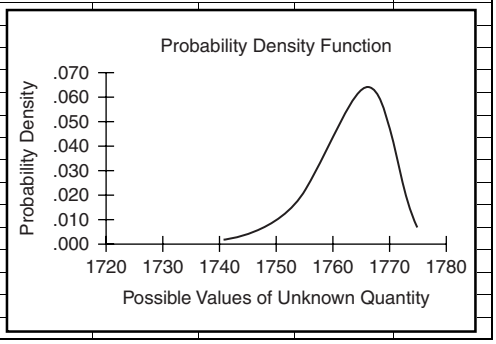


FIGURE **4.5**
**Subjectively assessed quartiles and a Generalized-Lognormal distribution**

estimated distribution was much too skewed, and so I tried increasing $Q_1$ and $Q_3$ each by 1, which brought the ratio of inner quartile widths up to

$$(Q_3 - Q_2)/(Q_2 - Q_1) = (1768 - 1764)/(1764 - 1759) = 0.8$$

The 0.01 and 0.99 cumulative probability values with these parameters became

GENLINV(0.01, 1759, 1764, 1768) = 1741

GENLINV(0.99, 1759, 1764, 1768) = 1775

Myerson_C04.qxd 09-09-2004 11:55 Page 131

which seemed compatible with my beliefs about these low and high ends of the distribution. Revisiting the questions that I originally used to assess the quartiles, I felt that these new estimates were not clearly less plausible than my original estimates, so I was willing to write

$$P(\mathbf{N} < 1759) = P(1759 \le \mathbf{N} < 1764) = P(1764 < \mathbf{N} \le 1768)$$
$$= P(1768 < \mathbf{N}) = 1/4$$

Thus, I concluded that the Generalized-Lognormal distribution with quartile boundaries at 1759, 1764, and 1768, as shown in Figure 4.5, was a reasonable description of my beliefs about Napoleon's birth year, given the information that I then had available. That is, I might have said that Napoleon's birth year was as likely to be in any given interval as the random variable GENLINV(RAND( ), 1759, 1764, 1768).

## 4.6 A Decision Problem with Discrete and Continuous Unknowns

In Part C of the Superior Semiconductor case, the decision analyst needed to assess quantitatively her colleagues' uncertainty about the development cost and potential market value of the proposed new product. For each of these unknown quantities, we may understand that she asked the best available expert some version of the quartile-assessment questions that were described in the previous section. That is, she may have asked the expert to specify three points that would divide the number line into four regions such that the expert could not say that any of these regions was more likely to contain the unknown quantity than any other region. Then by assuming a Generalized-Lognormal distribution for this unknown quantity, the decision analyst could easily extend these assessed quartiles to a full probability distribution. To check the fit of this assumed Generalized-Lognormal distribution, she might have asked the experts to think about all-but-1/100-probability lower and upper bounds for the unknown quantity, which she could compare to the 0.01 and 0.99 cumulative probability values returned by the GENLINV function.

As a result of such questioning, the chief production engineer assessed quartile boundary points for development cost equal to \$23, \$26, and \$29 million. Let us suppose that the Generalized-Lognormal assumption seemed to fit reasonably well the chief engineer's other beliefs about this unknown cost. These three quartile boundary points happen to satisfy the equal-difference condition that characterizes the Normal distributions among the Generalized-Lognormals, because

$$29 - 26 = 26 - 23$$

Myerson_C04.qxd 09-09-2004 11:55 Page 132

For a Normal distribution, the quartile boundary points $Q_1$, $Q_2$, $Q_3$ depend on the mean $\mu$ and standard deviation $\sigma$ by the formulas

$$Q_1 = \mu - 0.675*\sigma, \qquad Q_2 = \mu, \qquad Q_3 = \mu + 0.675*\sigma$$

So with these assessed quartiles, the mean and standard deviation of the development cost are

$$\mu = Q_2 = 26 \quad \text{and} \quad \sigma = (Q_3 - Q_2)/0.675 = (29 - 26)/0.675 = 4.444$$

(In Figure 4.1 we considered a Normal distribution with parameters close to these.) Thus, the development cost in this case could be simulated equally well by the formula

NORMINV(RAND( ), 26, 4.444)

or the formula

GENLINV(RAND( ), 23, 26, 29)

The quartile boundary points for total market value assessed by the marketing manager in this case are \$80, \$100, and \$125 million. Let us suppose that the Generalized-Lognormal assumption seemed also to fit reasonably well the marketing manager's other beliefs about this unknown market value. These quartiles happen to satisfy the equal ratio condition that characterizes the Lognormal distributions among the Generalized-Lognormals, because

125/100 = 100/80

For a Lognormal distribution with log-mean $m$ and log-standard-deviation $s$, the quartile boundary points $Q_1$, $Q_2$, $Q_3$ depend on the log-mean and log-standard-deviation by the formulas

$$\text{LN}(Q_1) = m - 0.6745*s, \qquad \text{LN}(Q_2) = m, \qquad \text{LN}(Q_3) = m + 0.6745*s$$

So with these assessed quartiles, the log-mean and log-standard-deviation of the total market value must be

$$m = \text{LN}(Q_2) = \text{LN}(100) = 4.605 \quad \text{and}$$
$$s = (\text{LN}(Q_3) - \text{LN}(Q_2))/0.6745 = (\text{LN}(125) - \text{LN}(100))/0.6745 = 0.3308$$

These log-mean and log-standard-deviation numbers are very hard to interpret, but they tell us that the total market value in this case could be simulated equally well by the formula

EXP(NORMINV(RAND( ), 4.605, 0.3308))

or the formula

GENLINV(RAND( ), 80, 100, 125)

Assuming that these unknown quantities are independent, the resulting simulation model is shown in Figure 4.6. The subjectively assessed

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | SUPERIOR SEMICONDUCTOR CASE, Part C | | | | | | | |
| 2 | Development-cost quartile points | | | | | P(Successful development) | | |
| 3 | 23 | 26 | 29 | | | 0.95 | | |
| 4 | | | | | | | | |
| 5 | Market-value quartile pts (if successful) | | | | | RiskTolerance | | |
| 6 | 80 | 100 | 125 | | | 36.49 | | |
| 7 | | | | | | | | |
| 8 | k | P(#competitors=k) | | | FORMULAS FROM RANGE A1:F27 | | | |
| 9 | 1 | 0.10 | | | A16. =GENLINV(RAND(),A3,B3,C3) | | | |
| 10 | 2 | 0.25 | | | A17. =IF(RAND()<F3,1,0) | | | |
| 11 | 3 | 0.30 | | | A18. =A17*GENLINV(RAND(),A6,B6,C6) | | | |
| 12 | 4 | 0.25 | | | A19. =DISCRINV(RAND(),A9:A13,B9:B13) | | | |
| 13 | 5 | 0.10 | | | A20. =A18/(1+A19)-A16 | | | |
| 14 | | | | | B27. =A20 | | | |
| 15 | Simulation model | | | | B23. =AVERAGE(B28:B828) | | | |
| 16 | 28.546 | Development cost | | | B24. =STDEV(B28:B828) | | | |
| 17 | 1 | Successful development? | | | E24. =COUNT(B28:B828) | | | |
| 18 | 156.379 | Total market value | | | F24. =B23-1.96*B24/(E24^0.5) | | | |
| 19 | 2 | Number of competitors | | | G24. =B23+1.96*B24/(E24^0.5) | | | |
| 20 | 23.580 | Profit | | | B25. =CE(B28:B828,F6) | | | |
| 21 | | | | | | | | |
| 22 | Analysis of simulation data | | | | | | | |
| 23 | | 1.772 | E(Profit) | | SampleSize | 95%ConfInt for E(Profit) | | |
| 24 | | 17.539 | Stdev(Profit) | | 801 | 0.557 | 2.986 | |
| 25 | | -1.7577 | CE(profit,RiskTol) | | | | | |
| 26 | | Profit ($millions) | | | | | | |
| 27 | SimTable | 23.580 | | | | | | |
| 28 | 0 | -37.2965 | | | | | | |
| 29 | 0.00125 | -35.7963 | | | | | | |
| 30 | 0.0025 | -33.4312 | | | | | | |
| 31 | 0.00375 | -31.807 | | | | | | |
| 32 | 0.005 | -31.4246 | | | | | | |
| 33 | 0.00625 | -31.2286 | | | | | | |
| 34 | 0.0075 | -31.2051 | | | | | | |
| 35 | 0.00875 | -31.0382 | | | | | | |
| 36 | 0.01 | -30.8203 | | | | | | |
| 37 | 0.01125 | -30.5732 | | | | | | |
| 38 | 0.0125 | -30.3894 | | | | | | |
| 39 | 0.01375 | -30.252 | | | | | | |
| 40 | 0.015 | -30.0745 | | | | | | |
| 41 | 0.01625 | -29.7856 | | | | | | |
| 42 | 0.0175 | -29.3954 | | | | | | |
| 43 | 0.01875 | -28.4644 | | | | | | |
| 44 | 0.02 | -28.165 | Simulation data in B28:B828 is sorted. | | | | | |
| 45 | 0.02125 | -28.1596 | Cumulative risk profile plots (A28:A828,B28:B828). | | | | | |

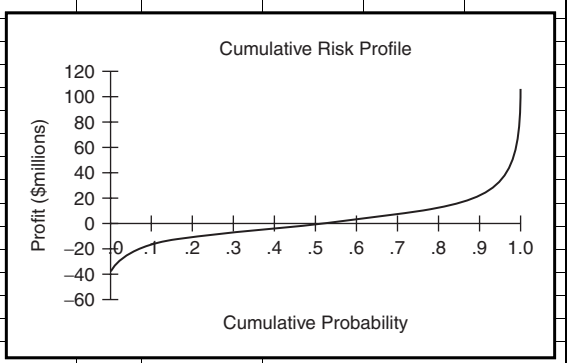


FIGURE **4.6**
**Simulation analysis of the Superior Semiconductor (C) Case**

probabilities and quartile points are listed in the top 13 rows of the spreadsheet. Then the simulation model is developed in rows 15–20. The development cost is simulated in cell A16 by the formula

=GENLINV(RAND( ),A3,B3,C3)

using the assessed quartile points for the development cost that are listed in cells A3:C3. The success or failure of the product development is simulated in cell A17 by the formula

=IF(RAND( )<F3,1,0)

Myerson_C04.qxd 09-09-2004 11:55 Page 134

where the value 1 denotes success, the value 0 denotes failure, and cell F3 contains the assessed probability of success (0.95). The total value of the market for this new product is simulated in cell A18 by the formula

=A17*GENLINV(RAND( ),A6,B6,C6)

using the assessed quartile points for the market value that are listed in cells A6:C6. In cell A18, the Generalized-Lognormal random variable is multiplied by A17, because the market value will actually be 0 if the new product cannot be developed. (If A18 is 1, then multiplying by it makes no difference, but if A18 is 0, then multiplying by it makes the total market value 0.) The number of competitors is simulated in cell A19 by the formula

=DISCRINV(RAND( ),A9:A13,B9:B13)

using the possible values listed in A9:A13 with corresponding probabilities in B9:B13. Finally, the bottom line for Superior Semiconductor is its profit, computed in cell B20 by the formula

=A18/(1+A19) – A16

That is, profit equals total market value divided by the number of firms in the market, minus development costs.

The use of simulation to estimate expected profit may have seemed unnecessary when we analyzed Part A of this case in Chapter 2, because we then had only one discrete random variable and we could easily compute the expected value directly. But now in Part C, we have four different random variables, some of which have infinitely many possible values. In examples like this, we really need simulation to estimate the expected profit because other direct methods of calculation may be difficult or impossible.

So the simulated profit was echoed in cell B27 of Figure 4.6 (by the formula =A20), and a simulation table contains the results of 801 independent recalculations of this simulated profit below in cells B28:B828. The sample mean and standard deviation from this data is computed in cells B23 and B24 to estimate the expected value and standard deviation of profit. These calculations suggest that we could still recommend this new product under the expected value criterion because the estimated expected profit is a positive number ($1.772 million). Analysis of more simulation data might give a different and more accurate estimate of the expected profit, but the 95% confidence interval computed in cells F24:G24 makes us quite confident that the expected profit is bigger than 0, because the lower bound in F24 (0.557) is positive. (A previous analysis using a sample size of 401 yielded a 95% confidence interval that included negative values, and so the sample size was increased to the 801 shown here to be able to confidently decide this question.)

So the expected value analysis for Part C of this case seems quite similar to Part A in Chapter 2, where uncertainties about development cost, development success, and market value were all ignored. But the measures of risk for Part C look very different from Part A. To begin with, the standard deviation

Myerson_C04.qxd 09-09-2004 11:55 Page 135

of profit is reported now to be about $17.5 million, more than the standard deviation of about $9 million computed for Part A. With the risk tolerance of $36.49 million that was assessed in Chapter 3 for Part B of the case, the estimated certainty equivalent becomes –$1.76 million, which is computed in cell B25 of Figure 4.6 by the formula =CE(B28:B828,F6). To get a positive certainty equivalent for this simulated profit data, the risk tolerance in cell F6 would have to be increased to at least $79 million (which could be assessed if the decision maker said that a lottery paying $20 million or –$10 million, each with probability 0.5, would be worth at least $3.6 million for sure). But if the risk tolerance of $36.49 million accurately measures the willingness to take risks at Superior Semiconductor, then the negative certainty equivalent that we found in cell B25 is an indication that we should not recommend the new product (because the alternative of not introducing it yields $0).

The cumulative risk profile in Figure 4.6 was made by sorting the simulated profit data in B28:B828 and plotting it in an *XY*-scatter chart against the percentile index in A28:A828. (Because of this sorting, Figure 4.6 shows only the worst profits that were sorted to the top of the data range.) Reading from this chart at the 0.05 cumulative probability level, we find a 5% value-at-risk of about –$23 million, which is substantially worse than the 5% value-at-risk that we found in our analysis of Part A (–$9.33 million). But the event of getting negative profit has probability about 0.53 here, which is less than its probability (0.65) in the model of Part A.

## 4.7 Certainty Equivalents of Normal Lotteries

We have seen in this chapter that Normality is a useful assumption in probability theory, and we saw in the previous chapter that constant risk tolerance is a useful assumption in utility theory. We now introduce a simple formula for computing certainty equivalents in situations where these two useful assumptions are both satisfied.

Consider a simple situation where a decision maker with constant risk tolerance $r$ has a lottery ticket that will pay a random amount of money $\mathbf{X}$ that is drawn from a Normal distribution with some mean $\mu$ and standard deviation $\sigma$. Then this decision maker's certainty equivalent of this lottery satisfies the following formula:

$$CE = \mu - 0.5*(\sigma\text{^}2)/r$$

This formula implies that the risk premium is

$$RP = \mu - CE = 0.5*(\sigma\text{^}2)/r$$

That is, the constant-risk-tolerant decision maker's risk premium for a Normal lottery is proportional to the lottery's variance ($\sigma$^2), where the constant of proportionality is $0.5/r$. (This formula also appeared in Section 3.3,

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | Constant risk tolerance | | | | | | | |
| 2 | 6000 | | | | | | | |
| 3 | | | | | | | | |
| 4 | Expected monetary value of a Normal lottery | | | | | | | |
| 5 | 9000 | | | | | | | |
| 6 | Std deviation | | FACT: | | | | | |
| 7 | 3000 | | If a lottery's net Payment ($) is drawn from a Normal distribution | | | | | |
| 8 | | | and the decision maker has constant risk tolerance, | | | | | |
| 9 | Certainty equiv | | then the DM's certainty equivalent of the lottery is | | | | | |
| 10 | 8250 | | CE = E($) - (0.5/RiskTol)*Stdev($)^2 | | | | | |
| 11 | | | | | | | | |
| 12 | Estimates from simulation | | | | | | | |
| 13 | CE | E($) | EU | CE from EU | | | | |
| 14 | 8177.37 | 8971.76 | -0.25592 | 8177.37 | | | | |
| 15 | | Stdev($) | Stdev(U) | | | | | |
| 16 | | 3135.07 | 0.136687 | | DataCount | 401 | | |
| 17 | | | | | 95% confidence interval | | | |
| 18 | | $Prize | Utility | | EU | -0.2693 | -0.24254 | |
| 19 | SimTable | 10232.49 | | | CE | 7871.64 | 8499.53 | |
| 20 | 0 | 6258.43 | -0.35237 | | | | | |
| 21 | 0.0025 | 10482.77 | -0.17427 | | FORMULAS FROM RANGE A1:G19 | | | |
| 22 | 0.005 | 6644.69 | -0.3304 | | A10. =A5-0.5*(A7^2)/A2 | | | |
| 23 | 0.0075 | 11308.80 | -0.15186 | | B19. =NORMINV(RAND(),A5,A7) | | | |
| 24 | 0.01 | 11507.81 | -0.14691 | | A14. =CE(B20:B420,A2) | | | |
| 25 | 0.0125 | 7968.94 | -0.26497 | | B14. =AVERAGE(B20:B420) | | | |
| 26 | 0.015 | 8136.84 | -0.25765 | | B16. =STDEV(B20:B420) | | | |
| 27 | 0.0175 | 8381.90 | -0.24734 | | C20. =UTIL(B20,$A$2) | | | |
| 28 | 0.02 | 7106.85 | -0.30591 | | C20 copied to C20:C420 | | | |
| 29 | 0.0225 | 8945.27 | -0.22517 | | C14. =AVERAGE(C20:C420) | | | |
| 30 | 0.025 | 11693.49 | -0.14243 | | D14. =UINV(C14,A2) | | | |
| 31 | 0.0275 | 6582.60 | -0.33384 | | C16. =STDEV(C20:C420) | | | |
| 32 | 0.03 | 8478.63 | -0.24339 | | F16. =COUNT(C20:C420) | | | |
| 33 | 0.0325 | 11168.13 | -0.15546 | | F18. =C14-1.96*C16/F16^0.5 | | | |
| 34 | 0.035 | 8254.67 | -0.25264 | | G18. =C14+1.96*C16/F16^0.5 | | | |
| 35 | 0.0375 | 11130.94 | -0.15643 | | F19. =UINV(F18,$A$2) | | | |
| 36 | 0.04 | 13282.94 | -0.10928 | | G19. =UINV(G18,$A$2) | | | |

FIGURE **4.7**
**Certainty equivalent of a Normal lottery**

where it was used to approximate the risk premiums of gambles that have small ranges of possible payoffs.)

This formula may also be written as $CE = \mu - (0.5*\sigma/r)*\sigma$, which indicates that $0.5*\sigma/r$ is the number of standard deviations by which the certainty equivalent is below the mean of the Normal lottery. Thus, the probability of this lottery paying less than its certainty equivalent is

$$P(\mathbf{X} < CE) = \text{NORMSDIST}(-0.5*\sigma/r)$$

So when the lottery's standard deviation $\sigma$ is much smaller than the decision maker's risk tolerance $r$, the certainty equivalent is very close to $\mu$, in the sense that the probability of the Normal lottery paying less than this certainty equivalent is almost 1/2. (Recall that $P(\mathbf{X} < \mu) = 1/2$ for a Normal random variable $\mathbf{X}$.) On the other hand, when the lottery's standard deviation $\sigma$ is much larger than the decision maker's risk tolerance $r$, then the certainty equivalent is very low in the Normal probability distribution, in the sense that the probability of the lottery paying less than this certainty

Myerson_C04.qxd 09-09-2004 11:55 Page 137

equivalent is small. (In practical applications, if σ is greater than 5∗*r*, then the probability of this certainty equivalent doing better than the lottery becomes so small that either the assumption of Normality or the assumption of constant risk tolerance should probably be called into question.)

The mathematical derivation of this formula for the certainty equivalents of Normal lotteries is a rather difficult calculus problem, but we can use simulation to demonstrate the validity of the formula, as shown in Figure 4.7. Cell A10 computes the certainty equivalent of a Normal lottery with mean \$9000 and standard deviation \$3000 for a decision maker with constant risk tolerance \$6000, using the formula

$$\mu - 0.5*(\sigma\text{^}2)/r = 9000 - 0.5*(3000\text{^}2)/6000 = \$8250$$

Cell B19 contains such a Normal random variable, generated by the formula NORMINV(RAND( ), 9000, 3000), and data from 401 independent simulations of B19 is stored in cells B20:B420 of this spreadsheet. Based on this simulation data, the formula *CE*(B20:B420, 6000) yields an estimated certainty equivalent \$8177, shown in cell A14. Cells F19 and G19 show a 95% confidence interval for the certainty equivalent based on this simulation data. Notice that this interval (from \$7871 to \$8499) contains the true certainty equivalent (\$8250) that we computed in cell A10.

## 4.8 Other Probability Distributions

In this section we describe some other probability distributions that are sometimes used in decision analysis models. These descriptions are intended mainly as a reference, and you may omit them at a first reading. But you may want to refer back to the appropriate subsection here when these special distributions are applied later.

### Uniform random variables

Given any two numbers *A* and *B* such that $A < B$, a continuous random variable has a *Uniform* distribution over the interval from *A* to *B* if the random variable is sure to be between the numbers *A* and *B* and the value of this random variable (rounded to any given number of decimal places) is equally likely to be any possible number between *A* and *B*. The probability density associated with this Uniform distribution is a flat constant $(1/(B-A))$ over the whole interval from the lower bound *A* and the upper bound *B*, and the probability density is 0 outside this interval.

The Excel function RAND( ) returns a random variable with a Uniform distribution over the interval from 0 to 1, which we use (with appropriate inverse cumulative functions) to create random variables with any other distribution. To make a Uniform random variable over some other

Myerson_C04.qxd 09-09-2004 11:55 Page 138

interval, from the lower bound $A$ to the upper bound $B$, we can use the linear formula

$$A + (B - A)*\text{RAND}()$$

For example, 95 + (105 – 95)*RAND( ) gives us a Uniform random variable over the interval from 95 to 105.

If **X** is a random variable with a Uniform distribution over the interval from $A$ to $B$, then **X** has quartile boundary points $Q_1 = (3*A + B)/4$, $Q_2 = (A + B)/2$, and $Q_3 = (A + 3*B)/4$. Its mean is $\mu = (A + B)/2$, and its standard deviation is approximately $\sigma = 0.2887*(B - A)$.

A Uniform random variable over the interval from 95 to 105 is (rounded to one decimal place) just as likely to be 104.9 as 97.2 or any other number in the interval, and yet 105.1 (just 0.2 more than 104.9) is completely impossible. Such a sudden break from "equally likely" to "impossible" seems unrealistic in most applications. Thus, there is rarely good reason to believe that any unknown quantity is well represented by a simple Uniform random variable.

## Triangular random variables

Like Uniform random variables (and Beta random variables, which are discussed in the next section), Triangular random variables are random variables that are guaranteed to stay within some bounded interval. A Triangular probability distribution is parameterized by its lower bound, by its mode (or "most likely" value), and by its upper bound. With Simtools, a random variable with a Triangular distribution, with lower bound $b_1$, mode $b_2$, and upper bound $b_3$, can be constructed by the formula

$$\text{TRIANINV}(\text{RAND}(), b_1, b_2, b_3)$$

The probability density for such a Triangular random variable has a graph that looks like a triangle over the interval from $b_1$ to $b_3$, with its peak at $b_2$. This simple triangular shape may be appealing to students, but statisticians consider such Triangular random variables to be rather crude objects. Still, this simplicity of Triangular random variables has made them popular for simulating unknown quantities where the decision maker is comfortable assessing a lower bound, a most-likely value, and an upper bound. We used such a Triangular distribution in Section 1.5.

## Beta random variables

Beta random variables are also bounded within a given interval, like Triangular and Uniform random variables, but Beta random variables are considered more sophisticated and more respectable to the refined tastes of statisticians.

For example, suppose that we are wondering about some unknown fraction, such as the fraction of MBA students who prefer to drink "Lite" beer.

Myerson_C04.qxd 09-09-2004 11:55 Page 139

This unknown quantity must be between 0 and 1 (that is, between 0% and 100%). So to describe our beliefs about this unknown fraction, we do not want to use a distribution that allows negative values or values that are greater than 1. The Beta distributions are such a family of distributions, yielding values that are always between 0 and 1. Beta distributions are often used by statisticians to describe beliefs about such unknown fractions.

The Beta distributions form a two-parameter family. Excel provides a function, BETAINV, that takes two parameters called $\alpha$ ("alpha") and $\beta$ ("beta"). The Excel formula

BETAINV(RAND( ), $\alpha$, $\beta$)

gives us a Beta random variable with parameters $\alpha$ and $\beta$, taking values between 0 and 1. The meaning of these parameters $\alpha$ and $\beta$ can be difficult to explain, so Simtools provides an alternative function called BETINV that is parameterized by the mean $\mu$ and standard deviation $\sigma$. That is,

BETINV(RAND( ), $\mu$, $\sigma$)

is a Beta random variable with mean $\mu$ and standard deviation $\sigma$, taking values between 0 and 1. The names of these two functions differ by just one letter, but you can remember that the Simtools function drops the letter "A" because it does not require you to understand a strange "Alpha" parameter. But an even better way to avoid confusion between these functions is to always use the Insert-Function dialog box to select either of them, because the dialog box will show you which function has "alpha" and "beta" as parameters and which has "mean" and "stdevn" as parameters. The relationship between the unintuitive $\alpha$ and $\beta$ parameters for this Beta distribution and the better-understood $\mu$ and $\sigma$ is given by the formulas

$$\mu = \alpha/(\alpha + \beta)$$

$$\sigma = (\alpha*\beta)\text{^}0.5/((\alpha + \beta) * (\alpha + \beta + 1)\text{^}0.5)$$

The Uniform distribution over the interval from 0 to 1 is actually a special case of a Beta distribution, using the special values $\alpha = 1$ and $\beta = 1$ (or $\mu = 0.5$ and $\sigma = 0.2887$).

The Beta distribution can be applied to any example in which we are wondering what fraction of a large population belongs to some subpopulation (such as the fraction of MBA students who drink Lite beer). Suppose that we began by thinking that this unknown fraction is drawn from a Uniform distribution over the interval from 0 to 1, but we then sampled $m$ members of this large population and found that $s$ of these members belong to the subpopulation. Then after observing the sample, our new probability beliefs about this unknown fraction should be a Beta distribution with parameters

$$\alpha = 1 + s \quad \text{and} \quad \beta = 1 + m - s$$

More generally, suppose that you thought that the unknown fraction of MBA students who prefer Lite beer has a Beta distribution with some

Myerson_C04.qxd 09-09-2004 11:55 Page 140

parameters α and β. Given these prior beliefs, you now learn that, among *m* randomly sampled students, *s* prefer to drink Lite beer. Then you should now assess a Beta distribution with parameters

$$\alpha' = \alpha + s \quad \text{and} \quad \beta' = \beta + m - s$$

for this unknown fraction.

One technical advantage of Beta distributions over Triangular distributions is that Beta distributions allow a wider range of standard deviations. A Triangular distribution with lower bound 0 and upper bound 1 could never have a standard deviation less than 0.2 or greater than 0.24. But a Beta distribution on the interval from 0 to 1 could have a standard deviation anywhere between 0 and 1/2. (More precisely, once the mean μ is specified, the standard deviation of a Beta distribution on the interval from 0 to 1 can be anywhere from 0 to (μ∗(1 – μ))^0.5.) We will see applications of Beta random variables in Chapter 6.

You can specify lower and upper bounds other than 0 or 1 by optional fourth and fifth parameters of both the BETINV and BETAINV functions. Thus, for example, BETINV(RAND( ), 24, 3, 20, 30) returns a Beta random variable that takes values between 20 and 30 and has mean 24 and standard deviation 3.

## Exponential random variables

Exponential random variables are often used to represent the length of time that we will have to wait for something to happen. Like the Lognormal random variables, an Exponential random variable is always a positive number. Unlike Normal and Lognormal random variables, which have the highest probability density at or near the mean, an Exponential random variable always has its highest probability density at 0.

If **X** is a random variable with an Exponential distribution, then, for any two nonnegative numbers *m* and *n*,

$$P(\mathbf{X} > m + n | \mathbf{X} > m) = P(\mathbf{X} > n)$$

To interpret this equation, suppose that **X** represents the number of minutes that we will have to wait for our first customer to arrive. This equation says that the conditional probability of having to wait more than *n* additional minutes, given that we have already waited any number of minutes (*m*), is independent of the number of minutes that we have already waited. This no-memory property is often quite natural to assume for waiting times. So the Exponential distribution is often used for random waiting times.

The Exponential distributions have one parameter: the mean. To construct an Exponential random variable with mean μ, you can use the Excel formula

$$-\mu*\text{LN}(1 - \text{RAND}( ))$$

or equivalently, with Simtools, you can use the formula

EXPOINV(RAND( ), μ)

This Exponential random variable has mean μ, and its standard deviation is also equal to μ. Its quartile points are 0.288*μ, 0.693*μ, and 1.386*μ. We will see applications of Exponential random variables in Chapter 9.

### Gamma random variables

The Gamma distribution is also used to represent beliefs about an unknown quantity that is a nonnegative number. The Gamma distribution has two parameters, often called α (alpha) and β (beta), which must also be nonnegative numbers. A Gamma random variable can be constructed in Excel by the formula

GAMMAINV(RAND( ), α, β).

Such a Gamma random variable has expected value (mean)

μ = α*β

and its standard deviation is

σ = (α^0.5)*β

Conversely, if we know the mean μ and standard deviation σ of a Gamma random variable, then its parameters α and β must be

α = (μ/σ)^2 and β = (σ^2)/μ

Thus, a Gamma random variable with mean μ and standard deviation σ can be constructed in Excel by the formula

GAMMAINV(RAND( ), (μ/σ)^2, (σ^2)/μ)

Simtools condenses this formula into one function called GAMINV. That is, the formula

GAMINV(RAND( ), μ ,σ)

returns a random variable that has a Gamma distribution with mean μ and standard deviation σ. To avoid confusion between these similarly named functions, you should always use the Insert-Function dialog box when you want to use either of them, because the dialog box will show you which function has "alpha" and "beta" as parameters, and which has "mean" and "stdevn" as parameters.

There are at least three common applications of Gamma distributions. The first is in statistics. Suppose that we have a range of *n* cells, in each of which the value is an independent Normal random variable with some mean μ and some standard deviation σ. Let **S** denote the value returned by Excel's STDEV function applied to this range. Then **S**^2 (the sample variance) is a Gamma random variable with mean σ^2 and standard deviation

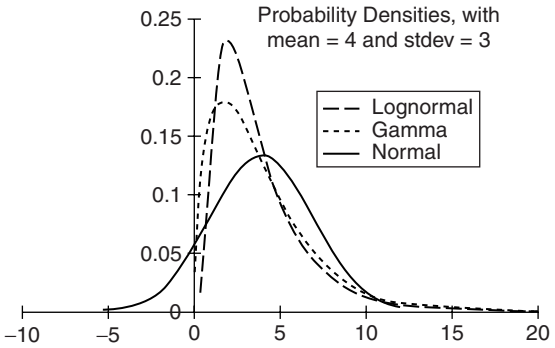


FIGURE **4.8**
**Comparison of Lognormal, Gamma, and Normal distributions**

(σ^2)*(2/(*n* – 1))^0.5. The sample standard deviation **S** itself is not a Gamma random variable, but when *n* is large its mean is close to σ and its standard deviation is close to σ/(2*(*n* – 1))^0.5. (So **S**/(2*(*n* – 1))^0.5 may be used as an estimate of the standard deviation of this statistic **S**.)

The second major application of Gamma distributions is in duration times. Recall that Exponential distributions may be applied to waiting times for the next customer's arrival. If $\mathbf{X}_1, \ldots, \mathbf{X}_n$ are $n$ independent random variables, each drawn from an Exponential distribution with mean τ, then the sum $\mathbf{X}_1 + \cdots + \mathbf{X}_n$ has a Gamma distribution with parameters $\alpha = n$ and $\beta = \tau$. So the time until the arrival of the $n$th customer may be a Gamma random variable. This fact implies that the Exponential distributions are actually a special subfamily of the Gamma distributions, being those for which the standard deviation is equal to the mean. We will see such applications of Gamma distributions in Chapter 9.

A third application of Gamma distributions arises when we have watched customers arriving at a store for a limited period of time and we are uncertain about the long-run mean arrival rate. Under some technical assumptions, it turns out to be natural to assume that our beliefs about this mean rate of arrival should be described by a Gamma distribution.

Gamma random variables are like a Lognormal random variable in that they can only take values greater than 0. But the Gamma distribution can have a positive probability density at 0, whereas the probability density of the Lognormal distribution cannot. In effect, Gamma random variables can come closer to 0 than Lognormals. Thus, the shape of the Gamma distribution often looks more similar to the Normal distribution with the same mean and standard deviation than does the corresponding Lognormal distribution. See, for example, Figure 4.8, which shows the probability densities for Gamma, Normal, and Lognormal distributions with mean 4 and standard deviation 3. When the mean is more than 5 times the standard deviation, the Gamma, Normal, and Lognormal distributions become very similar.

Myerson_C04.qxd 09-09-2004 11:55 Page 143

### Extreme-Value random variables

The Extreme-Value (or Gumbel) distribution arises when we look at the maximum of a large number of independent random variables that are drawn from any of a wide family of probability distributions, including the Normals. Simtools provides the function XTREMINV for making Extreme-Value random variables parameterized by their mean and standard deviation. With this function the formula

XTREMINV(RAND( ), μ, σ)

returns an Extreme-Value random variable with mean μ and standard deviation σ. Without Simtools, such a random variable can also be made in Excel by the formula

μ – σ*(0.45 + 0.78*LN(–LN(RAND( )))

Extreme-Value random variables can be positive or negative. But unlike the Normals, the Extreme-Value distributions are not symmetric about their means, but instead are skewed with a longer tail of likely values on the high side. XTREMINV allows you to reverse the direction of skewness by entering a negative value as the standard devation parameter. Such reverse-skewed distributions can arise when we look at the minimum of a large number of independent random variables from the same distribution.

A *Weibull random variable* with log-mean *m* and log-standard-deviation *s* is returned by the formula EXP(XTREMINV(RAND( ), *m*, –*s*)). Notice the negative sign on *s* to reverse skewness. Weibull random variables are often used as models of a device's operating lifetime.

### A nonparametric distribution fitting sample data

Probability distributions can also be constructed to more closely fit any given sample of data, if we are willing to go beyond simple families of distributions that have only a few parameters. Consider a situation where we have observed a sample of values that we believe to have been drawn independently from some fixed probability distribution, but we do not know anything else about this probability distribution. Suppose that this sample data is listed in a range of spreadsheet cells called *sampledata*. Then the Excel formula

PERCENTILE(*sampledata*, RAND( ))

returns a random variable that has a cumulative probability distribution that matches the cumulative percentiles in our sample data. In effect, this PERCENTILE formula tells Excel to sort the sample data from lowest to highest, then to randomly choose a pair of data values that are adjacent in this ordering, and then to return a number that is randomly drawn from a Uniform distribution on the interval between these two adjacent data

Myerson_C04.qxd 09-09-2004 11:55 Page 144

values. If the values in our sample are all different from each other, then this random variable is continuous. For each value $x$ in our sample data, the probability of this random variable being less than $x$ is exactly the fraction of other values in our sample that are less than $x$. But this random variable will never be lower than the lowest value in our sample or higher than the highest value in our sample. Without some stronger assumption about the probability distribution, we have no sensible way of assigning probabilities to values that are higher or lower than anything we have observed.

In the spreadsheet of Figure 4.4, for example, we could make a random variable that has a cumulative distribution matching the observed percentiles of the first strategy's outcomes by the formula PERCENTILE(F25:F525, RAND( )). You can compare this nonparametric random variable to the Normal or Generalized-Lognormal random variables that we would fit to the same data here with the formulas NORMINV(RAND( ), F17, F18) or GENLINV(RAND( ), F19, F20, F21), and you should find that these three random variables generate similar distributions of outcomes. Such nonparametric random variables will be applied in Section 7.5.

If *sampledata* is a range of $n$ cells in one column, then we could also select a random entry from the *sampledata* range by the Excel formula

INDEX(*sampledata*, INT($n$∗RAND( ) +1))

Unlike the PERCENTILE formula just described, this simple resampling technique will only repeat the values that actually appear in the given *sampledata*.

In addition to the families of continuous probability distributions that we have discussed, there are two other *families of discrete probability distributions* that are worth describing here.

## Binomial random variables

Suppose that we have an experiment whose outcome may be either a "success" or a "failure." Suppose we perform this experiment $n$ times independently, and let $p$ denote the probability of a success each time we perform this experiment. Then the number of successes is a Binomial random variable with parameters $n$ and $p$.

As we saw in Chapter 1, a Binomial random variable with parameters $n$ and $p$ can be simulated in a spreadsheet by taking the sum of $n$ cells, each of which has a random value generated by the formula

IF(RAND( ) < $p$, 1, 0)

Using Simtools, we can also simulate a Binomial random variable with parameters $n$ and $p$ by the formula

BINOMINV(RAND( ), $n$, $p$)

A Binomial random variable with parameters $n$ and $p$ has a mean equal to $n$∗$p$, and it has a standard deviation equal to ($n$∗$p$∗(1 – $p$))^0.5. If this standard deviation is larger than 3, then the Binomial random variable can also

Myerson_C04.qxd 09-09-2004 11:55 Page 145

be well approximated by a Normal random variable with the same mean and standard deviation. If you want make sure that such a Normal approximation to a Binomial is integer valued, then you can use the formula

ROUND(NORMINV(RAND( ), $n*p$, $(n*p*(1-p))$^0.5), 0)

(The Excel formula ROUND($x$, 0) rounds any number $x$ to the nearest integer.)

### Poisson random variables

The Poisson distribution is commonly used to simulate the number of times that something will happen when there is no clear upper bound on how many times it might happen. For example, the number of customers who will make purchases in a particular shop tomorrow could be described by a random variable with a Poisson distribution.

Suppose that the number of minutes from one customer's arrival until the next customer arrives is a random variable drawn from an Exponential distribution with mean $\tau$. Then the number of customers who arrive in any period of $L$ minutes is a Poisson random variable, and its mean is $L/\tau$. Thus, Poisson distributions are often used to represent our beliefs about the number of arrivals that will occur in a given period.

Poisson random variables always take values that are nonnegative integers. The Poisson distribution has one parameter: the mean. With Simtools, a Poisson random variable with mean $\mu$ can be simulated by the formula

POISINV(RAND( ), $\mu$).

A Poisson random variable with mean $\mu$ has a standard deviation that is equal to $\mu$^0.5. If $\mu$ is larger than 20, then a Poisson random variable with mean $\mu$ can also be well approximated by a normal random variable with mean $\mu$ and standard deviation $\mu$^0.5. That is, the formula

ROUND(NORMINV(RAND( ), $\mu$, $\mu$^0.5), 0)

gives values that are very similar to POISINV(RAND( ), $\mu$), when $\mu > 20$. A Poisson random variable with mean $\mu$ is also a good approximation for a Binomial random variable with parameters $n$ and $p$ if $n*p = \mu$ and $p$ is very small (say $p < 0.1$).

Appications of Poisson random variables will be discussed further in Chapter 9.

## 4.9 Summary

In this chapter we studied families of continuous probability distributions that are commonly used to describe uncertainty about unknown quantities which can take infinitely many possible values.

Myerson_C04.qxd 09-09-2004 11:55 Page 146

The Normal probability distributions are generally parameterized by their mean and standard deviation and can be applied to unknown quantities (like sample means) that can be expressed as the additive sum of many small terms drawn independently from some distribution. We saw how to compute the probability of a Normal random variable being in any given interval by using Excel's NORMSDIST function. The Excel function NORMINV is used to simulate Normal random variables and to plot their inverse cumulative curves. After introducing the concept of a probability density curve for a continous random variable, we showed how to estimate it from an inverse cumulative function like NORMINV.

Before introducting the Lognormals, we reviewed the exponential functions EXP and the natural logarithm function LN, showing how they arise when interest rates are compounded frequently. The advantages of using logarithmic rates of return were discussed. We also reviewed how EXP converts addition to multiplication, and how the inverse function LN converts multiplication to addition.

A Lognormal random variable is one whose natural logarithm is a Normal random variable. Lognormal random variables can be parameterized by their log-mean $m$ and log-standard-deviation $s$ or by their true mean $\mu$ and true standard deviation $\sigma$. With the first parameterization, they can be simulated by the Excel formula EXP(NORMINV(RAND( ), $m$, $s$)). With the second parameterization, they can be simulated by the Simtools formula LNORMINV(RAND( ), $\mu$, $\sigma$). Lognormal distributions can be applied to unknown quantities (such as growth ratios over some period of time) that can be expressed as the multiplicative product of many factors, each close to 1, which are drawn independently from some distribution.

The Generalized-Lognormals were introduced as a family of probability distributions that can be parameterized by the three quartile boundary points and that includes as special cases both the Normals (when the quartile points have equal differences) and Lognormals (when the quartile points have equal ratios). The Simtools function GENLINV was introduced to simulate Generalized-Lognormal random variables. Practical techniques were illustrated for subjectively assessing quartiles and fitting a Generalized-Lognormal distribution that can describe an individual's beliefs about a real unknown quantity for which there is no comparable statistical data. We then applied these techniques with simulation analysis to analyze a decision problem that involves several different unknown quantities, some continuous and some discrete.

A simple formula was introduced for computing certainty equivalents in the special case where the risky monetary payoff will be drawn from a Normal distribution and the decision maker has constant risk tolerance.

Finally, other families of probability distributions were briefly introduced and discussed.

Myerson_C04.qxd 09-09-2004 11:55 Page 147

## CASE FOR EXERCISE 1 ■ Lepton Inc.

Lepton Inc. is a bioengineering firm founded five years ago by a team of biologists. A venture capitalist has funded Lepton's research during most of this time in exchange for a controlling interest in the firm. Lepton's research during this period has been focused on the problem of developing a bacterium that can synthesize HP-273, a protein with great pharmacological and medical potential. This year, using new methods of genetic engineering, Lepton has at last succeeded in creating a bacterium that produces synthetic HP-273 protein.

Lepton can now get a patent for this bacterium, but the patented bacterium will be worthless unless the FDA approves the synthetic protein that it produces for pharmacological use. The process of getting FDA approval has two stages. The first stage involves setting up a pilot plant and testing the synthetic protein on animal subjects to get the preliminary certification of safety that is required before any human testing can be done. The second stage involves an extensive medical study for safety and effectiveness using human subjects.

Lepton's engineers believe that the cost of the first stage could equally likely be above or below $9 million, has probability 0.25 of being below $7 million, and has probability 0.25 of being above $11 million. The probability of successfully getting a preliminary certification of safety at the end of this first stage is 0.3.

The cost of taking the synthetic protein though the second-stage extensive medical study has a net present value that is equally likely to be above or below $20 million, has probability 0.25 of being below $16 million, and has probability 0.25 of being above $24 million. Given first-stage certification, if Lepton takes the synthetic HP-273 protein through this second stage, then the conditional probability of getting full FDA approval for its pharmacological use is 0.6.

If Lepton gets full FDA approval for pharmacological use of the synthetic HP-273 protein, then Lepton will sell its patent for producing the protein to a major pharmaceutical company. Given full FDA approval, the returns to Lepton from selling its patent would have a present value that is equally likely to be above or below $120 million, has probability 0.25 of being below $90 million, and has probability 0.25 of being above $160 million.

Lepton's management has been operating on the assumption that they would be solely responsible for the development of the protein throughout this process, which we may call Plan A. But there are two other alternatives that the owner wants to consider:

Plan B: Lepton has received an offer to sell its HP-273 patent for $4 million now, before undertaking any costs in the first stage of the approval process.

Plan C: If Lepton succeeds in getting preliminary FDA certification at the end of the first stage, then it could sell a 50% share of the HP-273 project to another bioengineering firm for $25 million. Lepton would then pay only half of the second-stage costs, but it would get only half of the income (if any) from selling the patent.

In your analysis of this situation, you may assume that each of the three unknown quantities described (the stage 1 cost, the stage 2 cost after a success at stage 1, and the value of the patent after a success at stage 2) has a probability distribution in the Generalized-Lognormal family.

Myerson_C04.qxd 09-09-2004 11:55 Page 148

## CASE FOR EXERCISE 5 ■ Global Sports

Century Publications is considering a proposal to launch a new magazine called *Global Sports*, which would be devoted to international sports, with features on competitive sports in different foreign countries that are not familiar to most Americans. The fixed costs of producing the magazine would be \$4 million in the first year, \$3 million in the second year, and then \$2 million per year thereafter. In addition, the variable costs of printing and distributing the magazine would be \$18 per subscriber per year. Subscriptions to the magazine would be sold for \$25 per year. Advertising revenue will depend on the number of subscriptions and is anticipated to be \$10 per subscriber each year.

Senior managers in Century's circulation department have tried to formulate some forecast of the potential for this new magazine, but predicting the number of subscribers for any new publication is always difficult. Whatever the potential market might be, subscriptions in the early years would almost certainly be substantially less than the long-run level. From past experience with other magazines, we can estimate that the logarithmic growth rate of subscriptions per year (defined in any year $t$ to be the LN($s_t/s_{t-1}$), where $s_t$ is subscriptions in year $t$, averaged over the year, and $s_{t-1}$ is subscriptions in the previous year $t-1$) have the statistics shown below:

| **Year *t* (from launch)** | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|
| ***E*(LN($s_t/s_{t-1}$))** | 0.5 | 0.3 | 0.2 | 0.1 | 0 | 0 | 0 | 0 | 0 |
| **Stdev(LN($s_t/s_{t-1}$))** | 0.3 | 0.2 | 0.15 | 0.07 | 0.04 | 0.04 | 0.04 | 0.04 | 0.04 |

The circulation managers felt comfortable assuming that growth of *Global Sports* would have statistics similar to these, with subscriptions climbing at first and then leveling off at some plateau level after five years. The question remained, How high would this plateau be? On this important question, there was great uncertainty, as it would depend on the unique and unproven appeal of this new magazine. Senior managers at Century estimated that circulation in year 5 after launch would be equally likely to be above or below 100,000. Working a bit harder to estimate upper and lower quartile boundaries for year-5 circulation, the circulation managers reached a consensus on 50,000 and 200,000. That is, they assessed a 25% probability of subscriptions in year 5 being above 200,000, and they assessed a 25% probability of subscriptions in year 5 being below 50,000.

In Century's analysis of the alternatives in this situation, the bottom-line quantity of interest to senior management is the net present value of profits from the magazine over the next 10 years, using a discount rate of 10% per year. In general, if the stream of profits during the next $n$ years will be $v_1, v_2, \ldots, v_n$, then the *net present value* of this profit stream at a discount rate $d$ is defined by the formula

$$\text{NPV}(d, v_1, \ldots, v_n) = \sum_{t=1}^{n} v_t/(1+d)\text{^}t$$

In this case, we want to evaluate the net present value of profits over $n = 10$ years, using a discount rate of $d = 0.10$. Excel has a built-in NPV function that we can use to do this calculation with the formula

NPV(0.10, *profitstream*)

where *profitstream* denotes a range of 10 cells (in a row, or in a column) that list the simulated profits over the next 10 years (with the first year's profit at the left or top of the range). For a fuller explanation of net present value, see Section 9.1.

Myerson_C04.qxd 09-09-2004 11:55 Page 149

## Exercises

**1.** Estimate the *expected value* and *standard deviation* of the current owner's profits from Lepton under Plan A and Plan C (see case on p. 147). Then make a *cumulative risk profile* for each plan. If the current owner of Lepton wants to maximize the expected value of profits, then what plan should we recommend and what is the *lowest price* at which Lepton should consider selling its HP-273 patent now?

Base your answers on simulation data that include enough so that the radius of your 95% confidence interval for the expected value of profit is less than \$2 million for both Plan A and Plan C. (That is, you should be able to generate 95% confidence intervals of the form $x \pm y$, where $y$ is less than \$2 million.) How many simulations did you use? What is your 95% confidence interval for expected profit under each plan?

**2.** **(a)** Among the three unknown quantities described in the Lepton case (the stage 1 cost, the stage 2 cost, and the value of the patent), which have Normal distributions? Compute the mean and the standard deviation for each Normal random variable.

**(b)** Among the three unknown quantities described in the Lepton case, which have Lognormal distributions? Compute the log-mean and the log-standard-deviation for each Lognormal random variable.

**(c)** Compute the probability that the stage 1 cost will be between \$8 and \$12 million. Also compute the probability that it will be more than \$12 million.

**3.** The venture capitalist who owns Lepton Inc. has said that a lottery paying either \$100 million or \$0, each with probability 1/2, would be worth just \$42 million to him. Assume that he has constant risk tolerance.

**(a)** What is his risk-tolerance constant?

**(b)** Based on your simulation data from Exercise 1, estimate his certainty equivalent for the HP-273 patent under Plan A and under Plan C. Which plan looks best for him, with such risk tolerance?

**(c)** What would be the highest risk tolerance for this decision maker such that you would recommend Plan B, selling the HP-273 patent now for \$4 million? At this level of risk tolerance, what would be his certainty equivalent for a lottery paying either \$100 million or \$0, each with probability 1/2?

**4.** Consider the following five unknown quantities:

**(a)** The number on the last numbered page of this book. (No peeking!)

**(b)** Age of Columbus when he died.

**(c)** Distance from Chicago to Seattle in miles.

**(d)** Melting point of aluminum (specify °F or °C).

**(e)** Population of Italy on January 1, 1973.

For each of these unknown quantities, assess subjective quartile boundary points $Q_1$, $Q_2$, $Q_3$, such that $Q_1 < Q_2 < Q_3$ and you would be indifferent among the following four lotteries:

*Lottery 1:* You get \$5000 if the unknown is less than $Q_1$, but get \$0 otherwise.

*Lottery 2:* You get \$5000 if the unknown is between $Q_1$ and $Q_2$, but get \$0 otherwise.

*Lottery 3:* You get \$5000 if the unknown is between $Q_2$ and $Q_3$, but get \$0 otherwise.

*Lottery 4:* You get \$5000 if the unknown is greater than $Q_3$, but get \$0 otherwise.

Then compute the points with 1% and 99% cumulative probability in the corresponding Generalized-Lognormal distribution by the formulas GENLINV(0.01, $Q_1$, $Q_2$, $Q_3$) and GENLINV(0.99, $Q_1$, $Q_2$, $Q_3$).

Check whether this Generalized-Lognormal distribution is a good fit to your actual beliefs by asking whether you think that the unknown has a 1/100 probability of being less than GENLINV(0.01, $Q_1$, $Q_2$, $Q_3$), and whether you think that the unknown has a 1/100 probability of being greater than GENLINV(0.99, $Q_1$, $Q_2$, $Q_3$). If not, then try to reassess the three quartile points until you also get a reasonable 1% point and 99% point from these GENLINV formulas.

Suggestions:

If GENLINV(0.01, $Q_1$, $Q_2$, $Q_3$) seems too low, then try increasing $Q_1$ or decreasing $Q_2$.

If GENLINV(0.01, $Q_1$, $Q_2$, $Q_3$) seems too high, then try decreasing $Q_1$ or increasing $Q_2$.

If GENLINV(0.99, $Q_1$, $Q_2$, $Q_3$) seems too low, then try increasing $Q_3$ or decreasing $Q_2$.

If GENLINV(0.99, $Q_1$, $Q_2$, $Q_3$) seems too high, then try decreasing $Q_3$ or increasing $Q_2$.

**5.** In the Global Sports case on page 148,

Myerson_C04.qxd 09-09-2004 11:55 Page 150

**(a)** Make a simulation model of this situation, use it to estimate the expected net present value of profits from the proposed magazine over the next 10 years, and show the cumulative distribution of this net present value, under the assumption that Century will publish this magazine for all of the 10 years regardless of its performance.

**(b)** Consider an alternative plan in which, after publishing the magazine for one year, the magazine will be terminated if fewer than 35,000 subscriptions are sold in the first year (in which case profits will go to $0 in all subsequent years). But if 35,000 or more subscriptions are sold in the first year, then the magazine will be continued, as in part (a), for the rest of the decade. Based on simulation data from your model, estimate the expected net present value of profits under this plan, and show the cumulative distribution of the net present value under this plan.

Myerson_C05.qxd 09-09-2004 10:34 Page 151

# 5

# Correlation and Multivariate Normal Random Variables

Recall that *independence* of two unknown quantities means that learning the value of one quantity would not change our beliefs about the value of the other quantity. It is easy to make independent random variables in Excel, but unknown quantities that arise in applied problems often do not satisfy independence. For example, the future values of stock market indexes like the DJIA and the S&P 500 at the end of next year are unknown quantities today, and we do not know whether they will be higher or lower than their current values. But if we learned that the DJIA index would increase by more than 30% between now and the end of next year, then we would consider it extremely unlikely that the S&P 500 would decrease between now and the end of next year. That is, learning the future DJIA would strongly affect our beliefs about the future S&P 500. So it would be a serious mistake to make a model in which the DJIA and S&P 500 were represented by independent random variables because such independence would be incompatible with the strong relationship between these two unknown quantities that we see in the real world.

Chapter 2 presented the expected value and the standard deviation as commonly used summary measures of any single random variable's probability distribution. But we must now focus more on the relationships among groups of random variables, and so this chapter introduces the mathematical concept of *correlation*, which is the most commonly used measure of relationships among pairs of random variables. We will see how to measure correlations from a joint probability distribution, from data, and from subjective assessments.

Chapter 4 presented the Normal distributions as the most commonly used family of probability distributions for modeling beliefs about a single random variable. But to discuss relationships between two or more random variables, we need to talk about their *joint probability distribution*. The most commonly used family of joint probability distributions for modeling beliefs about several nonindependent random variables is the *Multivariate Normal* family, which is a generalization of the Normal family for groups of

Myerson_C05.qxd 09-09-2004 10:34 Page 152

two or more random variables. These Multivariate Normal distributions are jointly parameterized by their means, standard deviations, and correlations. So we will also see how to work with Multivariate Normal random variables in this chapter, with an emphasis on applications to the analysis of financial portfolios. Two important Excel topics are also introduced in this chapter: the Solver optimization tool and array formulas that are entered into whole ranges of spreadsheet cells.

The analysis of relationships among random variables may be the most important topic in probability modeling, and there is more to say here than can fit into one chapter. So the discussion of other ways to measure relationships among random variables will be continued in Chapter 6, where the focus will be on conditional expectation and linear regression models.

## 5.1 Joint Distributions of Discrete Random Variables

Consider a young investor who owns some stock shares in two companies. She plans to sell these shares next year for a down payment on a condominium. The price per share at which she will be able to sell these stocks next year is now an unknown quantity. So she begins to think about probability distributions for these unknown quantities. For notation, let **X** denote next year's price per share of the first company that she has invested in, and let **Y** denote next year's price per share of the second company. Suppose that she assesses the following simple discrete probability distributions for these two unknown quantities:

| $x$ | $P(\mathbf{X} = x)$ | $y$ | $P(\mathbf{Y} = y)$ |
|---|---|---|---|
| 70 | 0.2 | 30 | 0.3 |
| 75 | 0.5 | 32 | 0.6 |
| 80 | 0.3 | 34 | 0.1 |

Let **W** denote the total value that she will get from selling her portfolio next year. Of course, **W** is also unknown now because it depends on these unknown prices **X** and **Y**. To be specific, suppose that she owns 100 shares of the stock that will be worth **X** and she owns 200 shares of the stock that will be worth **Y**. Then the total value of her portfolio **W** is given by the formula

$$\mathbf{W} = 100*\mathbf{X} + 200*\mathbf{Y}$$

With this formula and the investor's assessed probability distribution for **X** and **Y**, can we compute the probability distribution for **W**? Unfortunately, the answer is no. For example, try to compute the probability that her portfolio will be worth **W** = \$13,000, which (with the possible values of **X** and **Y** listed in the tables) can occur only if **X** = 70 and **Y** = 30. We know that $P(\mathbf{X} = 70)$ is 0.2 and $P(\mathbf{Y} = 30)$ is 0.3, but what is the probability of getting

Myerson_C05.qxd 09-09-2004 10:34 Page 153

$\mathbf{X} = 70$ and $\mathbf{Y} = 30$ at the same time? The answer to that question depends on what the investor believes about the relationship between **X** and **Y**. If she believes that **X** tends to be low when **Y** is also low, then the probability of getting both $\mathbf{X} = 70$ and $\mathbf{Y} = 30$ could be as large as 0.2. But if she believes that **X** will be low only when **Y** is not low, then the probability of getting both $\mathbf{X} = 70$ and $\mathbf{Y} = 30$ could be 0.

The symbol "$\cap$" is used to denote the *intersection* of two events, that is, the event that two other events both occur. So the expression "$\mathbf{X} = 70 \cap \mathbf{Y} = 30$" denotes the event that **X** is 70 and **Y** is 30 (in which case, $\mathbf{W} = 100*70 + 200*30 = 13{,}000$). With this notation, we may say that our problem is that knowing $P(\mathbf{X} = 70)$ and $P(\mathbf{Y} = 30)$ does not tell us enough to compute $P(\mathbf{X} = 70 \cap \mathbf{Y} = 30)$.

A table or function that specifies, for each pair of possible values $(x, y)$, the probability $P(\mathbf{X} = x \cap \mathbf{Y} = y)$ is called the *joint probability distribution* of **X** and **Y**. To compute the probability distribution, we need to assess the joint probability distribution of **X** and **Y**.

When **X** and **Y** are *independent* random variables, then their joint probability distribution can be computed by the simple formula

$$P(\mathbf{X} = x \cap \mathbf{Y} = y) = P(\mathbf{X} = x)*P(\mathbf{Y} = y)$$

So for this example, if **X** and **Y** are independent, then $P(\mathbf{X} = 70 \cap \mathbf{Y} = 30)$ is $0.2*0.3 = 0.06$. With this multiplicative formula, specifying the separate distributions of independent random variables implicitly gives us their joint distribution. That is why, in Chapter 4, we did not worry about joint probability distributions of the multiple unknown quantities that arose in Part C of the Superior Semiconductor case, because we were assuming that these unknown quantities were all independent of each other.

But stock prices often move up and down together, following general market trends. Thus, there is good reason to suppose that our investor might not think of **X** and **Y** as independent random variables. The discrete distributions shown allow nine possible pairs of values for **X** and **Y** together:

$$\{(70, 30), (70, 32), (70, 34), (75, 30), (75, 32), (75, 34), (80, 30),$$
$$(80, 32), (80, 34)\}$$

and we may ask our investor to subjectively assess a probability for each pair. She may have some difficulty assessing these nine probabilities (which must sum to 1 because exactly one pair will occur), but let us suppose that her answers are as summarized in the joint probability table shown in the range A3:D6 of Figure 5.1.

The separate probability distributions for **X** and **Y** can be computed from this joint probability distribution by summing across each row (to get the probability distribution of **X**) or down each column (to get the probability distribution of **Y**). These computations are performed in cells F4:F6 and B8:D8 of Figure 5.1, where they yield the probability distributions for **X** and **Y** that we have described, and so this joint distribution is consistent with the earlier assessment of the separate probability distributions of **X**

Myerson_C05.qxd 09-09-2004 10:34 Page 154

| | A | B | C | D | E | F |
|---|---|---|---|---|---|---|
| 1 | JointProbys P(X=x∩Y=y) | | | | | |
| 2 | | y= | | | | |
| 3 | x= \ | 30 | 32 | 34 | | P(X=x) |
| 4 | 70 | 0.1 | 0.1 | 0 | | 0.2 |
| 5 | 75 | 0.2 | 0.3 | 0 | | 0.5 |
| 6 | 80 | 0 | 0.2 | 0.1 | | 0.3 |
| 7 | | P(Y=y) | | | | sum |
| 8 | | 0.3 | 0.6 | 0.1 | | 1 |
| 9 | | | | | | |
| 10 | X | Y | | | | |
| 11 | 75.5 | 31.6 | Expected Value | | | |
| 12 | 3.5 | 1.2 | Standard deviation | | | |
| 13 | | | | | | |
| 14 | FORMULAS FROM RANGE A1:F12 | | | | | |
| 15 | F4. =SUM(B4:D4) | | | | | |
| 16 | F4 copied to F4:F6. | | | | | |
| 17 | B8. =SUM(B4:B6) | | | | | |
| 18 | B8 copied to B8:D8. | | | | | |
| 19 | F8. =SUM(B4:D6) | | | | | |
| 20 | A11. =SUMPRODUCT(A4:A6,F4:F6) | | | | | |
| 21 | B11. =SUMPRODUCT(B3:D3,B8:D8) | | | | | |
| 22 | A12. =STDEVPR(A4:A6,F4:F6) | | | | | |
| 23 | B12. =STDEVPR(B3:D3,B8:D8) | | | | | |

FIGURE **5.1**
**Joint probability distribution of two random variables**

and **Y**. The probability distribution of a single random variable is sometimes called its *marginal* probability distribution when it is computed from a joint distribution in this way (so called simply because it is often displayed next to the margin of the joint probability table, as in Figure 5.1).

In cells B4:D6 of Figure 5.1, notice that the probability numbers in the top-left and bottom-right corners, where the unknowns are either both greater than their expected values or both less than their expected values, are greater than the probability numbers in the top-right and bottom-left corners, where one unknown is greater than expected and the other is less than expected. Thus, the joint probability table in this example exhibits a kind of positive relationship between **X** and **Y**, as stock prices that tend to go up or down together. To provide a quantitative measure of such relationships or tendencies of random variables to covary, mathematicians have defined two concepts called *covariance* and *correlation*.

But before introducing these concepts, it will be helpful to switch to a different way of representing joint probability distributions, because the two-dimensional table shown in Figure 5.1 is not so convenient for computations in spreadsheets. A more convenient (if less intuitive) way of displaying a joint probability distribution is to list each possible combination of values of **X** and **Y** in a separate row, with all the joint probabilities listed down one column. Such a one-dimensional joint probability table is shown in the range A4:C13 of Figure 5.2. The nine possible pairs of values of **X** and **Y** are listed in rows 5 through 13, with **X** values in A5:A13 and **Y** values in B5:B13, and the probabilities for these (**X**,**Y**) pairs are listed in cells C5:C13. (A Simtools procedure called Combine Rows can be used for automatically generating all possible combinations of two lists for a range like A5:B13

Myerson_C05.qxd 09-09-2004 10:34 Page 155

| | A | B | C | D | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | | | | | | | Portfolio shares | | | | |
| 2 | | | | | | | of X | of Y | | | |
| 3 | | | JointProbys | | | | 100 | 200 | | | |
| 4 | x | y | P(X=x∩Y=y) | | ProdDevsFromMean | | | | W = Portfolio value | | |
| 5 | 70 | 30 | 0.1 | | 8.8 | | | | 13000 | | |
| 6 | 70 | 32 | 0.1 | | -2.2 | | | | 13400 | | |
| 7 | 70 | 34 | 0 | | -13.2 | | | | 13800 | | |
| 8 | 75 | 30 | 0.2 | | 0.8 | | | | 13500 | | |
| 9 | 75 | 32 | 0.3 | | -0.2 | | | | 13900 | | |
| 10 | 75 | 34 | 0 | | -1.2 | | | | 14300 | | |
| 11 | 80 | 30 | 0 | | -7.2 | | | | 14000 | | |
| 12 | 80 | 32 | 0.2 | | 1.8 | | | | 14400 | | |
| 13 | 80 | 34 | 0.1 | sum | 10.8 | | | | 14800 | | |
| 14 | | | | 1 | | | | | | | |
| 15 | X | Y | | Covariance | | Correlation | | | W | =100*X+200*Y | |
| 16 | 75.5 | 31.6 | E | | 2.2 | 0.524 | | | 13870 | E | 13870 |
| 17 | 3.5 | 1.2 | Stdev | | 2.2 | 0.524 | | | 517.78 | Stdev | 517.78 |
| 18 | | | | | | | | | 268100 | Var | 268100 |
| 19 | PRODS(Stdevs) | | | CorrelArray | | | PRODS(Shares) | | | | |
| 20 | 12.25 | 4.2 | | 1 | 0.524 | | 10000 | 20000 | | | |
| 21 | 4.2 | 1.44 | | 0.524 | 1 | | 20000 | 40000 | | | |
| 22 | | | | | | | | | | | |
| 23 | | | | CovarArray | | | | | Simulation model | | |
| 24 | | | | 12.25 | 2.2 | | | | 0.09 | (rand) | |
| 25 | | | | 2.2 | 1.44 | | | | X | Y | W |
| 26 | | | | | | | | | 70 | 30 | 13000 |
| 27 | FORMULAS | | | | | | | | | | |
| 28 | E5. =(A5-$A$16)*(B5-$B$16) | | | | | | | E5 copied to E5:E13 | | | |
| 29 | I5. =SUMPRODUCT($G$3:$H$3,A5:B5) | | | | | | | I5 copied to I5:I13 | | | |
| 30 | D14. =SUM(C5:C13) | | | | | | | | | | |
| 31 | A16. =SUMPRODUCT(A5:A13,$C$5:$C$13) | | | | | | | A16 copied to B16,E16,I16 | | | |
| 32 | A17. =STDEVPR(A5:A13,$C$5:$C$13) | | | | | | | A17 copied to B17,I17 | | | |
| 33 | F16. =E16/(A17*B17) | | | | | | | | | | |
| 34 | K16. =SUMPRODUCT(G3:H3,A16:B16) | | | | | | | | | | |
| 35 | E17. =COVARPR(A5:A13,B5:B13,C5:C13) | | | | | | | | | | |
| 36 | F17. =CORRELPR(A5:A13,B5:B13,C5:C13) | | | | | | | | | | |
| 37 | I18. =I17^2 | | | | | | | | | | |
| 38 | A20:B21. {=PRODS(A17:B17)} | | | | | | | | | | |
| 39 | E20,D21. =F16 | | | | | | | | | | |
| 40 | G20:H21. {=PRODS(G3:H3)} | | | | | | | | | | |
| 41 | D24. =A20*D20 | | | | | | | D24 copied to D24:E25. | | | |
| 42 | K18. =SUMPRODUCT(G20:H21,A20:B21,D20:E21) | | | | | | | | | | |
| 43 | K17. =SUMPRODUCT(PRODS(G3:H3),PRODS(A17:B17),D20:E21)^0.5 | | | | | | | | | | |
| 44 | I24. =RAND() | | | | | | | | | | |
| 45 | I26. =DISCRINV($I$24,A5:A13,$C$5:$C$13) | | | | | | | | | | |
| 46 | J26. =DISCRINV($I$24,B5:B13,$C$5:$C$13) | | | | | | | | | | |
| 47 | K26. =SUMPRODUCT(G3:H3,I26:J26) | | | | | | | | | | |

FIGURE **5.2**
**Correlation of two discrete random variables and portfolio analysis**

here, but it is also easy to just type them in.) You should verify that the joint probabilities shown in cells C5:C13 of Figure 5.2 are really the same as those shown in cells B4:D6 of Figure 5.1.

Cells I26:J26 in Figure 5.2 show how to make a simulation of **X** and **Y** with this joint probability distribution. The formulas in I26 and J26 are, respectively,

=DISCRINV($I$24,A5:A13,$C$5:$C$13)

=DISCRINV($I$24,B5:B13,$C$5:$C$13)

Myerson_C05.qxd 09-09-2004 10:34 Page 156

where cell I24 contains the formula =RAND( ). Because the same RAND( ) in I24 is used in both of these formulas, the two DISCRINV functions select the same row in the C5:C13 probabilities (the lowest row above which the sum of probabilities is less than I24) and then return the values for **X** and **Y** from the A and B cells in that row. It is important that we used the same RAND( ) in both formulas. If we had used different RANDs (instead of I24) in these formulas, then cells I25 and I26 would have been independent random variables that could have taken the values 70 and 34 at the same time, even though $\mathbf{X} = 70 \cap \mathbf{Y} = 34$ has 0 probability in the joint distribution.

## 5.2 Covariance and Correlation

In general, the *covariance* of any two random variables **X** and **Y** is defined to be the expected value of the product of their deviations from their respective means. That is,

(1) $Covar(\mathbf{X}, \mathbf{Y}) = E((\mathbf{X} - \mu_x)*(\mathbf{Y} - \mu_y))$

For our example, the mean $\mu_x$ (that is, the expected value $E(\mathbf{X})$) is computed in cell A16 of Figure 5.2 by the formula

```
=SUMPRODUCT(A5:A13,$C$5:$C$13)
```

The mean $\mu_y$ is similarly computed in cell B16, by copying from A16. The products of deviations from these means $(\mathbf{X} - \mu_x)*(\mathbf{Y} - \mu_y)$ are computed for the nine possible pairs of (**X**, **Y**) values by entering the formula

```
=(A5–$A$16)*(B5–$B$16)
```

into cell E5 of Figure 5.2, and then copying E5 to E5:E13. Now to compute the expected product of deviations from the means, we multiply each possible product of deviations by its corresponding probability and sum the results. So the covariance can be computed by the formula

```
=SUMPRODUCT(E5:E13,$C$5:$C$13)
```

in cell E16 of Figure 5.2.

How should we interpret this covariance? Recall that the product of two numbers is negative only when one is negative and one is positive, and the product of two negative numbers is positive. So the product of deviations from means is positive in the cases the two random variables deviate from their respective means in the same direction (that is, when both are greater than their means or both are less than their means). The products of deviations from means is negative only when the two random variables deviate from their respective means in opposite directions (one greater than its

Myerson_C05.qxd 09-09-2004 10:34 Page 157

mean while the other is less). So the covariance is a positive number here because there is more probability in the cases where the two random variables are varying from their means in the same direction, and there is less probability in the cases where the two random variables are varying from their means in opposite directions. In this sense, the covariance is a measure of the tendency of two random variables to covary.

Notice also that the covariance of a random variable with itself is the random variable's variance, which is the square of the standard deviation. That is,

$$Covar(\mathbf{X}, \mathbf{X}) = E((\mathbf{X} - \mu_x)*(\mathbf{X} - \mu_x)) = E((\mathbf{X} - \mu_x)\text{^}2) = Var(\mathbf{X}) = Stdev(\mathbf{X})\text{^}2$$

**NOTE** (Recall that the variance of **X** was defined as the expected value of the squared deviation of **X** from its own expected value, and the standard deviation was defined as the square root of the variance. Here "*Var*" stands for "variance.")

The actual numerical value of the covariance (2.2 in this example) is hard to interpret because (like the variance) it has a very strange unit of measurement. When our random variables **X** and **Y** are measured in dollars, their means and standard deviations are also measured in dollars. But the products of deviations then involve multiplying dollar amounts by other dollar amounts, and so the covariance is measured in dollars times dollars or dollars squared. (If we had measured the **X** price in yen instead of dollars, keeping **Y** in dollars, then the units of the covariance would be even stranger: dollars times yen.)

But when we divide the covariance of any two random variables by their standard deviations, we get a unit-free number called the *correlation* (or the *Pearson correlation coefficient*), which will be easier to interpret. That is, the correlation of two random variables **X** and **Y** is defined by the equation

(2) $$Correl(\mathbf{X}, \mathbf{Y}) = Covar(\mathbf{X}, \mathbf{Y})/(Stdev(\mathbf{X})*Stdev(\mathbf{Y}))$$

For our example, the correlation of **X** and **Y** is computed by the formula =E16/(A17*B17) in cell F16 in Figure 5.2. (Simtools gives us special functions COVARPR and CORRELPR for computing the covariance and correlation directly from this kind of joint probability table. These functions are illustrated by the formulas =COVARPR(A5:A13,B5:B13,C5:C13) and =CORRELPR(A5:A13,B5:B13,C5:C13) in cell E17 and F17 of Figure 5.2, where they return the same values that we got in cells E16 and F16.)

To interpret the correlation of two random variables, you should know several basic facts. The correlation of any two random variables is always a number between –1 and 1. If **X** and **Y** are independent random variables, then their correlation is 0. On the other hand, an extreme correlation of 1 or –1 occurs only when **X** depends on **Y** by a simple linear equation of the form $\mathbf{X} = c*\mathbf{Y} + d$ (where $c$ and $d$ are nonrandom numbers). Such a linear relation yields correlation 1 when the constant $c$ is positive, so that **X**

Myerson_C05.qxd 09-09-2004 10:34 Page 158

increases linearly as **Y** increases; but it yields correlation –1 when the constant $c$ is negative, so that **X** decreases linearly as **Y** increases. In particular, the correlation of any random variable with itself is always 1 (being the case of $c = 1$ and $d = 0$).

So when we find that **X** and **Y** have a correlation equal to 0.524 in this example, we may infer the **X** and **Y** have a relationship that is, in some sense, approximately midway between independence (correlation = 0) and a perfect positively sloped linear relationship (correlation = 1). A higher correlation would imply a relationship that looks more like a perfect linear relationship, whereas a smaller positive correlation would imply a relationship that looks more like independence.

## 5.3 Linear Functions of Several Random Variables

We introduced this chapter with an investor who owns 100 shares of the stock that will have price **X** per share and 200 shares of the stock that will have price **Y** per share. So the total value of her portfolio, **W**, will depend on **X** and **Y** by the linear formula $\mathbf{W} = 100*\mathbf{X} + 200*\mathbf{Y}$. In Figure 5.2, the numbers of shares that the investor owns in these two stocks have been entered into cells G3 and H3, and the possible portfolio values corresponding to each possible (**X**, **Y**) pair are shown in column I. To compute the value of **W** that corresponds to the (**X**, **Y**) values in cells A5:B5, cell I5 contains the formula

```
=SUMPRODUCT($G$3:$H$3,A5:B5)
```

and this formula has been copied from I5 to cells I5:I13. Then $E(\mathbf{W})$ and $Stdev(\mathbf{W})$ can be computed by the formulas

```
=SUMPRODUCT(I5:I13,$C5:C13)
=STDEVPR(I5:I13,$C5:C13)
```

in cells I16 and I17, respectively. The variance $Var(\mathbf{W})$, as the square of the standard deviation, is computed in cell I18 of Figure 5.2 by the formula

```
=I17^2
```

But the expected value and variance of **W** can also be calculated another way, by using some general formulas about the means and variances of linear functions of random variables. These formulas are complicated but they are important in statistics and finance.

Given any number $n$, suppose that $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n$ are a collection of $n$ random variables. Consider any random variable **W** that is defined by the linear equation

$$\mathbf{W} = \alpha_0 + \alpha_1*\mathbf{X}_1 + \cdots + \alpha_n*\mathbf{X}_n = \alpha_0 + \sum_{i=1}^{n} \alpha_i*\mathbf{X}_i$$

Myerson_C05.qxd 09-09-2004 10:34 Page 159

where $\alpha_i$ are nonrandom numbers for each $i$ in $\{0,1, \ldots ,n\}$. Then the expected value and variance of the random variable **W** can be computed by the formulas

$$(3) \quad E(\mathbf{W}) = \alpha_0 + \alpha_1 * E(\mathbf{X}_1) + \cdots + \alpha_n * E(\mathbf{X}_n) = \alpha_0 + \sum_{i=1}^{n} \alpha_i * E(\mathbf{X}_i)$$

$$(4) \quad Var(\mathbf{W}) = \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_i * \alpha_j * Covar(\mathbf{X}_i, \mathbf{X}_j)$$

You may find the double sum in equation (4) hard to read, but before explaining how to use this equation, we must make it even a bit more complicated. Recall the standard deviation of any random variable is the square root of the variance, that is $Stdev(\mathbf{W}) = Var(\mathbf{W})$^0.5, and $Covar(\mathbf{X}_i, \mathbf{X}_j)$ is equal to $Stdev(\mathbf{X}_i) * Stdev(\mathbf{X}_j) * Correl(\mathbf{X}_i, \mathbf{X}_j)$. So equation (4) implies

$$(4') \quad Var(\mathbf{W}) = \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_i * \alpha_j * Stdev(\mathbf{X}_i) * Stdev(\mathbf{X}_j) * Correl(\mathbf{X}_i, \mathbf{X}_j)$$

$$(5) \quad Stdev(\mathbf{W}) = \left( \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_i * \alpha_j * Stdev(\mathbf{X}_i) * Stdev(\mathbf{X}_j) * Correl(\mathbf{X}_i, \mathbf{X}_j) \right) \text{^}0.5$$

Equations (3) and (5) tell us that if we know the expected values $E(\mathbf{X}_i)$, standard deviations $Stdev(\mathbf{X}_i)$, and correlations $Correl(\mathbf{X}_i, \mathbf{X}_j)$ for all $i$ and $j$, then we can compute the expected value and standard deviation of any such linear function **W** without knowing anything more about the joint distribution of these $\mathbf{X}_i$ random variables. But these formulas are complicated, and so we should now see how they can be implemented in a spreadsheet.

To see how to use these equations in practical calculations, let us return to our example in Figure 5.2. In this example, there are just two asset prices. So $n = 2$ here, and **X** and **Y** here are taking the place of $\mathbf{X}_1$ and $\mathbf{X}_2$ in the general equations given earlier. We are interested in the future portfolio value $\mathbf{W} = 100*\mathbf{X} + 200*\mathbf{Y}$. So here $\alpha_0$ is 0, and the coefficients $\alpha_1$ and $\alpha_2$ represent the numbers of shares of the two assets that the investor owns, and so $\alpha_1 = 100$ and $\alpha_2 = 200$.

Equation (3) tells us that the expected value of **W** here can be computed by the simple linear formula

$$E(\mathbf{W}) = 100*E(\mathbf{X}) + 200*E(\mathbf{Y})$$

With the share coefficients 100 and 200 listed in G3:H3 of Figure 5.2, and with the expected values $E(\mathbf{X})$ and $E(\mathbf{Y})$ listed in A16:B16, this linear formula for $E(\mathbf{W})$ can be computed in cell K16 by the formula

```
=SUMPRODUCT(G3:H3,A16:B16)
```

To check, notice cells K16 and I16 are equal in Figure 5.2. More generally, equation (3) tells us that any portfolio's expected future value can be computed with the spreadsheet formula

```
=SUMPRODUCT(shares,means)
```

where *shares* denotes a range that lists the numbers of shares of each financial asset that we are holding in this portfolio, and *means* denotes a range that lists the expected values of each asset's future price per share.

Myerson_C05.qxd 09-09-2004 10:34 Page 160

Using equation (5) to compute *Stdev*(**W**) is much more complicated, because equation (5) refers to correlations *Correl*($\mathbf{X}_i$, $\mathbf{X}_j$) for all pairs of assets $\mathbf{X}_i$ and $\mathbf{X}_j$ that we are using in our portfolio. When there are $n$ assets, we have $n$ choices for $\mathbf{X}_i$ and $n$ choices for $\mathbf{X}_j$, and so we have a matrix of $n*n$ correlations appearing in this formula. To do computations in a spreadsheet, we will regularly list all these correlations together in a *correlation array* (or correlation matrix). Such a correlation array for our example is shown in the range D20:E21 of Figure 5.2. With **X** as our first asset price and **Y** as our second asset price, the correlation array is a table listing the following four numbers:

| | |
|---|---|
| *Correl*(**X**, **X**) = 1 | *Correl*(**X**, **Y**) = 0.524 |
| *Correl*(**Y**, **X**) = 0.524 | *Correl*(**Y**, **Y**) = 1 |

In any correlation array, the elements on the diagonal from top left to bottom right are always 1s because the correlation of any random variable with itself is 1. Below this diagonal and to its right, the elements of the correlation array are always symmetric because *Correl*(**X**, **Y**) equals *Correl*(**Y**, **X**).

Inside the summation in equation (5), we are told to multiply each number in the correlation array by a corresponding product of share coefficients $\alpha_i*\alpha_j$ and by a corresponding product of standard deviations *Stdev*($\mathbf{X}_i$)*Stdev*($\mathbf{X}_j$). The four products of the share coefficients $\alpha_1$=100 and $\alpha_2$=200 are

| | |
|---|---|
| 100∗100 = 10,000 | 100∗200 = 20,000 |
| 200∗100 = 20,000 | 200∗200 = 40,000 |

as shown in cells G20:H21 in Figure 5.2. The four products of the standard deviations *Stdev*(**X**) = 3.5 and *Stdev*(**Y**) = 1.2 are

| | |
|---|---|
| 3.5∗3.5 = 12.25 | 3.5∗1.2 = 4.2 |
| 1.2∗3.5 = 4.2 | 1.2∗1.2 = 1.44 |

as shown in cells A20:B21 of Figure 5.2. It would not be difficult to make a table that computes all these products, by entering multiplication formulas into each cell, but Simtools gives us a function called PRODS that makes it easier. But before introducing PRODS, we first have to learn about array formulas in Excel.

An *array formula* in Excel is a formula that is shared by a whole range of cells at once. To enter an array formula, first select the desired range (say, A20:B21), then type the desired formula (say, =PRODS(A17:B17)), and finish with the special keystroke combination [Ctrl][Shift][Enter]. (That is, while holding down the [Ctrl] and [Shift] keys, press [Enter].) Now if you look at the formula for any cell in the range that you selected, you should see the array formula that you entered within a pair of special braces { }, which were added by Excel to indicate that this is an array formula. (Typing the braces yourself will not work; you must use [Ctrl][Shift][Enter]. From any

Myerson_C05.qxd 09-09-2004 10:34 Page 161

selected cell that has an array formula, you can identify the whole range of cells that share the formula by the special keystroke [Ctrl]/.)

The Simtools function PRODS is designed to be used in such an array formula. If *values* denotes a range containing numbers in a column or row of a spreadsheet, then PRODS(*values*) returns a square array containing the results of multiplying each pair of numbers in the values range. The *i*th row and *j*th column of the array returned by PRODS will be the result of multiplying the *i*th entry in the *values* range times the *j*th entry in the *values* range. So in Figure 5.2, the array formula

{=PRODS(A17:B17)}

has been entered into the range A20:B21 to return the array of products of the standard deviations from A17 and B17. Similarly, the array formula

{=PRODS(G3:H3)}

has been entered into the range G20:H21 to return the array of products of the shares from G3 and H3.

The variance of the linear portfolio **W** can now be computed in cell K18 of Figure 5.2 by the formula

=SUMPRODUCT(G20:H21,A20:B21,D20:E21)

(Notice that Excel's SUMPRODUCT function can be applied to any number of ranges.) Here SUMPRODUCT multiplies each product of standard deviations in A20:B21 by the corresponding correlation in D20:E21 and by the corresponding product of share coefficients in G20:H21, and the results of all these various multiplications are summed, to return the same variance 268100 ($$) that we found before in cell I18.

The standard deviation can now be computed as the square root (that is, the power ^0.5) of the variance. So the formula for *Stdev*(**W**) in cell K17 of Figure 5.2 is

=SUMPRODUCT(PRODS(G3:H3),PRODS(A17:B17),D20:E21)^0.5

Notice here that the PRODS function can be used inside a regular formula as well as in an array formula. That is, the $2 \times 2$ array returned by PRODS(A17:B17) can be interpreted directly in the SUMPRODUCT function, and it does not actually need to be entered into any other $2 \times 2$ range in the spreadsheet. More generally, equation (5) tells us that the standard deviation of any portfolio's value can be computed in a spreadsheet cell with the formula

=SUMPRODUCT(PRODS(*shares*),PRODS(*stdevns*),*correls*)^0.5

Here *shares* denotes a range that lists the number of shares in each financial asset that we are holding in this portfolio, *stdevns* denotes a range that lists the standard deviations of each asset's price per share, and *correls* denotes a square range that lists the pairwise correlations among these assets' prices.

Myerson_C05.qxd 09-09-2004 10:34 Page 162

## 5.4 Estimating Correlations from Data

Excel provides the statistical function CORREL for estimating the correlations among random variables drawn from some joint distribution that we do not know but from which we have observed repeated independent draws. To illustrate the use of this statistical function, Figure 5.3 shows price data from the decade of the 1980s for six large stock funds that are tied to major market indexes. (Funds 1, 2, 3, and 4 represent broad portfolios of stocks in the New York Stock Exchange. Fund 1 contains stocks of companies that are broadly characterized as industrial, Fund 2 contains stocks of transportation companies, Fund 3 contains stocks of utility companies, and Fund 4 contains stocks of financial companies. Funds 5 and 6 are well-known selected portfolios of leading industrial companies.) The end-of-year average price-per-share data for these stock funds is shown in the range B3:G13 of Figure 5.3. Then the annual growth ratio for the values of these funds is then computed in range B17:G26 by entering the formula =B4/B3 into cell B17 and then copying cell B17 to B17:G26. So the 1980 Fund 1 growth ratio 1.22 in cell B17, for example, tells us that each dollar invested in Fund 1 just before 1980 began would have been worth $1.22 at the end of 1980.

Our goal is to use this growth ratio data from the 1980s to build simple forecasting models that could have helped to make predictions about the

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | End-of-year price data for six stock funds | | | | | | |
| 2 | | Fund 1 | Fund 2 | Fund 3 | Fund 4 | Fund 5 | Fund 6 |
| 3 | 1979 | 64.76 | 47.34 | 38.20 | 61.42 | 84.44 | 103.01 |
| 4 | 1980 | 78.70 | 60.61 | 37.35 | 64.25 | 89.14 | 118.78 |
| 5 | 1981 | 85.44 | 72.61 | 38.91 | 73.52 | 93.29 | 128.05 |
| 6 | 1982 | 78.18 | 60.41 | 39.75 | 71.99 | 88.44 | 119.71 |
| 7 | 1983 | 107.45 | 89.36 | 47.00 | 95.34 | 119.03 | 160.41 |
| 8 | 1984 | 108.01 | 85.63 | 46.44 | 89.28 | 117.85 | 160.46 |
| 9 | 1985 | 123.79 | 104.11 | 56.75 | 114.21 | 132.82 | 186.84 |
| 10 | 1986 | 155.85 | 119.87 | 71.36 | 147.20 | 179.28 | 236.34 |
| 11 | 1987 | 195.31 | 140.39 | 74.30 | 146.48 | 227.60 | 286.83 |
| 12 | 1988 | 180.95 | 134.12 | 71.77 | 127.26 | 206.08 | 265.79 |
| 13 | 1989 | 216.23 | 175.28 | 87.43 | 151.88 | 250.89 | 322.84 |
| 14 | | | | | | | |
| 15 | Annual growth ratios | | | | | | |
| 16 | | Fund 1 | Fund 2 | Fund 3 | Fund 4 | Fund 5 | Fund 6 |
| 17 | 1980 | 1.22 | 1.28 | 0.98 | 1.05 | 1.06 | 1.15 |
| 18 | 1981 | 1.09 | 1.20 | 1.04 | 1.14 | 1.05 | 1.08 |
| 19 | 1982 | 0.92 | 0.83 | 1.02 | 0.98 | 0.95 | 0.93 |
| 20 | 1983 | 1.37 | 1.48 | 1.18 | 1.32 | 1.35 | 1.34 |
| 21 | 1984 | 1.01 | 0.96 | 0.99 | 0.94 | 0.99 | 1.00 |
| 22 | 1985 | 1.15 | 1.22 | 1.22 | 1.28 | 1.13 | 1.16 |
| 23 | 1986 | 1.26 | 1.15 | 1.26 | 1.29 | 1.35 | 1.26 |
| 24 | 1987 | 1.25 | 1.17 | 1.04 | 1.00 | 1.27 | 1.21 |
| 25 | 1988 | 0.93 | 0.96 | 0.97 | 0.87 | 0.91 | 0.93 |
| 26 | 1989 | 1.19 | 1.31 | 1.22 | 1.19 | 1.22 | 1.21 |
| 27 | | | | | | | |
| 28 | FORMULAS | | | | | | |
| 29 | B17. =B4/B3 | | | | | | |
| 30 | B17 copied to B17:G26. | | | | | | |

FIGURE **5.3**
**Price data and annual growth ratios in the 1980s for six mutual funds**

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | | Fund 1 | Fund 2 | | Products of deviations from means | | | |
| 2 | 1980 | 1.22 | 1.28 | | 0.0098 | | | |
| 3 | 1981 | 1.09 | 1.20 | | -0.0022 | | | |
| 4 | 1982 | 0.92 | 0.83 | | 0.0718 | | | |
| 5 | 1983 | 1.37 | 1.48 | | 0.0768 | | | |
| 6 | 1984 | 1.01 | 0.96 | | 0.0260 | | | |
| 7 | 1985 | 1.15 | 1.22 | | 0.0005 | | | |
| 8 | 1986 | 1.26 | 1.15 | | -0.0004 | | | |
| 9 | 1987 | 1.25 | 1.17 | | 0.0019 | | | |
| 10 | 1988 | 0.93 | 0.96 | | 0.0421 | | | |
| 11 | 1989 | 1.19 | 1.31 | | 0.0087 | | | |
| 12 | | | | | | | | |
| 13 | Means | 1.1375 | 1.1548 | | 0.0261 | Sum/(n-1) [Covariance] | | |
| 14 | StDevns | 0.1520 | 0.1922 | | | | | |
| 15 | Correl | 0.8937 | | | 0.8937 | Correlation | | |
| 16 | n=Count | 10 | | | | | | |
| 17 | | | | | FORMULAS | | | |
| 18 | Squared deviations from means | | | | B13. =AVERAGE(B2:B11) | | | |
| 19 | | 0.0060 | 0.0157 | | B14. =STDEV(B2:B11) | | | |
| 20 | | 0.0027 | 0.0019 | | B13:B14 copied to C13:C14 | | | |
| 21 | | 0.0495 | 0.1042 | | B15. =CORREL(B2:B11,C2:C11) | | | |
| 22 | | 0.0561 | 0.1052 | | B16. =COUNT(B2:B11) | | | |
| 23 | | 0.0175 | 0.0386 | | E2. =(B2-$B$13)*(C2-$C$13) | | | |
| 24 | | 0.0001 | 0.0037 | | E2 copied to E2:E11 | | | |
| 25 | | 0.0148 | 0.0000 | | E13. =SUM(E2:E11)/(B16-1) | | | |
| 26 | | 0.0134 | 0.0003 | | E15. =E13/(B14*C14) | | | |
| 27 | | 0.0445 | 0.0398 | | B19. =(B2-B$13)^2 | | | |
| 28 | | 0.0033 | 0.0231 | | B19 copied to B19:C28 | | | |
| 29 | | | | | B30. =SUM(B19:B28)/($B$16-1) | | | |
| 30 | Sum/(n-1) | 0.0231 | 0.0370 | | B31. =B30^0.5 | | | |
| 31 | SqrRoot | 0.1520 | 0.1922 | | B30:B31 copied to C30:C31 | | | |
| 32 | | | | | B35:C35. {=CORAND(B15)} | | | |
| 33 | | | | | B38. =NORMINV(B35,B13,B14) | | | |
| 34 | Simulation Model: | | | | C38. =NORMINV(C35,C13,C14) | | | |
| 35 | CORANDs | 0.0309 | 0.0610 | | | | | |
| 36 | Multivariate Normals with means, stdevs, and correl as above. | | | | | | | |
| 37 | | Fund 1 | Fund 2 | | | | | |
| 38 | | 0.85 | 0.86 | | | | | |
| 39 | | | | | A bad model with independent funds | | | |
| 40 | | | | | Fund 1 | Fund 2 | | |
| 41 | | | | | 0.92 | 1.44 | | |
| 42 | | | | | FORMULAS FROM RANGE E39:F41 | | | |
| 43 | | | | | E41. =NORMINV(RAND(),B13,B14) | | | |
| 44 | | | | | F41. =NORMINV(RAND(),C13,C14) | | | |

FIGURE **5.4**
**Computing the correlation of Funds 1 and 2**

possible returns to various investment portfolios in the subsequent decade. The basic assumption of our analysis will be that the annual growth ratios of these funds are jointly drawn each year out of some joint probability distribution. With this assumption, we can try to estimate the means, standard deviations, and correlations of this distribution from the past data. Then we will use these estimates to make a simulation model of future returns from any portfolio of investments in these six funds.

But six funds is a lot to think about, so let us begin considering just Funds 1 and 2 alone, as shown in Figure 5.4.

We have already seen (in Chapter 2) that the unknown means and standard deviations for each random variable can be estimated from our statistical data by the sample average and the sample standard deviation, which

Myerson_C05.qxd 09-09-2004 10:34 Page 164

can be computed in Excel by the functions AVERAGE and STDEV. The annual growth ratio data for Fund 1 is exhibited in the range B2:B11 of Figure 5.4. So cell B13 estimates the mean annual growth ratio for Fund 1 by the formula =AVERAGE(B2:B11), and cell B14 estimates the standard deviation of Fund 1's annual growth ratio by the formula =STDEV(B2:B11). Cells B18:B31 in Figure 5.4 just remind you of how the STDEV function computes this estimated standard deviation, as the square root of an adjusted average of squared deviations from the sample mean. (The average in cell B30 is "adjusted" in that it is divided by one less than the number of data points, which are counted in cell B16 by the formula =COUNT(B2:B11).)

The annual growth ratio data for Funds 2 is exhibited in the range C2:C11 of Figure 5.4. So copying cells B13 and B14 to C13 and C14 similarly yields our statistical estimates of the mean and standard deviation of the annual growth ratio for Fund 2.

Now, to estimate the correlation between the annual growth ratios of Funds 1 and 2, the formula

=CORREL(B2:B11,C2:C11)

has been entered into cell B15 in Figure 5.4. The E column in Figure 5.4 shows in more detail how Excel's CORREL function computes this estimated correlation. First, the product of the two random variables' deviations from their respective means is computed, which is done in the spreadsheet of Figure 5.4 by entering the formula

=(B2 – $B$13)*(C2 – $C$13)

into cell E2 and then copying cell E2 to E2:E11. Next, the adjusted average (the sum divided by $n - 1$, where $n$ is the number of observations that we have of each random variable) of these products of deviations is computed, as shown in cell E13, by the formula

=SUM(E2:E11)/(B16–1)

This adjusted average of the observed products of deviations is a good estimate of the covariance of these two random variables (which has been defined as the expected product of their deviations from their respective means). But the correlation is the covariance divided by both of the standard deviations. So we can compute the estimated correlation from the estimated covariance in cell E15 by the formula

=E13/(B14*C14)

Notice that this formula in cell E15 returns the same value as the CORREL function in cell B15. The high value of this correlation, 0.8937, tells us that these data points are almost perfectly linearly related (see also Figure 5.6).

By the law of large numbers, if we get a very large sample of paired random variables (like the growth ratios of Funds 1 and 2 here) that are drawn together from some fixed joint distribution, with each pair drawn independently of all other pairs, then the sample correlation computed by

| | A | B | C | D | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Annual growth ratios for six stock funds | | | | | | | | | | |
| 2 | | Fund1 | Fund2 | Fund3 | Fund4 | Fund5 | Fund6 | | | | |
| 3 | 1980 | 1.22 | 1.28 | 0.98 | 1.05 | 1.06 | 1.15 | | | | |
| 4 | 1981 | 1.09 | 1.20 | 1.04 | 1.14 | 1.05 | 1.08 | | | | |
| 5 | 1982 | 0.92 | 0.83 | 1.02 | 0.98 | 0.95 | 0.93 | | | | |
| 6 | 1983 | 1.37 | 1.48 | 1.18 | 1.32 | 1.35 | 1.34 | | | | |
| 7 | 1984 | 1.01 | 0.96 | 0.99 | 0.94 | 0.99 | 1.00 | FORMULAS | | | |
| 8 | 1985 | 1.15 | 1.22 | 1.22 | 1.28 | 1.13 | 1.16 | B14. =AVERAGE(B3:B12) | | | |
| 9 | 1986 | 1.26 | 1.15 | 1.26 | 1.29 | 1.35 | 1.26 | B15. =STDEV(B3:B12) | | | |
| 10 | 1987 | 1.25 | 1.17 | 1.04 | 1.00 | 1.27 | 1.21 | B14:B15 copied to B14:G15 | | | |
| 11 | 1988 | 0.93 | 0.96 | 0.97 | 0.87 | 0.91 | 0.93 | B19:G24. {=MCORRELS(B3:G12)} | | | |
| 12 | 1989 | 1.19 | 1.31 | 1.22 | 1.19 | 1.22 | 1.21 | G28. =SUM(B27:G27) | | | |
| 13 | | | | | | | | B30:G30. {=CORAND(B19:G24)} | | | |
| 14 | Mean | 1.138 | 1.155 | 1.092 | 1.106 | 1.126 | 1.129 | B33. =NORMINV(B30,B14,B15) | | | |
| 15 | StDev | 0.152 | 0.192 | 0.115 | 0.162 | 0.162 | 0.140 | B33 copied to B33:G33 | | | |
| 16 | | | | | | | | I34. =SUMPRODUCT(B34:G34,$B$27:$G$27) | | | |
| 17 | Correlation coefficients | | | | | | | I34 copied to I34:I534 | | | |
| 18 | | Fund1 | Fund2 | Fund3 | Fund4 | Fund5 | Fund6 | I28. =AVERAGE(I34:I534) | | | |
| 19 | Fund1 | 1 | .8937 | .5980 | .7280 | .9178 | .9846 | I29. =STDEV(I34:I534) | | | |
| 20 | Fund2 | .8937 | 1 | .5389 | .7404 | .7216 | .8804 | I30. =PERCENTRANK(I34:I534,A28) | | | |
| 21 | Fund3 | .5980 | .5389 | 1 | .8992 | .7627 | .7260 | I31. =PERCENTILE(I34:I534,0.05) | | | |
| 22 | Fund4 | .7280 | .7404 | .8992 | 1 | .7597 | .8109 | I32. =CE(I34:I534,I25) | | | |
| 23 | Fund5 | .9178 | .7216 | .7627 | .7597 | 1 | .9550 | K36. =CORREL(B34:B534,C34:C534) | | | |
| 24 | Fund6 | .9846 | .8804 | .7260 | .8109 | .9550 | 1 | | | | |
| 25 | | | | | | | | | 6000 | RiskTol for CE | |
| 26 | Portfolio: $ to invest now, for W next year | | | | | | | | | | |
| 27 | | 3000 | 3000 | 1000 | 1000 | 1000 | 1000 | | W stats from sim data | | |
| 28 | | | | | | TotalNow | 10000 | | 11353.32 | E(W) | |
| 29 | | CORANDs | | | | | | | 1491.17 | Stdev(W) | |
| 30 | | 0.118 | 0.119 | 0.141 | 0.139 | 0.066 | 0.099 | | 0.188 | P(W<G28) | |
| 31 | MultivariateNormals with above parameters | | | | | | | | 8876.97 | 5%ile(W) | |
| 32 | | Fund1 | Fund2 | Fund3 | Fund4 | Fund5 | Fund6 | | 11168.60 | CE(W) | |
| 33 | SimTab | 0.96 | 0.93 | 0.97 | 0.93 | 0.88 | 0.95 | | W | | |
| 34 | 0 | 1.28 | 1.22 | 1.22 | 1.27 | 1.34 | 1.27 | | 12610.64 | From sim data | |
| 35 | .002 | 1.22 | 1.20 | 1.05 | 1.08 | 1.19 | 1.19 | | 11784.22 | Correl(1,2) | |
| 36 | .004 | 0.98 | 1.06 | 0.86 | 0.92 | 0.88 | 0.94 | | 9737.27 | | 0.9 |
| 37 | .006 | 1.18 | 1.12 | 1.00 | 1.04 | 1.19 | 1.14 | | 11275.56 | | |
| 38 | .008 | 1.00 | 0.91 | 0.85 | 0.73 | 0.96 | 0.96 | | 9233.78 | | |
| 39 | .010 | 1.03 | 1.08 | 1.08 | 1.00 | 1.00 | 1.04 | | 10434.28 | | |
| 40 | .012 | 0.96 | 0.91 | 1.00 | 0.98 | 0.95 | 0.97 | | 9518.15 | | |
| 41 | .014 | 1.01 | 0.89 | 0.96 | 0.84 | 1.02 | 0.99 | | 9495.47 | | |
| 42 | .016 | 1.13 | 1.18 | 1.24 | 1.21 | 1.15 | 1.16 | | 11693.57 | | |
| 43 | .018 | 1.26 | 1.27 | 1.33 | 1.36 | 1.29 | 1.28 | | 12868.70 | | |
| 44 | .020 | 1.05 | 1.20 | 0.96 | 0.95 | 0.90 | 1.02 | | 10551.77 | | |

FIGURE **5.5**
**Making a model of the six funds' annual growth ratios**

CORREL is very likely to be very close to the true correlation of the joint probability distribution from which we have sampled.

Figure 5.5 takes up the problem of estimating means, standard deviations, and correlations for the annual growth ratios of all six funds. The observed growth ratios from Figure 5.3 are repeated in cells B3:G12 in Figure 5.5, with the 10 observed growth ratios for Fund 1 in B3:B12, the 10 observed growth ratios for Fund 2 in C3:C12, and so on. The expected growth ratio for each fund is estimated in row 14 by the AVERAGE of its growth data. The standard deviation of each fund's growth ratio is estimated in row 15 by the sample standard deviation (STDEV) of its growth data. But with six funds, we have 15 different correlations to compute

Myerson_C05.qxd 09-09-2004 10:34 Page 166

between every pair of distinct funds, to fill a 6 × 6 correlation array (which will be symmetric across a diagonal of 1s). The correlations for each of these pairs can be computed using Excel's CORREL, but typing in 15 different CORREL formulas can become tedious. So Simtools provides a function called MCORRELS that allows us to compute all these correlations with one array formula. (The "M" in this function name stands for *matrix*, which is another mathematical term for an array.)

MCORRELS returns the correlations among the columns of the data range, which is specified as its input parameter. If *dataRange* is a range that contains *n* columns of data, then {=MCORRELS(*dataRange*)} should be entered as an array formula into an *n*-by-*n* range. The value returned by MCORRELS to the *i*th row and *j*th column of its output range is the sample correlation that would be computed by CORREL between the *i*th and *j*th columns of the data range. In Figure 5.5, the data range is B3:G12, with one data column for each of the six funds. So the array formula

{=MCORRELS(B3:G12)}

has been entered into the 6 × 6 range B19:G24 in Figure 5.5. (To enter this formula, first the range B19:G24 was selected, then the formula =MCORRELS(B19:G24) was typed, and then the [Ctrl] and [Shift] keys were held down while pushing the [Enter] key.) So in the *i*th row and *j*th column of this range B19:G24, we find the correlation between Fund *i* and Fund *j*. For example, cell C19 is in the first row and second column of this MCORRELS array in B19:G24, and so the value of cell C19 in Figure 5.5 is 0.8937, which is the correlation between Fund 1 and Fund 2 that we already computed in Figure 5.4.

Looking at the correlation array in cells B19:G24 of Figure 5.5, you can see that the lowest correlation is 0.5389, between Fund 2 (the transportation stocks) and Fund 3 (the utility stocks). The highest correlation between distinct funds is 0.9846, between Fund 1 and Fund 6, both of which are mixes of industrial stocks. Cells along the diagonal of the correlation array have the value 1 because the sample correlation of any data column with itself is always 1. The correlation array is always symmetric around this diagonal of 1s, because

CORREL(*column1*, *column2*) = CORREL(*column2*, *column1*)

for any two columns of data.

## 5.5 Making Multivariate Normal Random Variables with CORAND and NORMINV

When we have a mean and a standard deviation for some unknown quantity, it is often analytically convenient (and not unreasonable) to assume that this unknown quantity has a Normal probability distribution with the

Myerson_C05.qxd 09-09-2004 10:34 Page 167

given mean and standard deviation. This convenient Normality assumption is commonly extended to a collection of random variables by assuming that they jointly have a Multivariate Normal distribution.

The *Multivariate Normal* distributions are a family of joint probability distributions for collections of two or more random variables. To specify a Multivariate Normal distribution for any collection of random variables $\mathbf{X}_1, \ldots, \mathbf{X}_n$, we must specify the mean $E(\mathbf{X}_i)$ and the standard deviation $Stdev(\mathbf{X}_i)$ for each of these random variables, and the correlation $Correl(\mathbf{X}_i, \mathbf{X}_j)$ for each pair of these random variables. That is, the Multivariate Normal distributions for $n$ random variables are parameterized by $n$ means, $n$ standard deviations, and an $n \times n$ correlation array. Given these parameters, there is a (complicated) mathematical formula that defines the joint probability density function of the corresponding Multivariate Normal distribution. You can find this formula in any advanced probability text, but it will not tell you how these distributions are really used. So here, instead, we will discuss some of the important properties of Multivariate Normal random variables. Most importantly, we will see how to make Multivariate Normal random variables, because the philosophy of this book is that a good way to learn about any kind of random variables is to make them in a spreadsheet.

If $(\mathbf{X}_1, \ldots, \mathbf{X}_n)$ are such Multivariate Normal random variables (that is, their joint probability distribution is in this Multivariate Normal family), then each of these random variables $\mathbf{X}_i$ is (by itself) a Normal random variable. But more importantly, any linear function of such Multivariate Normal random variables will also be a Normal random variable. That is, given any nonrandom numbers $\alpha_0, \alpha_1, \ldots, \alpha_n$, the linear function

$$\mathbf{W} = \alpha_0 + \alpha_1 * \mathbf{X}_1 + \ldots + \alpha_n * \mathbf{X}_n$$

will also have a Normal probability distribution when $(\mathbf{X}_1, \ldots, \mathbf{X}_n)$ are Multivariate Normal.

This fact, that every linear function of Multivariate Normals is also Normal, makes the Multivariate Normal joint distribution very convenient for financial analysis. We have seen how the mean and standard deviation of a portfolio's value can be computed from the means, standard deviations, and correlations of the various asset prices (by using equations (3) and (5)). So if we can assume that these asset prices have a Multivariate Normal joint distribution, then our portfolio's value will have the Normal distribution with the computed mean and standard deviation, and so we can easily compute its probability of being greater or less than any given number.

In the special case where the correlations $Correl(\mathbf{X}_i, \mathbf{X}_j)$ are all 0 (for all $i \neq j$), the Multivariate Normal random variables are just independent Normal random variables. So any collection of independent Normal random variables is included in the Multivariate Normal family.

But we must now see how to construct Multivariate Normal random variables in general, when correlations are not 0. Simtools provides an array function called CORAND that makes it easy to construct random variables with any Multivariate Normal distribution. (An equivalent construction

Myerson_C05.qxd 09-09-2004 10:34 Page 168

can be done in Excel without an add-in, but it requires more mathematical sophistication and will be discussed later in Section 5.10.)

Simtools's CORAND function has one required parameter, which is a correlation array. When we want to make two random variables, however, the correlation array can also be just a single number that denotes the desired correlation between these two random variables. In Figure 5.4, for example, CORAND has been used to make two random variables that have the correlation 0.8937, which we found in cell B15 for the growth ratios of Funds 1 and 2 in the 1980s data. To make two random variables, CORAND must be entered as an array formula into two cells. This has been done in cells B35:C35 of Figure 5.4 with the array formula

{=CORAND(B15)}

To enter this array formula, we select the range B35:C35 in this spreadsheet, type the formula =CORAND(B15), and then use the special [Ctrl][Shift][Enter] keystroke.

In a range that contains such a CORAND formula, each cell contains a random variable that is drawn from a Uniform distribution on the interval from 0 to 1, just like a RAND( ) formula. But unlike RANDs, the values in a CORAND array are not independent of each other. Because cell B15 contains a high correlation value that is close to 1, the values in B35 and C35 that are generated by this array formula {=CORAND(B15)} tend to stay rather close together whenever we recalculate the spreadsheet. (You should try this yourself in a spreadsheet now. But notice that if you enter separate CORAND array formulas in different ranges of a spreadsheet, then the CORAND values are independent across these separate arrays. The correlation among CORAND values only applies to cells that are part of the same array formula.)

This CORAND function has been designed so that when the values from a CORAND array formula are used as the first-parameter inputs into a collection of NORMINV functions (taking the place ordinarily filled by RANDs), the resulting NORMINV values are Multivariate Normal random variables. A pair of such Multivariate Normal random variables has been constructed in cells B38 and C38 of Figure 5.4, which contain the formulas

=NORMINV(B35,B13,B14) and =NORMINV(C35,C13,C14)

The result is that cells B38:C38 contain Multivariate Normal random variables with the means computed in B13 and C13, the standard deviations computed in B14 and C14, and with the correlation computed in B15, because B15 is the correlation parameter of the CORAND array formula in B35:C35. Recall that cells B13 and B14 contain the mean and standard deviation of Fund 1's observed growth ratios, and cells C13 and C14 contain the mean and standard deviation of Fund 2's observed growth ratios. So cells B38 and C38 here are Multivariate Normal random variables that have the same means, standard deviations, and correlations as we observed in the data for Funds 1 and 2.

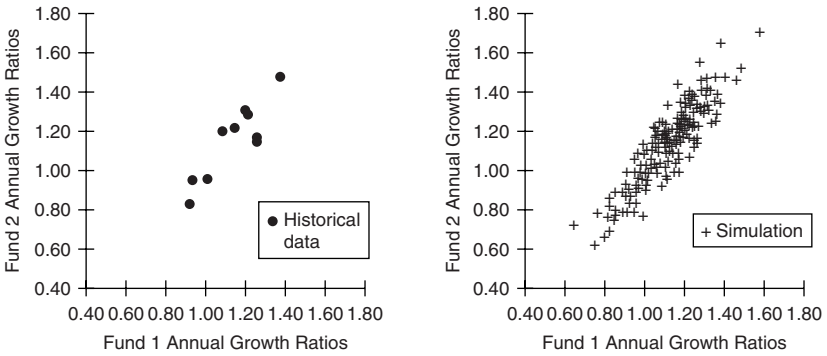


FIGURE **5.6**
**Extending historical data by simulation**

Figure 5.6 compares the results of 200 simulations of these Multivariate Normal random variables (on the right) with the 10 historical data points (on the left). Comparing these two scatterplots, you can see the power of a simulation model. The model enables us to generate large simulated data sets that exhibit the same basic patterns and relationships as our real data and thus may enable us to see these patterns and relationships more clearly.

The simulation model can also give us more realistic answers to detailed questions than we might otherwise pull out of our data. Consider, for example, the event that Fund 1's growth ratio is less than 1 while Fund 2's growth ratio is greater than 1. Based on this data, how should we estimate the probability that this event will occur in the coming year? This event never happened in our 10 years of observation, but such a failure to observe the event in 10 years does not prove that this event cannot happen next year. Indeed, the event "almost" occurred in 1988, for example. In the simulation data shown in Figure 5.6, this event occurred 7 out of 200 times. So the simulation data suggests 7/200 = 0.035 as an estimated probability of this event, which seems more reasonable than 0.

Consider now what would happen if we did not use CORAND in this model, but instead simulated the funds' growth ratios as independent Normal random variables. Such a model is shown in cells E41 and F41 of Figure 5.4, which contain the formulas

=NORMINV(RAND( ),B13,B14) and =NORMINV(RAND( ),C13,C14)

So cells E41 and F41 here are Normal random variables with the same means and standard deviations as the Multivariate Normal random variables in cells B38 and C38, but E41 and F41 are independent Normals with correlation 0 because they are not linked through any CORAND array. Figure 5.7 plots 201 simulated values of E41:F41, and the roundness of their scatter looks very different from the patterns that we saw in Figure 5.6. In this scatterplot, the event that Fund 1's growth ratio is less than 1 while Fund 2's growth ratio is greater than 1 occurs about 14% of the time, and extreme cases where Fund 1 growth ratio is less than 1 while Fund 2's growth ratio is greater than 1.30 are not rare (as illustrated in Figure 5.4 itself).

Myerson_C05.qxd 09-09-2004 10:34 Page 170

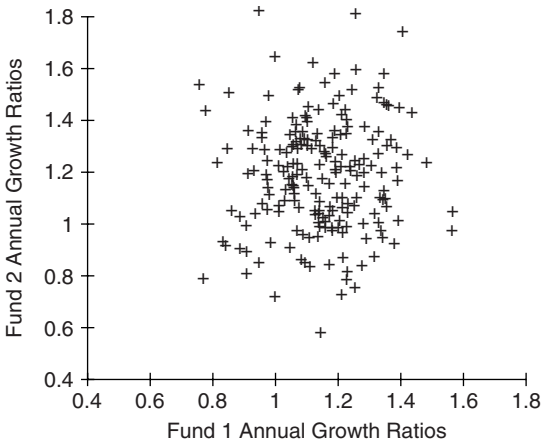


FIGURE **5.7**
**Simulations with independent funds**

If you had money invested in both Funds 1 and 2, your portfolio would look significantly less risky in this E41:F41 model than in the B38:C38 model because many occurrences of low values for Fund 1 in E41 could be counterbalanced by simultaneous high values for Fund 2 in F41. That is, the E41:F41 model would underestimate the probability of your investments both doing badly at the same time. So such a model of the funds as independent random variables would be seriously misleading and inconsistent with the relationships among the funds' annual performances that we observed in the 1980s data.

Figure 5.5 illustrates the construction of larger collections of Multivariate Normal random variables. To make more than two correlated random variables with CORAND, the required parameter of the CORAND array formula must be a square correlation array that contains a row and a column for each random variable that we are making. Recall that the range B19:G24 in Figure 5.5 contains the $6 \times 6$ correlation array for the annual growth ratios of our six stock funds. So to make random variables that have the correlations in this array, the array formula

{=CORAND(B19:G24)}

has been entered (with [Ctrl][Shift][Enter]) into the range B30:G30. Then these CORAND values have been used as inputs to NORMINV functions, along with the means from row 14 and the standard deviations from row 15, by entering the formula

=NORMINV(B30,B14,B15)

into cell B33, and then copying cell B33 to B33:G33. Thus, cells B33:G33 in Figure 5.5 contain six random variables that have a Multivariate Normal joint probability distribution with the same means, standard deviations, and correlations that we found in the 1980s data for the annual growth ratios of these six funds.

Myerson_C05.qxd 09-09-2004 10:34 Page 171

Data from 501 simulations of these six random variables is contained in cells B34:G534 of the spreadsheet in Figure 5.5. As usual, each row of this simulation table represents an independent simulation of the six random variables, but within each row the six random variables have been correlated according to the array in cells B19:G24. Of course, the actual correlation between the various columns of the simulation data may differ slightly from the theoretical Multivariate Normal distribution that generated it. For example, cells B33 and C33 are driven by the CORAND array here to have correlation 0.8937 (from cells C19 and B20). But the simulation data generated by these cells in this spreadsheet actually has correlation 0.900, as shown in cell K36 with the formula =CORREL(B34:B534,C34:C534).

## 5.6 Portfolio Analysis with Multivariate Normal Asset Returns

To introduce the methods of portfolio analysis, consider the portfolio of investments shown in cells B27:G27 of Figure 5.5. In this portfolio, we invest \$3000 now in Funds 1 and 2, and we invest \$1000 now in Funds 3, 4, 5, and 6. So our total investments now are \$10,000. We want to analyze the potential returns from this portfolio of investments next year. If we let $\mathbf{X}_i$ denote the growth ratio of Fund 1 in the coming year, then the total value of this portfolio next year will be

$$\mathbf{W} = 3000*\mathbf{X}_1 + 3000*\mathbf{X}_2 + 1000*\mathbf{X}_3 + 1000*\mathbf{X}_4 + 1000*\mathbf{X}_5 + 1000*\mathbf{X}_6$$

In the spreadsheet of Figure 5.5, the range B34:G534 contains data from 501 simulations of the six funds' growth ratios. With the investment levels listed in cells B27:G27, the portfolio value **W** that would be returned in each of these simulations is computed in cells I34:I534 of this spreadsheet by entering the formula

```
=SUMPRODUCT(B34:G34,$B$27:$G$27)
```

into cell I34 and copying I34 to I34:I534. The sample average of these simulated portfolio values is computed in cell I28 by the formula

```
=AVERAGE(I34:I534)
```

This sample average (\$11,353) can be interpreted as an estimate of $E(\mathbf{W})$, the expected value of the portfolio next year. The sample standard deviation of these simulated portfolio values is computed in cell I29 by the formula

```
=STDEV(I34:I534)
```

This sample standard deviation (\$1491) can be used to estimate $Stdev(\mathbf{W})$, the standard deviation of the portfolio's value next year. With \$10,000 (the

Myerson_C05.qxd 09-09-2004 10:34 Page 172

current investment value) entered into cell G28, the fraction of the simulated portfolio values that are less than $10,000 is computed in cell I30 by the formula

=PERCENTRANK(I34:I534,G28)

This fraction (0.188) can be used to estimate $P(\mathbf{W}<10000)$, the probability that this portfolio will be worth less than $10,000 next year. The value that exceeds 5% of simulated portfolio values is computed in cell I31 by the formula

=PERCENTILE(I34:I534,0.05)

So we may estimate next year's portfolio value **W** has about a 0.05 cumulative probability of being less than this 0.05 percentile value ($8877). With a risk tolerance of $6000 entered into cell I25, the formula

=CE(I34:I534,I25)

in cell I32 yields as an estimate ($11,169) of the certainty equivalent of this risky portfolio for an investor who has a constant risk tolerance of $6000.

Notice that these calculations in I28:I32 of Figure 5.5 are just statistical estimates that may be some distance from the actual expected values and probabilities for this random variable **W**. We could get more accurate estimates by using more simulation data. But by using equations (3) and (5) from Section 5.3, we can exactly compute the expected value $E(\mathbf{W})$ and standard deviation $Stdev(\mathbf{W})$ of the portfolio's future value. Then using the Multivariate Normality assumption, we can exactly compute the portfolio's cumulative probability of being less than any amount, and we can exactly compute the portfolio's certainty equivalent for an investor with any constant risk tolerance. These exact calculations are shown in Figure 5.8.

The means, standard deviations, and correlations that we are assuming (based on the 1980s data) for the annual growth ratios of the six stock funds are shown in the top 12 rows of Figure 5.8. The funds' means are in cells B3:G3, their standard deviations are in cells B4:G4, and their $6 \times 6$ correlation array is in cells B7:G12. (All these parameter values are taken from the calculations shown in Figure 5.5.) Multivariate Normal random variables with these means, standard deviations, and correlations have been constructed in cells B16:G16 of Figure 5.8, using NORMINV formulas driven by a CORAND array in cells B14:G14.

Cells B20:G20 in Figure 5.8 list the levels of investment in each fund, that is, the *shares* of each fund in the portfolio that we are considering. The future value of this portfolio,

$$\mathbf{W} = 3000*\mathbf{X}_1 + 3000*\mathbf{X}_2 + 1000*\mathbf{X}_3 + 1000*\mathbf{X}_4 + 1000*\mathbf{X}_5 + 1000*\mathbf{X}_6$$

can be simulated in cell I22 here by the formula

=SUMPRODUCT(B20:G20,B17:G17)

| | A | B | C | D | E | F | G | H | I | J |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | MultivariateNormal Parameters of Fund Returns | | | | | | | | | |
| 2 | | Fund 1 | Fund 2 | Fund 3 | Fund 4 | Fund 5 | Fund 6 | | | |
| 3 | Means | 1.138 | 1.155 | 1.092 | 1.106 | 1.126 | 1.129 | | | |
| 4 | StDevs | 0.152 | 0.192 | 0.115 | 0.162 | 0.162 | 0.140 | | | |
| 5 | Correlation coefficients | | | | | | | | | |
| 6 | | Fund 1 | Fund 2 | Fund 3 | Fund 4 | Fund 5 | Fund 6 | | | |
| 7 | Fund 1 | 1 | 0.8937 | 0.5980 | 0.7280 | 0.9178 | 0.9846 | | | |
| 8 | Fund 2 | 0.8937 | 1 | 0.5389 | 0.7404 | 0.7216 | 0.8804 | | | |
| 9 | Fund 3 | 0.5980 | 0.5389 | 1 | 0.8992 | 0.7627 | 0.7260 | | | |
| 10 | Fund 4 | 0.7280 | 0.7404 | 0.8992 | 1 | 0.7597 | 0.8109 | | | |
| 11 | Fund 5 | 0.9178 | 0.7216 | 0.7627 | 0.7597 | 1 | 0.9550 | | | |
| 12 | Fund 6 | 0.9846 | 0.8804 | 0.7260 | 0.8109 | 0.9550 | 1 | | | |
| 13 | | | | | | | | | | |
| 14 | CORANDs | 0.757 | 0.655 | 0.634 | 0.682 | 0.802 | 0.752 | | | |
| 15 | Multivariate Normals with above parameters | | | | | | | | | |
| 16 | | 1.24 | 1.23 | 1.13 | 1.18 | 1.26 | 1.22 | | | |
| 17 | | | | | | | | | | |
| 18 | Portfolio: $ to invest now, for W next year | | | | | | | | | |
| 19 | | Fund 1 | Fund 2 | Fund 3 | Fund 4 | Fund 5 | Fund 6 | | Total $ now | |
| 20 | | 3000 | 3000 | 1000 | 1000 | 1000 | 1000 | | 10000 | |
| 21 | | | | | | Sim of portfolio $ next year (W) | | | | |
| 22 | | | | | | | | | 12226.26 | |
| 23 | FORMULAS | | | | | | Computed from parameters | | | |
| 24 | B14:G14. {=CORAND(B7:G12)} | | | | | | | | 11328.90 | E(W) |
| 25 | B16. =NORMINV(B14,B3,B4) | | | | | | | | 1490.50 | Stdev(W) |
| 26 | B16 copied to B16:G16 | | | | | | | | 0.186 | P(W<I20) |
| 27 | I20. =SUM(B20:G20) | | | | | RiskTol | | | 8877.25 | 5%ile(W) |
| 28 | I22. =SUMPRODUCT(B20:G20,B16:G16) | | | | | 6000 | | | 11143.77 | CE |
| 29 | I24. =SUMPRODUCT(B20:G20,B3:G3) | | | | | | | | | |
| 30 | I25. =SUMPRODUCT(PRODS(B20:G20),PRODS(B4:G4),B7:G12)^0.5 | | | | | | | | | |
| 31 | I26. =NORMSDIST((I20-I24)/I25) | | | | | | | | | |
| 32 | I27. =NORMINV(0.05,I24,I25) | | | | | | | | | |
| 33 | I28. =I24-0.5*(I25^2)/G28 | | | | | | | | | |

FIGURE **5.8**
**Portfolio calculations for investments in the six funds**

Equation (3) tells us that the actual expected value of $\mathbf{W}$ is

$$E(\mathbf{W}) = 3000*E(\mathbf{X}_1) + 3000*E(\mathbf{X}_2) + 1000*E(\mathbf{X}_3) + 1000*E(\mathbf{X}_4) + 1000*E(\mathbf{X}_5) + 1000*E(\mathbf{X}_6)$$

Thus, the expected value $E(\mathbf{W})$ = \$11,328.90 can be computed in cell I24 of Figure 5.8 by the formula

=SUMPRODUCT(B20:G20,B3:G3)

because cells B20:G20 contain the shares of each fund in the portfolio and cells B3:G3 contain the expected future values (per unit of investment now) for each fund.

Now recall the big equation for a portfolio's standard deviation (5):

$$Stdev(\mathbf{W}) = \left(\sum_{i=1}^{n}\sum_{j=1}^{n} \alpha_i*\alpha_j*Stdev(\mathbf{X}_i)*Stdev(\mathbf{X}_j)*Correl(\mathbf{X}_i, \mathbf{X}_j)\right)\text{^}0.5$$

Here the $\alpha_i$ are the levels of investment or shares invested now in each fund, and the number of funds $n$ here is 6. So equation (5) asks us here to compute a sum of 36 (6*6) terms, each of which is the multiplicative product of a product of shares, a product of standard deviations, and a correlation. Recall that Simtools function PRODS generates square arrays of such

Myerson_C05.qxd 09-09-2004 10:34 Page 174

pairwise products. So as we saw at the end of Section 5.3, the standard deviation of the portfolio's future value can be computed in Excel with Simtools by a formula of the form

=SUMPRODUCT(PRODS(*shares*),PRODS(*stdevns*),*correls*)^0.5

In Figure 5.8, the *shares* that are being invested now in the various funds are listed in cells B20:G20, the *stdevns* of the funds' future values (per unit invested now) are listed in cells B4:G4, and the *correls* (correlation array) of the six funds' future values are listed in cells B7:G12. So the formula PRODS(B20:G20) returns a $6 \times 6$ array of all pairwise products of the shares ($\alpha_i * \alpha_j$), and the formula PRODS(B4:G4) returns a $6 \times 6$ array of all pairwise products of the funds' standard deviations ($Stdev(\mathbf{X}_i) * Stdev(\mathbf{X}_j)$). Thus, the portfolio's standard deviation $Stdev(\mathbf{W})$ = \$1490.50 can be computed in cell I25 of Figure 5.8 by the formula

=SUMPRODUCT(PRODS(B20:G20),PRODS(B4:G4),B7:G12)^0.5

The SUMPRODUCT function here multiplies corresponding terms in the three $6 \times 6$ arrays, and then adds up these 36 products. The square root (^0.5) at the end of the formula must not be forgotten; without it, we would get the variance instead.

Now because we are assuming that the funds' returns are jointly Multivariate Normal random variables, the total return of any portfolio (as a linear function of the funds' individual returns) is also Normal. Thus, once we know the expected value (or mean) and standard deviation of the portfolio's future value, we can compute the probability of the portfolio's future value being in any interval, using the NORMSDIST function as discussed in Chapter 4. With $E(\mathbf{W})$ in cell I24 and $Stdev(\mathbf{W})$ in cell I25, the probability $P(\mathbf{W} < 10000) = 0.186$ can be computed in cell I26 of Figure 5.8 by the formula

=NORMSDIST((10000–I24)/I25)

The value-at-risk that has 5% cumulative probability below it for this portfolio, that is, the portfolio's 0.05 percentile value (\$8877.25), can be computed in cell I27 by the formula

=NORMINV(0.05,I24,I25)

(Recall that NORMINV is the inverse cumulative function, and so it can be used for computing inverse cumulative values as well as for making random variables.) Because **W** is a Normal random variable, we can also use the special formula for computing certainty equivalents of Normal lotteries from Section 4.7. With the constant risk tolerance \$6000 entered into cell G28, the certainty equivalent of the random return **W** next year (\$11,143.77) is computed in cell I28 of Figure 5.8 by the formula

=I24-0.5*(I25^2)/G28

We should emphasize that cells I24:I28 of Figure 5.8 contain exact formulas for these expectations and probabilities, based on an assumption of

Myerson_C05.qxd 09-09-2004 10:34 Page 175

Multivariate Normality, whereas cells I28:I32 of Figure 5.5 contain statistical estimates of these expectations and probabilities, based on a sample of 501 simulations.

Let us now discuss the Normality assumption for a bit. In fact, some sort of apology should be offered for this assumption, because Chapter 4 previously suggested that growth ratios of financial assets should be considered Lognormal random variables in our models. (It was suggested that logarithmic growth rates could be considered Normal, which would make the growth ratios Lognormal.) But for each of the six funds' growth ratios in this example, the mean is more than five times the standard deviation (compare rows 3 and 4 in Figure 5.8, or rows 14 and 15 in Figure 5.5), and in such cases the Lognormal and Normal distributions are almost the same. So it is not very inconsistent for us to use Normal distributions instead of Lognormals here. Indeed, such high ratios of mean over standard deviation are common in short-term financial analysis, where the values of portfolios are being evaluated a year or less from now, and so the Normality assumption is commonly used in such analyses. It typically becomes more important to use Lognormal growth ratios in long-term financial analysis, where we consider the future values of investments after several years.

What do we gain by assuming Normality here? We could make Lognormal random variables with the appropriate means, standard deviations, and correlations by simply replacing the NORMINV functions by LNORMINV functions in cells B33:G33 of Figure 5.5 and in cells B17:G17 of Figure 5.8. If we made this replacement and regenerated the simulation data in Figure 5.5, all the simulation estimates in cells I28:I32 would still be equally valid for the new Lognormal model. Furthermore, our exact calculations of $E(\mathbf{W})$ and $Stdev(\mathbf{W})$ in cells I24 and I25 of Figure 5.8 would also be equally valid for this new Lognormal model, because equations (3) and (5) did not require any Normality assumption. But the probability and certainty equivalent calculations in cells I26:I28 of Figure 5.8 would not be valid without the Normality assumption, and this is what we gain from Normality. Even if we replaced the NORMINV in cell I27 with an LNORMINV function, we still would not have a valid calculation, because an assumption that the individual funds' future values have Lognormal probability distributions would not imply that our portfolio's future value has a Lognormal probability distribution.

In fact, Multivariate Normality is the only mathematical assumption that statisticians have to guarantee that the mean and standard deviation of any portfolio's future value fully characterizes the whole probability distribution of the portfolio's future value. We do not have any similarly nice theory of "Multivariate Lognormality." So financial analysts like to assume Multivariate Normal asset values in short-term financial models. When they switch to a more appropriate Lognormality assumption (or any other distributional assumption) in longer-term financial analysis, probability and certainty-equivalent calculations like those in cells I26:I28 of Figure 5.8 must be replaced by simulation estimates like those in cells I30:I32 of Figure 5.5.

Myerson_C05.qxd 09-09-2004 10:34 Page 176

## 5.7 Excel Solver and Efficient Portfolio Design

All the portfolio analysis in Figure 5.5 and Figure 5.8 has been set up to depend on the investment levels in the portfolio (listed in cells B27:G27 of Figure 5.5 or cells B20:G20 of Figure 5.8). So we can now apply these models to compare different portfolios and answer questions about portfolio design.

Suppose that we have been hired as consultants to advise an investor who has $10,000 to invest for the coming year, and who is considering the six funds that we modeled in the previous section. We might think about suggesting the investment plan shown in cells B20:G20 of Figure 5.8, where the dollar amounts invested now in Funds 1 through 6 are, respectively

3000, 3000, 1000, 1000, 1000, 1000

Next year's value of this investment portfolio would have expected value and standard deviation

$E$ = \$11,329 and *Stdev* = \$1490

as shown in cells I24 and I25 of Figure 5.8. If we were to recommend this portfolio, we might point out that it offers an expected rate of return better than 13%, and that putting the whole $10,000 in either of the two individual funds that offer better expected rates of return would be more risky. (Putting the whole $10,000 in Fund 1 would yield standard deviation $1520. Putting the whole $10,000 in Fund 2 would yield standard deviation $1920.) We are assuming here that this investor is risk averse and does not simply want to maximize the expected portfolio value regardless of risk. That is, the investor might not prefer a portfolio that had a higher expected value if it also involved substantially more risk, as measured by the standard deviation.

But, in response to such a recommendation, our client might well ask if there are any other portfolios that could yield the same expected returns with an even lower standard deviation. If so, the client might ask, which of these portfolios would be the least risky? With this question, we are facing an analytical problem that is quite different from what we have done up to now in our spreadsheets. If our client had asked about any specific portfolio, we could have easily evaluated it by entering the new investment levels into cells B20:G20 of Figure 5.8 and reading the expected value and standard deviation of next year's returns in cells I24 and I25 of this spreadsheet. But the client is asking us to search among all possible ways of investing the $10,000 and to find the one that has lowest standard deviation, subject to the constraint that it must offer expected return not lower than $11,329. Such problems are called *optimization problems*.

Microsoft Excel comes with an add-in called Solver.xla that can solve optimization problems in spreadsheets. The default installation of Excel often does not install this xla file, and so you may need to redo a custom

Myerson_C05.qxd 09-09-2004 10:34 Page 177

installation to get Solver.xla into the appropriate library folder on your computer's hard disk. Once it is there, you need to use the Tools:AddIns menu to activate Solver.xla as a regular part of Excel. Once installed and added in, Solver should appear as a part of your Tools menu in Excel.

In the spreadsheet of Figure 5.8, we can apply Solver to answer our investor's question about the least risky portfolio with the given expected return. With this spreadsheet activated, we launch Solver by the menu command sequence Tools:Solver, which opens the Solver-Parameters dialog box. This dialog box looks complicated, but it basically requires us to specify three things.

First, we must specify a goal or target that we are directing Solver to pursue, by specifying in the Solver-Parameters box a *target cell,* which is either to be *maximized*, or to be *minimized*, or to be made equal to some *value*. In this case, we have been asked to look for a portfolio that has as little risk as possible, and so we should select as our target cell I25 in Figure 5.8, where the standard deviation of the portfolio's value is computed. We want to make this standard deviation as small as possible, and so we should select the Minimize option for this target cell.

Next, we must specify the cells that Solver is allowed to change or adjust in pursuit of this minimization target. We do so by clicking on the *changing cells* line (just above the middle of the Solver-Parameters box) and selecting there (by typing an address or by selecting with the mouse) the range of cells that Solver is allowed to change or adjust. Each of these cells must currently contain a simple numerical value, not a formula. In this case, the numbers that we want Solver to consider changing are the planned investments in the six stock funds, which have been entered into cells B20:G20 in Figure 5.8. So we should specify the range B20:G20 as the changing cells that we want Solver to adjust.

Finally, we must specify what constraints, if any, should restrict the values that Solver is allowed to use in the changing cells. In these constraints, various cells can be required to be less than some values, or equal to some values, or greater than some values. Specifying the constraints is often the hardest part in using Solver, because you can get nonsensical results if you misspecify your constraints. In this case, our mission is to invest $10,000 in these six funds, and we are not interested in any investment plan that allocates more or less money. So we should require Solver, as it changes B20:G20, to keep the sum of the six investments equal to $10,000. This sum of B20:G20 is computed by the formula in cell I20 in Figure 5.8, and so we should give Solver the constraint

I20 = 10000

To enter this Solver constraint, click the Add button in the Solver-Parameters dialog box, which launches the Add-Constraint dialog box. Then enter the address I20 in the left-hand side of this Add-Constraint box, select the "=" option in the drop-down option list in the center of the Add-Constraint box, and enter the number 10000 in the right-hand side of this

Myerson_C05.qxd 09-09-2004 10:34 Page 178

Add-Constraint box. Then click OK to close the Add-Constraint box and return to the Solver-Parameters box.

We had been asked to find the least risky portfolio that has an expected return not less than the $11,329 offered by the portfolio shown in Figure 5.8. The expected value returned by any new B20:G20 portfolio would be computed in cell I24 in Figure 5.8. So we should also give Solver the constraint

I24 >= 11329

To do so in the Solver-Parameters dialog box, click the Add button to launch the Add-Constraint dialog box again, and then enter the address I24 in the left-hand side of the Add-Constraint box, select the ">=" option in the center of the box, and enter the number 11329 in the right-hand side of the box. Then click OK to return again to the Solver-Parameters box.

So now we have specified the following Solver parameters:

Minimize the target cell I25 by changing cells B20:G20

subject to the constraints I20 = 10000, I24 >= 11329

Now we can click the Solve button in the Solver-Parameters box, and Solver will get to work, searching for new ways to invest $10,000 that can yield lower risk without decreasing our expected value.

After some delay, the Solver-Results dialog box should announce, "Solver has found a solution. All constraints and optimality are satisfied." This is good news, that Solver thinks that it has done its best to solve the problem that we gave it. Then with "Keep Solver Solution" checked, we should click OK to see Solver's solution. Unfortunately, in this case the solution may be one that our client will not appreciate, because the investments in Funds 1 through 6 that Solver has recommended in cells B20:G20 are, respectively,

–17512, –1256, 2313, –10295, –18211, 54961

The negative numbers in this solution might be interpreted as instructions to sell short Funds 1, 2, 4, and 5. In simpler terms, Solver is suggesting that we should make a net investment of $54,961 in Fund 6 and $2,313 in Fund 3, which adds up to $57,274, of which $10,000 will come from the money that our client actually has, and the extra $47,274 will be raised by selling shares in Fund 1 that are worth $17,512, shares in Fund 2 that are worth $1256, shares in Fund 4 that are worth $10,295, and shares in Fund 5 that are worth $18,211. If our investor does not have these shares to sell, then perhaps some friend might loan them, in exchange for a written promise to return them next year, which our investor could do by buying back the borrowed shares out of the returns from the sale of the positive investments. If our investor does not have a friend (or broker) willing to make such a nice loan, then we must admit that a solution with negative values in cells B20:G20 is not really feasible for our investor.

Myerson_C05.qxd 09-09-2004 10:34 Page 179

To prevent such negative values in cells B20:G20, we should launch Solver again, go back to the Add-Constraint dialog box, and enter as our third constraint

B20:G20 >= 0

(Solver will accept a constraint with a range of cells on the left-hand side and a single number on the right-hand side, in which case each cell's value is required to be above or below the given number.) Now when we click Solver's Solve button, it returns the following solution in cells B20:G20:

6902, 1518, 1580, 0, 0, 0

This solution tells us that we can reduce our standard deviation to a minimum of $1432 with nonnegative investments totaling $10,000 now, and with an expected return next year that is not less than $11,329, by investing $6902 in Fund 1, $1518 in Fund 2, and $1580 in Fund 3.

Thus, Solver can be used to design efficiently diversified portfolios that reduce financial risks. But this is just one example to illustrate the remarkable power of Solver to search for ways to improve a target cell by changing other cells, subject to well-specified constraints. We should consider other variations on this optimization problem with other targets that may be of interest and different constraints.

For example, suppose that our investor becomes a bit more greedy and asks for the least risky portfolio that would offer an expected return (in cell I24) of at least $12,000 next year (instead of just $11,329). With the spreadsheet in Figure 5.8 still activated, we can launch Solver again, select the I24 >= 11329 constraint, click the "Change" button, and change the constraint to

I24 >= 12000

But now when we ask Solver to solve this new optimization problem, it will announce in its Solver-Results box that "Solver could not find a feasible solution." That is, our optimization problem is *infeasible*, which means that our constraints are impossible to satisfy and must be relaxed in some way. With nonnegative investments that sum to $10,000, we cannot get a higher expected return than $11,548, which is achieved by putting the whole $10,000 into Fund 2. To verify this fact, you should launch Solver again and give it the optimization problem:

maximize the target I24 (the expected return) by changing cells

B20:G20

subject to the constraints I20 = 10000, B20:G20 >= 0

(You can eliminate a selected constraint by using the Delete button in the Solver-Parameters dialog box.) For this new problem, Solver should return the solution

0, 10000, 0, 0, 0, 0

in cells B20:G20.

Myerson_C05.qxd 09-09-2004 10:34 Page 180

To see one of the more puzzling outcomes that you can get from running Solver, try the following optimization problem in the spreadsheet from Figure 5.8:

Maximize the target cell I24 by changing cells B20:G20

subject to the constraint I20 = 10000

That is, we are asking Solver to find the highest possible expected return for a portfolio that costs $10,000 now, but now we are allowing some funds to have negative investments (selling them short) because the nonnegativity constraint has been deleted. When you ask Solver to solve this problem, the Solver-Results box will report "The Set Target Cell values do not converge." This obscure message actually means that Solver has found an *unbounded optimal solution*. That is, Solver has found a way to go on improving your target cell infinitely. In this case, Solver is telling us that it can find ways to earn expected returns in the millions or billions of dollars next year, with a net investment of only $10,000 now. So whenever you get this "Set Target Cell values do not converge" message, it means either that Solver has found a way for you to make an infinite amount of money, or else (as is more likely) you have misspecified the problem. In this case, the problem is that we are ignoring risk and credit constraints by assuming that we can raise millions of dollars now by short selling the low-expected-return funds (such as Fund 3) and that we can invest these millions into the high-expected-return funds (such as Fund 2). That is, if our optimization problem is to describe the investor's real situation, then it must include some constraint that rules out such enormous risk taking with funds raised from short sales or loans.

Cell I28 in Figure 5.8 computes the investor's certainty equivalent for the portfolio next under the assumption that the investor has a utility function with constant risk tolerance of $6000. So the portfolio with the highest certainty equivalent can be found in this spreadsheet by asking Solver to maximize this target cell I28 by changing cells B20:G20 subject to the constraints B20:G20 >= 0 and I20 = 10000. Solver's optimal solution yields a calculated certainty equivalent of $11,240, which is achieved by investing now about $585 in Fund 1, $9415 in Fund 2, and $0 in the other funds. So with this risk tolerance, an investor should dilute the risks of Fund 2 by some diversification into Fund 1.

Now there is one more thing you need to know about Solver, which may explain why Microsoft omits it from the default installation of Excel: Unlike most other features in Excel, Solver can often give the wrong answer. To illustrate the problem, consider the simple spreadsheet shown in Figure 5.9. In a blank spreadsheet, let us start by entering any positive value (say the value 1) into cell A1. Then enter the formula =A1^2 into cell B1. Now launch Solver and give it the optimization problem

Maximize the target cell B1 by changing cell A1

subject to the constraints A1 <= 5, A1 >= –10

Myerson_C05.qxd 09-09-2004 10:34 Page 181

| | A | B | C | D |
|---|---|---|---|---|
| 1 | 5 | 25 | | |
| 2 | | | | |
| 3 | FORMULAS FROM RANGE A1:B1 | | | |
| 4 | B1. =A1^2 | | | |
| 5 | | | | |
| 6 | Solver: Maximize B1 by changing A1 | | | |
| 7 | subject to A1<=5, A1>=-10. | | | |

FIGURE **5.9**
**A simple example of how Solver can fail**

Solver will find and report the solution that is shown in Figure 5.9, where the value of cell A1 has been changed to 5, yielding the target-cell value of 25. But we could get a much higher target-cell value of 100 if we entered the value –10 into cell A1, which also satisfies the required constraints. Why did Solver miss this obviously better solution?

To understand what went wrong, you need to know something about how Solver works. Whenever we launch Solver, it starts with some numerical values that we have already entered into the changing cells. Suppose first that our initial values do not violate any constraints. Then Solver considers a variety of very small changes in these changing-cell values, looking for ways to make small changes that will improve the target-cell value without violating any constraints. If it finds any way to do this, then it continues making larger changes in the same direction for as far as it can get further improvement in the target cell without violating any constraints. Then Solver repeats this process, looking around the new changing-cell values to find a direction in which small changes can improve the target-cell value without violating any constraints, and it again moves the changing-cell values in this new direction as long as further improvement is possible. When Solver finds changing-cell values from which there are no small changes that could offer improvement without violating constraints, then it stops and reports that an optimum has been found. Such a solution where small feasible changes cannot improve the target is called a *local optimal solution* or a *local optimum*. (If our initial changing-cell values do not satisfy the constraints, then Solver works similarly to minimize the violation of the constraints, turning to target-cell improvement after all constraint violations have been eliminated.)

For the example in Figure 5.9, Solver has found a local optimum at A1 = 5, where small increases in A1 = 5 would not be allowed and small decreases below A1 = 5 would not improve the target. But this local optimum is not really the optimal solution because decreasing A1 by a larger amount (subtracting 10 or more) could increase the target cell without violating any constraints. And indeed, if we had started with any negative value in A1 when we launched Solver, then it would have moved swiftly to the true optimal solution at A1 = –10.

Solver may be compared to a robot that searches for the highest point in the world by continually walking in an uphill direction until it reaches the top of a hill or mountain, where it stops. Such a robot is not likely to find the top of Mt. Everest unless it starts in the right part of Nepal or Tibet!

Myerson_C05.qxd 09-09-2004 10:34 Page 182

Some optimization problems may resemble a smooth landscape dominated by one large hill, and Solver will be reliably successful in such cases. But other optimization problems may resemble a rough landscape with many separate mountain ranges, and then Solver's results may be very different, depending critically on its initial position.

So whenever you use Solver to analyze an optimization problem, you should run Solver several times with a variety of different initial values for the changing cells. If Solver always comes back reporting the same optimal solution, then you may have some confidence that it is truly optimal. But if Solver's optimal solution changes each time, then you should run it many times and keep a record of the changing-cell and target-cell values for the best solutions that it reports. Unless the optimization problem has some special mathematical structure, however, you might never be sure that one more application of Solver would not find an even better solution.

## 5.8 Subjective Assessment of Correlations

In Section 4.5, we introduced subjective assessment of quartiles as a practical way of assessing probability distributions for unknown quantities, even when relevant data for statistical analysis is not available. But quartile boundary points, like means and standard deviations, only measure beliefs about one unknown quantity at a time. It is just as important to measure beliefs about the relationships between unknown quantities, such as are expressed by their pairwise correlations. So in this section, we introduce a practical method for subjectively assessing correlations.

Let us consider an example like that of Section 5.1, where **X** and **Y** represent the prices next year (in dollars per share) of two stocks that our investor owns. But unlike Section 5.1, where a simple discrete probability distribution was rather unrealistically assumed, let us assume now that the investor's beliefs about these future stock prices can be described by continuous probability distributions in the Generalized Lognormal family (which includes the Normals). To build a model that represents these beliefs about the unknown quantities **X** and **Y**, we may begin by asking her questions to assess her quartile boundary points for each of these unknowns. To be specific, suppose that this investor would assess equally likely quartiles with boundary points 73, 75.5, and 78 for **X**, and she would assess equally likely quartiles with boundary points 30.8, 31.6, and 32.4 for **Y**. (These quartiles happen to be compatible with **X** and **Y** having Normal distributions with the same means, 75.5 and 31.6, and the same standard deviations, 3.5 and 1.2, as the discrete distributions that we assumed in Section 5.1.) These quartile boundaries for **X** and **Y** have been entered in cells B5:B7 and C5:C7 of Figure 5.10.

Now suppose that our plan is to simulate **X** and **Y** by Generalized Lognormal random variables that are driven (in the first parameter for their

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | WORKSHEET FOR SUBJECTIVE ASSESSMENT OF CORRELATIONS | | | | | | |
| 2 | Suppose X and Y have Generalized-Lognormal distributions | | | | | | |
| 3 | with quartile boundary points as listed below: | | | | | | |
| 4 | | X | Y | | | | |
| 5 | Q1 (0.25) | 73 | 30.8 | | | | |
| 6 | Q2 (0.50) | 75.5 | 31.6 | | | | |
| 7 | Q3 (0.75) | 78 | 32.4 | | | | |
| 8 | | | | | | | |
| 9 | When Correlation(X,Y) = | | | 0.9 | | | |
| 10 | 0.27191169 | = Conditional median of one CORAND given other is 0.25 | | | | | |
| 11 | 0.72808831 | = Conditional median of one CORAND given other is 0.75 | | | | | |
| 12 | | X | | | | | |
| 13 | | 73.24999 | = Conditional median of X given Y=30.8 | | | | |
| 14 | | 77.75001 | = Conditional median of X given Y=32.4 | | | | |
| 15 | | | Y | | | | |
| 16 | | | 30.88 | = Conditional median of Y given X=73 | | | |
| 17 | | | 32.32 | = Conditional median of Y given X=78 | | | |
| 18 | Simulation model: | | | | | | |
| 19 | (corands) | 0.048584 | 0.068766 | | | | |
| 20 | | X | Y | | | | |
| 21 | | 69.35188 | 29.83862 | | | | |
| 22 | | | | | | | |
| 23 | FORMULAS | | | | | | |
| 24 | A10. =MIDRAND($D$9,0.25) | | | | | | |
| 25 | A11. =MIDRAND($D$9,0.75) | | | | | | |
| 26 | B13. =GENLINV(A10,B5,B6,B7) | | | C16. =GENLINV(A10,C5,C6,C7) | | | |
| 27 | B14. =GENLINV(A11,B5,B6,B7) | | | C17. =GENLINV(A11,C5,C6,C7) | | | |
| 28 | B19:C19. {=CORAND(D9)} | | | | | | |
| 29 | B21. =GENLINV(B19,B5,B6,B7) | | | C21. =GENLINV(C19,C5,C6,C7) | | | |
| 30 | C13. =" = Conditional median of X given Y="&C5 | | | | | | |
| 31 | C14. =" = Conditional median of X given Y="&C7 | | | | | | |
| 32 | D16. =" = Conditional median of Y given X="&B5 | | | | | | |
| 33 | D17. =" = Conditional median of Y given X="&B7 | | | | | | |

FIGURE **5.10**
**Computing conditional medians for subjective correlation assessment**

GENLINV formulas) by a pair of linked CORAND values. Such a model is shown in cells B21:C21 of Figure 5.10, which contain the formulas

=GENLINV(B19,B5,B6,B7) and =GENLINV(C19,C5,C6,C7)

In this spreadsheet, the range B19:C19 contains the array formula

{=CORAND(D9)}

So in this model, B21 is a Generalized Lognormal random variable with the quartiles that we have assessed for **X**, and C21 is a Generalized Lognormal random variable with the quartiles that we have assessed for **Y**. (Actually, both are Normal random variables because their quartile boundaries satisfy the equal differences property.) The CORAND array makes these two random variables correlated in a way that is driven by the number that we enter into cell D9 of this spreadsheet. But what correlation number should we enter there? We need to ask our investor some question that can be used to estimate the correlation parameter for these CORANDs.

To assess the relationship between these two unknown quantities, we should ask some question about how her beliefs about one unknown quantity

Myerson_C05.qxd 09-09-2004 10:34 Page 184

would be affected by information about the other quantity. If these unknown quantities were believed to be independent, then learning about one would have no effect on her beliefs about the other, but if they are highly correlated, then learning about one would greatly change her beliefs about the other. One such question that we may ask our investor is:

> If you learned that the quantity **X** was actually equal to the value that you have assessed for **X**'s upper quartile boundary point, but you got no additional information about the unknown **Y**, then for what number $\hat{y}$ would you say that **Y** was equally likely to be above or below $\hat{y}$?

That is, if we let $x_3$ denote this upper quartile boundary point for **X**, then we are asking for a number $\hat{y}$ such that

$$P(\mathbf{Y} < \hat{y} | \mathbf{X} = x_3) = 1/2$$

This number $\hat{y}$ can be called the *conditional median* of **Y** given that **X** is equal to $x_3$.

In our example, we might put this question to the investor as follows:

> You have told us that the price of the second stock next year is equally likely to be above or below \$31.60 per share. But now suppose for a moment that the price per share of the first stock turned out next year to be \$78. Given this good news about the first stock, you might revise your beliefs about the second stock. So if you knew that first stock's price would be \$78, for what number would $\hat{y}$ you think that the second stock's price was equally likely to be above or below $\hat{y}$ ?

Suppose that, after some thought, the investor responded to this question by reporting 32 as her subjectively assessed conditional median of **Y** given **X** = 78. The fact that this conditional median of \$32.00 is greater than the unconditional median of \$31.60 indicates that, in this investor's opinion, a stock market where the first stock performs well is more likely to be a stock market where the second stock also performs better than expected.

Simtools provides a function, MIDRAND, that can be used to compute the appropriate correlation parameter for our CORANDs from such a conditional median assessment. When a pair of CORAND values is generated with some correlation parameter ρ, then for any number $q$ between 0 and 1, the Simtools formula

MIDRAND(ρ, $q$)

gives us the conditional median for one of these CORAND values given that the other CORAND is equal to $q$. That is, when one ρ correlated CORAND value is very close to $q$, the other value is equally likely to be above or below MIDRAND(ρ, $q$). Notice that the first parameter of MIDRAND is the correlation and the second parameter of MIDRAND is the given value of the other CORAND. (This MIDRAND value could also be equivalently computed by the Excel formula NORMSDIST(ρ∗NORMINV($q$, 0, 1)).)

Myerson_C05.qxd 09-09-2004 10:34 Page 185

In Figure 5.10, cell A11 contains the formula

=MIDRAND($D$9,0.75)

When cell D9 is equal to 0.9, this formula returns the value 0.728 in cell A11. So if we made a table of thousands of simulations of the CORANDs in cells B19:C19 with the correlation parameter D9 equal to 0.9, and if we then selected from this simulation data the rows where the B cell had a value very close to 0.75, then the value of the C cell should be less than cell A11 (0.728) in about half of these selected simulations. (To get 200 rows where the value of B19 was between 0.74 and 0.76, we would need about 10,000 total simulations.) But these CORANDs in B19:C19 are being used as the random first parameter driving the simulated values of **X** and **Y** in cells B21 and C21. So if the CORAND value in cell B19 were near 0.75, then the simulated value of **X** in cell B21 would be near

GENLINV(0.75, 73, 75.5, 78) = 78

because 78 is the quartile boundary point with cumulative probability 0.75 for the unknown quantity **X**. But any CORAND value in cell C19 that is less than the MIDRAND value in cell A11 would generate a simulated value of **Y** in cell C21 that is less than

GENLINV(A11, 30.8, 31.6, 32.4)

which is computed in cell C17 of Figure 5.10. When A11 is 0.728, this formula in cell C17 takes the value 32.32. So if the correlation parameter in our model were 0.9, then having **X** near 78 would imply that **Y** was equally likely to be above or below $32.32. Thus, this value of cell C17 (32.32) would be the conditional median of **Y** given **X** = 78 if the correlation parameter were 0.9.

But our investor subjectively assessed $32.00 as the conditional median of **Y** given **X** = 78, and so the correlation 0.9 is too high. By adjusting the value of cell D9 in this spreadsheet and watching how cell C11 changes, you can verify that the correlation D9 = 0.5 makes the computed conditional median in cell C17 very close to the subjectively assessed conditional median of 32.00. With the spreadsheet in Figure 5.10, we could ask Excel's GoalSeek tool to set cell C17 to the value 32 by changing the correlation parameter in cell D9, and GoalSeek would return the value 0.5 in cell D9. (Try the menu command Tools:GoalSeek.) So entering a correlation parameter of 0.5 into cell D9 of Figure 5.10 would give us simulated **X** and **Y** values in cells B21 and C21 (driven by the CORANDs in B19:C19) that fit this investor's subjective assessments of conditional median and quartiles.

Cell C16 in Figure 5.10 similarly computes the conditional median of our simulated **Y** when the simulated **X** is very close to 73 (the quartile point with cumulative probability 0.25 for the unknown quantity **X**). The formula in cell C16 is

=GENLINV(A10,C5,C6,C7)

Myerson_C05.qxd 09-09-2004 10:34 Page 186

where cell A10 contains the formula

=MIDRAND($D$9,0.25)

When the correlation in cell D9 is 0.5, the value of cell C16 is 31.2. So to check our correlation assessment, we might ask the expert, "If I had told you the first stock's price next year would be only $73, might you say then that the second stock's price would be just as likely to be below $31.20 as to be above $31.20?" If the investor felt comfortable with this statement, then we would have further support for estimating 0.5 to be the correlation between these two stock prices. Otherwise, we might try to modify the correlation in D9 until the computed values in C16 and C17 both seem approximately correct to the expert as conditional medians of **Y** given the **X** values of $73 and $78, respectively.

The spreadsheet in Figure 5.10 can be applied to assess correlation parameters for CORAND in simulation models where the random variables have any Generalized Lognormal distributions, even when they are not Normal. But the correlation parameter in CORAND and the correlation that is statistically estimated by CORREL may have some subtle differences in such non-Normal models, as discussed in the following section.

## 5.9 Using CORAND with Non-Normal Random Variables

We have learned to make many types of random variables (Normal, Lognormal, Triangular, Gamma, etc.) by using RAND( )s as the first parameter in the appropriate inverse-cumulative function (NORMINV, LNORMINV, GENLINV, TRIANINV, etc.). In each case, we can use a CORAND value in place of the RAND( ) value, and we will get a random variable drawn from the same probability distribution because each CORAND value has the same uniform distribution as a RAND( ). But with CORAND, these random variables can be correlated (positively or negatively) with other random variables that are made using other random values in the same CORAND array. So CORAND can be used as a versatile tool for creating nonindependent random variables with different kinds of probability distributions.

But when CORAND values are used in this way to make random variables, the correlation of the resulting random variables can be guaranteed to equal the correlation parameter in the CORAND function only when we are making Normal random variables. For non-Normal random variables, the actual correlation of the resulting random variables may be somewhat different from the correlation parameter of the CORANDs. In particular, when a pair of positively correlated CORANDs are used to make two random variables that are skewed in opposite directions, these random variables may have a correlation that is substantially smaller than the CORANDs' parameter.

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Correlation parameter | | | | Quartile parameters | | | | | | FORMULAS | | | | |
| 2 | 0.9 | | | Q1 | -2 | -3 | | | | | C9:D9. {=CORAND(A2)} | | | | |
| 3 | | | | Q2 | 0 | 0 | | | | | E9. =GENLINV(C9,E2,E3,E4) | | | | |
| 4 | | | | Q3 | 3 | 2 | | | | | F9. =GENLINV(D9,F2,F3,F4) | | | | |
| 5 | | | | | | | | Normalized rank | | | SimTable data fills C10:F110. | | | | |
| 6 | | | CORREL(Corands) | | CORREL(GenLogNmls) | | | correlation | | | D7. =CORREL(C10:C110,D10:D110) | | | | |
| 7 | | | | 0.8963 | | 0.6963 | | | 0.8956 | | F7. =CORREL(E10:E110,F10:F110) | | | | |
| 8 | | | Corands | | GenLognormals | | | | H | | 10:H110. {=NORMIZE(E10:E110)} | | | | |
| 9 | SimTable | | 0.1975 | 0.2056 | -2.402 | -3.832 | | Normalized ranks | | | I10:I110. {=NORMIZE(F10:F110)} | | | | |
| 10 | MidRand | 0 | 0.1268 | 0.2475 | -2.98 | -3.042 | | -1.445 | -0.796 | | I7. =CORREL(H10:H110,I10:I110) | | | | |
| 11 | 0.0181 | 0.01 | 0.1437 | 0.1038 | -2.834 | -6.799 | | -1.378 | -1.316 | | A11. =MIDRAND($A$2,B11) | | | | |
| 12 | 0.0323 | 0.02 | 0.8658 | 0.6683 | 5.6713 | 1.3811 | | 0.9034 | 0.1999 | | A11 copied to A11:A109 | | | | |
| 13 | 0.0453 | 0.03 | 0.4989 | 0.715 | -0.01 | 1.7359 | | -0.124 | 0.3548 | | | | | | |
| 14 | 0.0576 | 0.04 | 0.8881 | 0.88 | 6.4674 | 3.0389 | | 1.0217 | 0.9413 | | | | | | |
| 15 | | | | | | | | | | | | | | | |
| 16 | | | | | | | | | | | | | | | |
| 17 | | | | | | | | | | | | | | | |
| 18 | | | | | | | | | | | | | | | |
| 19 | | | | | | | | | | | | | | | |
| 20 | | | | | | | | | | | | | | | |
| 21 | | | | | | | | | | | | | | | |
| 22 | | | | | | | | | | | | | | | |
| 23 | | | | | | | | | | | | | | | |
| 24 | | | | | | | | | | | | | | | |
| 25 | | | | | | | | | | | | | | | |
| 26 | | | | | | | | | | | | | | | |
| 27 | | | | | | | | | | | | | | | |
| 28 | | | | | | | | | | | | | | | |
| 29 | | | | | | | | | | | | | | | |
| 30 | 0.2244 | 0.2 | 0.931 | 0.9723 | 8.6356 | 4.1039 | | 1.3776 | 1.9642 | | Chart series: | | | | |
| 31 | 0.234 | 0.21 | 0.3343 | 0.5225 | -1.362 | 0.2001 | | -0.576 | -0.2 | | Left (C10:C110,D10:D110) | | | | |
| 32 | 0.2435 | 0.22 | 0.9193 | 0.8343 | 7.9241 | 2.6537 | | 1.2593 | 0.698 | | and (B11:B109,A11:A109) | | | | |
| 33 | 0.253 | 0.23 | 0.7371 | 0.9355 | 2.7851 | 3.591 | | 0.5187 | 1.4448 | | Middle (F10:F110,G10:G110) | | | | |
| 34 | 0.2625 | 0.24 | 0.892 | 0.8936 | 6.6221 | 3.1627 | | 1.0644 | 1.0644 | | Right (I10:I110,J10:J110) | | | | |

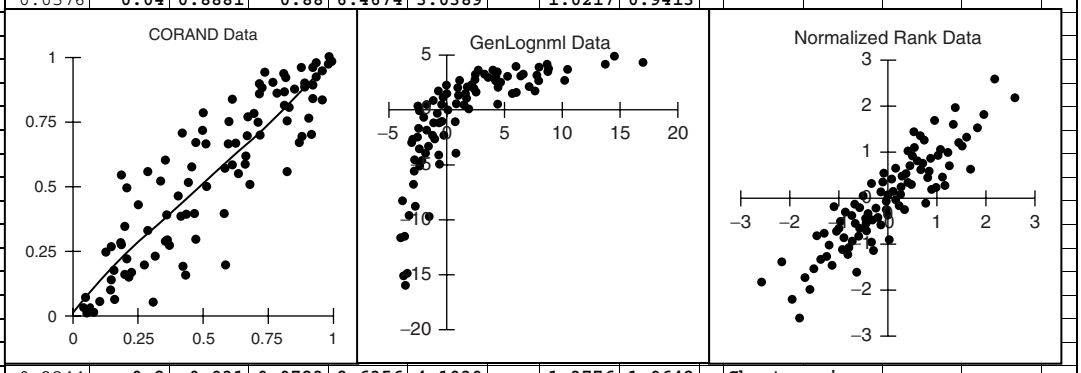


FIGURE **5.11** **Non-Normal random variables and normalized rank correlations**

To illustrate this effect, cells C9:D9 in Figure 5.11 contain the array formula

{=CORAND(A2)}

and cell A2 contains the value 0.9. The CORAND function was designed so that two Normal random variables that are driven by these two CORANDs would have the correlation 0.9. But cells E9 and F9 contain non-Normal random variables that are driven by the C9:D9 CORANDs with the formulas

=GENLINV(C9,E2,E3,E4) and =GENLINV(D9,F2,F3,F4)

The values in (E2, E3, E4) are (–2, 0, 3), and the values in (F2, F3, F4) are (–3, 0, 2). So these two Generalized Lognormal random variables are not Normal, because their quartiles do not satisfy the equal differences property. Cell E9 has a positively skewed probability distribution, with a long tail on the high side. Cell F9 has a negatively skewed probability distribution, with a long tail on the low side.

Myerson_C05.qxd 09-09-2004 10:34 Page 188

Cells C10:D110 in Figure 5.11 contain data from 101 recalculations of these CORANDs in cells C9:D9 and Generalized Lognormal random variables in E9:F9. That is, each row in the data range C10:F110 contains in the (C and D columns) a pair of 0.9 correlated CORAND values and (in the E and F columns) the corresponding values of the Generalized Lognormal random variables that would be generated by these CORANDs. But the formula

```
=CORREL(E10:E110,F10:F110)
```

in cell F7 returns the value 0.6963. This value is just a statistical estimate from 101 sampled pairs, but it is actually an accurate estimate of the true correlation of the two random variables in cells E9 and F9.

The scatter charts at the bottom of Figure 5.11 show why this surprisingly low correlation is generated. When two random variables have a high correlation close to 1, it means that the values of these random variables tend to be scattered close to a positively sloped line when they are plotted together in an XY scatter chart. But the different skews of the random variables in the E and F columns introduce a nonlinearity that is quite evident in the middle scatterplot. Because of the high positive correlation parameter in the underlying CORANDs, the E and F values (which both have median 0) tend to be either both greater than 0 or both less than 0. But the E value tends to go much farther from 0 than the F value when they are positive, and the F value tends to go much farther from 0 than the E value when they are negative. As a result, the (E, F) pairs in the middle chart are scattered along an obtuse angle. Because these points do not stay close to any one straight line, their correlation is much less than 1.

But now we have a tricky problem: Suppose that we have data for two continuous random variables that are clearly not independent and not Normal, and we want to simulate these two random variables using CORAND. What parameter for CORAND should we use to get our simulated output to be correlated like our given data? Figure 5.11 shows that, to get Generalized Lognormals with quartile points (–3, 0, 2) and (–2, 0, 3) to have a statistical correlation of about 0.7, we may need the CORANDs' input parameter to be around 0.9. But, if we cannot estimate the appropriate parameter for CORAND by applying the CORREL function to our data for non-Normal random variables, then how can we estimate it? The answer is that for any two continuous random variables, the appropriate parameter for CORAND can be estimated by applying the CORREL function, not to our original data for these two random variables, but to order-preserving normalizations of these two data sets that are returned by the NORMIZE array function in Simtools.

Simtools's NORMIZE array function is illustrated in columns H and I of Figure 5.11. The range H10:H110 was selected and filled with the array formula

```
{=NORMIZE(E10:E110)}
```

This array formula fills cells H10:H110 with values that are ordered down the column in the same way as the data in E10:E110 (for example,

Myerson_C05.qxd 09-09-2004 10:34 Page 189

E12 > E13 > E11 > E10 and H12 > H13 > H11 > H10), but the values in H10:H110 are designed to look like values of a standard Normal random variable that has mean 0 and standard deviation 1. For example, the probability of a Normal random variable with mean 0 and standard deviation 1 being less than the value of cell H10 (–1.445) is about 0.07, which is the fraction of data values in E10:E110 that are less than the value of the corresponding cell E10. Similarly, the probability of a standard Normal random variable being less than the value of cell H12 (0.9034) is about 0.82, which is the fraction of data values in E10:E110 that are less than the value of cell E12. In effect, the array function NORMIZE(E10:E110) returns a normalization of the data in cells E10:E110.

**NOTE** (In a column of data where the values are all distinct, the NORMIZE number that corresponds to a particular value in the data column could also be computed by the formula

NORMINV((RANK(*value*, *datacolumn*, 1) –0.5)/COUNT(*datacolumn*), 0, 1)

in Excel without Simtools.)

The range I10:I110 in Figure 5.11 has been similarly filled with a normalization of the data in F10:F110 by the array formula

{=NORMIZE(F10:F110)}

As before, the normalized values in I10:I110 have the same ordering as the data values in F10:F110 (for example, I13 > I12 > I10 > I11 and F13 > F12 > F10 > F11). The statistical correlation of these two normalized data arrays in the H and I columns is computed in cell I7 by the formula

=CORREL(H10:H110,I10:I110)

This correlation of the normalized data is called the *normalized rank correlation* of the original data in the E and F columns, and it is our statistical estimate of the parameter for CORAND that could generate such data. Notice that the value of this normalized rank correlation in cell I7 (0.8956) is indeed quite close to the original CORAND parameter in A2 (0.9). Thus, when we use a CORAND array to drive a pair of non-Normal continuous random variables, we may refer to the correlation parameter of the CORANDs as a *normalized rank correlation* of the resulting non-Normal random variables, because it could be estimated from data in this way.

The basic concept of correlation, as defined in Section 5.2, is also known more fully as the *Pearson correlation coefficient*. The Pearson correlation coefficient and the normalized rank correlation are equal for Normal random variables, and the difference between them is small in many other cases. So the fine terminological distinction between these different concepts of correlation may often be ignored, and we may simply use the statistical correlations of our (untransformed) data to estimate the CORAND parameters for making random variables that are approximately Normal. But for very skewed random variables, it may be worthwhile also to look at the normalized rank correlation.

Myerson_C05.qxd 09-09-2004 10:34 Page 190

## 5.10 More About Linear Functions of Random Variables

Consider any collection of $n$ random variables ($\mathbf{X}_1, \ldots, \mathbf{X}_n$), and let $\mathbf{V}$ and $\mathbf{W}$ be two linear functions of these random variables:

$$\mathbf{W} = \alpha_0 + \alpha_1 * \mathbf{X}_1 + \cdots + \alpha_n * \mathbf{X}_n = \alpha_0 + \sum_{i=1}^{n} \alpha_i * \mathbf{X}_i$$
$$\mathbf{V} = \beta_0 + \beta_1 * \mathbf{X}_1 + \cdots + \beta_n * \mathbf{X}_n = \beta_0 + \sum_{i=1}^{n} \beta_i * \mathbf{X}_i$$

where $\alpha_i$ and $\beta_i$ are nonrandom numbers for each $i$ in $\{0,1, \ldots, n\}$. Equations (3) and (4) from Section 5.3 told us how to compute the mean and variance of such linear combinations of random variables:

(3) $$E(\mathbf{W}) = \alpha_0 + \sum_{i=1}^{n} \alpha_i * E(\mathbf{X}_i)$$

(4) $$Var(\mathbf{W}) = \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_i * \alpha_j * Covar(\mathbf{X}_i, \mathbf{X}_j)$$

The covariance of $\mathbf{W}$ and $\mathbf{V}$ satisfies the following equation:

(6) $$Covar(\mathbf{W},\mathbf{V}) = \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_i * \beta_j * Covar(\mathbf{X}_i, \mathbf{X}_j)$$

Equation (6) is a generalization of equation (4), because $Var(\mathbf{W}) = Covar(\mathbf{W},\mathbf{W})$.

Consider the special case where the random variables $\mathbf{X}_1, \ldots, \mathbf{X}_n$ are independent of each other. In this independent case, $Covar(\mathbf{X}_i, \mathbf{X}_j)$ and $Correl(\mathbf{X}_i, \mathbf{X}_j)$ are 0 whenever $i \neq j$. So our equations (4) and (6) in the independent case become

(4i) $$Var(\mathbf{W}) = (\alpha_1\text{^}2) * Var(\mathbf{X}_1) + \cdots + (\alpha_n\text{^}2) * Var(\mathbf{X}_n)$$

(6i) $$Covar(\mathbf{W},\mathbf{V}) = (\alpha_1 * \beta_1) * Var(\mathbf{X}_1) + \cdots + (\alpha_n * \beta_n) * Var(\mathbf{X}_n)$$

As one important application of these linear formulas, consider the case where $\mathbf{W}$ is the average of $n$ independent random variables $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n$, that is,

$$\mathbf{W} = (\mathbf{X}_1 + \cdots + \mathbf{X}_n)/n = (1/n) * \mathbf{X}_1 + \cdots + (1/n) * \mathbf{X}_n$$

Suppose that $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n$ have been drawn from the same probability distribution and have mean $\mu$ and standard deviation $\sigma$ (variance $\sigma\text{^}2$). With $\mathbf{W}$ defined as above, equations (3) and (4i) give us

$$E(\mathbf{W}) = (1/n) * E(\mathbf{X}_1) + \cdots + (1/n) * E(\mathbf{X}_n) = n * (1/n) * \mu = \mu$$
$$Var(\mathbf{W}) = ((1/n)\text{^}2) * Var(\mathbf{X}_1) + \cdots + ((1/n)\text{^}2) * Var(\mathbf{X}_n)$$
$$= n * ((1/n)\text{^}2) * (\sigma\text{^}2) = (\sigma\text{^}2)/n$$
$$Stdev(\mathbf{W}) = Var(\mathbf{W})\text{^}0.5 = \sigma/(n\text{^}0.5)$$

Thus, we have derived the formulas for the expected value and standard deviation of a sample average that we used in Chapter 2 for computing 95% confidence intervals.

Now suppose that $\mathbf{X}_1, \ldots, \mathbf{X}_n$ are independent Normal random variables. Then any collection of linear functions of these random variables is

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | MAKING MULTIVARIATE-NORMAL RANDOM VARIABLES | | | | | | | |
| 2 | | | Constants | | | | | |
| 3 | Independent | | 10 | 20 | 30 | | | |
| 4 | (0,1)-Normals | | Linear coefficients | | | | | |
| 5 | 0.847022 | | 5 | 7 | 0 | | | |
| 6 | 0.937496 | | 3 | 4 | -4 | | | |
| 7 | 1.190924 | | 6 | 0 | 9 | | | |
| 8 | 0.020443 | | -2 | 5 | 2 | | | |
| 9 | | Multivariate-Normal random variables | | | | | | |
| 10 | | | U | V | W | | | |
| 11 | | | 24.15225 | 29.78135 | 37.00922 | | | |
| 12 | | | | | | | | |
| 13 | Expected values | | 10 | 20 | 30 | | | |
| 14 | | | | | | | | |
| 15 | Covariances | | U | V | W | | | |
| 16 | | U | 74 | 37 | 38 | | | |
| 17 | | V | 37 | 90 | -6 | | | |
| 18 | | W | 38 | -6 | 101 | | | |
| 19 | | | | | | | | |
| 20 | Standard devns | | 8.602325 | 9.486833 | 10.04988 | | | |
| 21 | | | | | | | | |
| 22 | Correlations | | U | V | W | | | |
| 23 | | U | 1 | 0.453382 | 0.439549 | | | |
| 24 | | V | 0.453382 | 1 | -0.06293 | | | |
| 25 | | W | 0.439549 | -0.06293 | 1 | | | |
| 26 | | | | | | | | |
| 27 | Making similar Multivariate-Normals with CORAND: | | | | | | | |
| 28 | | CORANDs | 0.947834 | 0.655694 | 0.577054 | | | |
| 29 | | Multivariate-Normal random variables | | | | | | |
| 30 | | | U | V | W | | | |
| 31 | | | 23.97192 | 23.80175 | 31.95332 | | | |
| 32 | FORMULAS | | | | | | | |
| 33 | A5. =NORMINV(RAND(),0,1) | | | | | 1 | 0.453382 | 0.439549 |
| 34 | A5 copied to A5:A8 | | | | | 0.453382 | 1 | -0.99 |
| 35 | C11. =C3+SUMPRODUCT(C5:C8,$A$5:$A$8) | | | | | 0.439549 | -0.99 | 1 |
| 36 | C11 copied to C11:E11 | | | | | | | |
| 37 | C13. =C3 | | C13 copied to C13:E13 | | | #VALUE! | #VALUE! | #VALUE! |
| 38 | C16. =SUMPRODUCT($C$5:$C$8,C5:C8) | | | | | | | |
| 39 | C17. =SUMPRODUCT($D$5:$D$8,C5:C8) | | | | | | | |
| 40 | C18. =SUMPRODUCT($E$5:$E$8,C5:C8) | | | | | | | |
| 41 | C16:C18 copied to D16:D18 and E16:E18 | | | | | | | |
| 42 | C20. =C16^0.5 | | D20. =D17^0.5 | | E20. =E18^0.5 | | | |
| 43 | C23. =C16/(C20*$C$20) | | C24. =C17/(C20*$D$20) | | | | | |
| 44 | C25. =C18/(C20*$E$20) | | C23:C25 copied to D23:D25 and E23:E25 | | | | | |
| 45 | C28:E28. {=CORAND(C23:E25)} | | | | F37:H37. {=CORAND(F33:H35)} | | | |
| 46 | C31. =NORMINV(C28,C13,C20) | | | | C31 copied to C31:E31 | | | |

FIGURE **5.12**
**Making Multivariate Normal random variables**

a Multivariate Normal collection. This fact provides a way to make Multivariate Normal random variables in Excel without Simtools, as illustrated in Figure 5.12.

Cells A5:A8 of Figure 5.12 contain the formula

=NORMINV(RAND( ),0,1)

to give us four independent Normal random variables ($\mathbf{X}_1, \ldots, \mathbf{X}_4$), each of which has mean 0 and standard deviation 1. Entering the formula

=C3+SUMPRODUCT(C5:C8,$A$5:$A$8)

Myerson_C05.qxd 09-09-2004 10:34 Page 192

into cell C11, and copying C11 to C11:E11, we get a collection of three nonindependent Multivariate Normal random variables (**U**, **V**, and **W**) in cells C11:E11, each generated by a different linear function of the independent Normals in A5:A8.

Because each of the underlying independent Normals in cells A5:A8 in Figure 5.12 has an expected value equal to 0, equation (3) tells us that the expected values of the three Multivariate Normals in cells C11:E11 are just the constant terms from C3:E3. The covariances among the three Multivariate Normals in C11:E11 are computed in cells C16:E18 using equation (4i), simplified by the fact that each of the independent Normals in A5:A8 has a variance equal to 1 (because its standard deviation is 1 and 1^2 = 1). With all $Var(\mathbf{X}_i) = 1$, equation (4i) just tells us to sum the products of the coefficients. Thus, for example, the covariance between the random variables in cells D11 and C11 is computed in cell C17 by the formula

```
=SUMPRODUCT($D$5:$D$8,C5:C8)
```

The standard deviations of the random variables in C11:E11 are computed in cells C20:E20 by taking the square root of each random variable's covariance with itself. Then the correlations among the random variables in C11:E11 are computed in cells C23:E25 by dividing the covariance of each pair of random variables by the product of their standard deviations.

The correlation numbers that we see in C23:E25 of Figure 5.12 may have some intuitive interpretation. The random variables in cells C11 and D11 have a positive correlation (0.453) mainly because C11 and D11 both depend with positive coefficients on the independent Normals in cells A5 and A6, so any change in A5 or A6 would push C11 and D11 in the same direction. But the coefficients in cells C8 and D8 have different signs, and so any change in cell A8 would push C11 and D11 in opposite directions, which decreases the correlation of C11 and D11. The correlation between cells D11 and E11 is negative (–0.0629) mainly because the random cell A6 has opposite effects on D11 and E11 due to the coefficients 4 and –4 in cells D6 and E6.

Cells C31:E31 of Figure 5.12 contain a collection of Multivariate Normal random variables made with a CORAND array in C28:E28 to have the same means, standard deviations, and correlations as the Multivariate Normals in cells C11:E11. That is, the random variables in C28:E28 have exactly the same joint distribution as the random variables in C11:E11, and so no statistical test could ever distinguish between simulation data from these two ranges.

The correlations in C23:E25 of Figure 5.12 ultimately depend on the linear coefficients above in cells C5:E8. To make a computer program that generates Multivariate Normal random variables for any given correlation array, the hard part is to find coefficients like those in C5:E8 that will generate the desired correlation array. The internal computations of the

Myerson_C05.qxd 09-09-2004 10:34 Page 193

CORAND function actually involve solving this mathematical problem (by a method known as Cholesky factorization).

But you should know that some square arrays may not be "possible" as correlation arrays. We have already seen that a correlation array must have all 1s on the diagonal from top left to bottom right, and there must be a symmetry between the values to the right of this diagonal and below this diagonal. But there can also be other, more subtle restrictions. For example, cells F33:H25 in Figure 5.12 show a square array that looks the same as the correlation array in C23:E25, except that the correlation among the second and third random variables has been changed to –0.99. If we tried to make Multivariate Normals with this array of correlations, the CORAND (F33:H25) formula would return an error message, as shown in cells F37:H37 of Figure 5.12. The problem is that you would be asking CORAND to make three random variables in which the first has a substantial positive correlation with both the second and the third, whereas the second and third have very strong negative correlation. This is impossible, because the first random variable cannot be kept close enough to both of the others when they are being kept so far away from each other. Thus, there can be subtle mathematical constraints on the correlations of three or more random variables. But fortunately, these feasibility constraints will always be satisfied by the correlation arrays that we estimate from statistical data.

## 5.11 Summary

We began this chapter by considering joint distributions of discrete random variables. The covariance of two random variables is defined as the expected value of the product of their deviations from their respective means, and their correlation is defined as their covariance divided by the product of their standard deviations. Between these two measures of the relationship between random variables, we emphasized the correlation here because it is a unit-free number between –1 and 1 and it is somewhat easier to interpret. We learned to use correlation arrays for two or more random variables, which display the correlations among all pairs of these random variables, and always have 1s along the diagonal from top left to bottom right.

We studied some general formulas for linear functions of random variables. If we know the means, standard deviations, and correlations of a collection of random variables, then the expected values, variances, standard deviations, and covariances of any linear functions of these random variables can be computed by applications of some general equations ((1)–(6), (4i), (6i), listed in Sections 5.2, 5.3, and 5.10). These equations

Myerson_C05.qxd 09-09-2004 10:34 Page 194

were applied to financial portfolio analysis. For a portfolio of investments that contains given numbers of shares of various financial assets, if means, standard deviations, and correlations of the future values per share of these assets are specified in appropriately named ranges of our spreadsheet, then the expected future value of the portfolio can be computed by the formula

SUMPRODUCT(*shares*, *means*)

and the standard deviation of the portfolio's future value can be computed by the formula

SUMPRODUCT(PRODS(*shares*), PRODS(*stdevns*), *correlations*)^0.5

We introduced the family of Multivariate Normal random variables, which have a joint probability distribution that is parameterized by their means, standard deviations, and correlations. It is very convenient to work with Multivariate Normal random variables because any linear function of them has a Normal probability distribution. Multivariate Normal random variables can be constructed by using Simtools's CORAND array function to provide appropriately correlated random inputs to NORMINV.

Correlations are estimated from sample data by the CORREL function in Excel. The Simtools function MCORRELS can also be used to make correlation arrays from sample data for several random variables. We also introduced a technique for subjective assessment of the correlation between two unknown quantities based on questions about how the conditional median for one unknown quantity would depend on given information about the other unknown quantity.

CORAND can be used to generate correlated random variables with other probability distributions, but the statistical correlation of non-Normal random variables may be different from the correlation parameter of the CORANDs that drive them. In general, the CORANDs' correlation parameter actually corresponds to a normalized rank correlation of the resulting random variables. The difference between this normalized rank correlation and the usual statistical correlation is often quite small, however, and it vanishes for Normal random variables.

Two important Excel techniques were introduced in this chapter: array formulas and Solver. Array formulas are entered into a selected range of cells with the special keystroke [Ctrl][Shift][Enter]. To use Solver, we must specify the target cell to be maximized or minimized, the changing cells, and the constraints that make up an optimization problem. We learned to recognize solution difficulties such as infeasibility, unbounded optimality, and local optimality. Because of the possibility of a local optimal solution that is not the true optimum, we saw the importance of rerunning Solver with different initial conditions.

Myerson_C05.qxd 09-09-2004 10:34 Page 195

The Excel function CORREL was introduced in this chapter. We also used several Simtools functions to make correlation analysis easier: COVARPR, CORRELPR, MIDRAND, and the array functions PRODS, CORAND, MCORRELS, and NORMIZE.

## Exercises

**1.** Suppose we are planning to make investments in three mutual funds. Suppose the annual growth ratios of the three funds are Multivariate Normal random variables with the following means, standard deviations, and correlations:

The annual growth ratio of Fund 1 has expected value 1.12 and standard deviation 0.15.

The annual growth ratio of Fund 2 has expected value 1.14 and standard deviation 0.18.

The annual growth ratio of Fund 3 has expected value 1.10 and standard deviation 0.19.

The annual growth ratios of Funds 1 and 2 have correlation 0.8.

The annual growth ratios of Funds 1 and 3 have correlation 0.5.

The annual growth ratios of Funds 2 and 3 have correlation 0.3.

We have $50,000 in cash that we can invest now. Any cash that we do not invest in these mutual funds will be put into risk-free short-term bonds that pay 5% annual interest (not compounded).

**(a)** Set up a spreadsheet to compute, for any level of investments in these three funds:

**(a1)** the expectation of our portfolio's value one year from now,

**(a2)** the standard deviation of our portfolio's value one year from now,

**(a3)** the probability that our portfolio's value one year from now will be less than $50,000,

**(a4)** the certainty equivalent of the portfolio one year from now, if RiskTolerance = $10,000.

Compute these quantities for the plan of investing $20,000 in Fund 1, $15,000 in Fund 2, $10,000 in Fund 3, and $5000 in the risk-free bonds.

**(b)** Find the portfolio that maximizes the certainty equivalent of our portfolio's value next year if we evaluate such gambles with a constant risk tolerance of $10,000.

**(c)** How does your answer to part (b) change if we add the constraint that the whole $50,000 must be invested now in the mutual funds? (That is, nothing can be in bonds.)

**2.** An investor is considering two mutual funds. In each year, the logarithmic rates of return for these two funds are Multivariate Normal random variables. The logarithmic rate of return for Fund 1 each year has expected value 0.09 and standard deviation 0.18. The logarithmic rate of return for Fund 2 each year has expected value 0.07 and standard deviation 0.20. The correlation coefficient between the two funds' logarithmic rates of return each year is 0.4.

Suppose that the rates of return in any one year are independent of all past and future rates of return of these funds.

**(a)** Make a simulation model of the values that will be returned five years from now per dollar invested today in each of these two funds.

**(b)** Suppose an investor with $20,000 to invest puts $15,000 in Fund 1 and $5000 in Fund 2 today and plans to hold these investments for five years. Using at least 400 simulations of your model, show the cumulative risk profile for the total value of the portfolio in five years.

**(c)** Now consider alternative investment plans of putting $\$X$ into Fund 1 and $\$20{,}000 - X$ into Fund 2, for $X = 0$, $X = 1000$, $X = 2000$, . . . , up to $X = 20{,}000$. Make a table and a chart showing, for each of these plans, the *expected value* and the *standard deviation* of the total value of the portfolio after five years. Also, for each of these investment plans, estimate the *certainty equivalent* of the portfolio five years from now if the investor evaluates such gambles with a constant risk tolerance of $15,000.

**(d)** Two money managers have $1 million to invest over five years. Manager #1 plans to invest everything in Fund 1. Manager #2 plans to

invest everything in Fund 2 at first but will switch everything to Fund 1 at the end of any year when the portfolio value has grown to be at least 5% more than manager #1's portfolio value. What is the probability that #2's portfolio will have a higher value than #1's at the end of five years? Compare their cumulative risk profiles.

**3.** Let **X** denote the dollar value per share that will be returned by Stock 1 next year, and let **Y** denote the dollar value per share that will be returned by Stock 2 next year. In a simple discrete model, we figure that the possible values of **X** are \$40, \$45, and \$50, and the possible values of **Y** are \$8, \$9, and \$10, and the joint probability distribution for **X** and **Y** is

| $x$ | $y$ | $P(\mathbf{X} = x \cap \mathbf{Y} = y)$ |
|---|---|---|
| 40 | 8 | 0.15 |
| 40 | 9 | 0.05 |
| 40 | 10 | 0 |
| 45 | 8 | 0.10 |
| 45 | 9 | 0.25 |
| 45 | 10 | 0.10 |
| 50 | 8 | 0 |
| 50 | 9 | 0.15 |
| 50 | 10 | 0.20 |

**(a)** Calculate the expected values $E(\mathbf{X})$ and $E(\mathbf{Y})$.

**(b)** Calculate the standard deviations *Stdev*(**X**) and *Stdev*(**Y**).

**(c)** Calculate the covariance *Covar*(**X**,**Y**) and the correlation *Correl*(**X**,**Y**).

**(d)** Suppose that Stock 1 is currently selling for \$42.50 per share and Stock 2 is currently selling for \$8.50 per share. An investor currently owns 600 shares of Stock 1 and has no shares of Stock 2 but is thinking about selling some of the shares of Stock 1 and investing the proceeds in Stock 2 now. Make a table showing how the expected value and standard deviation of the portfolio next year would depend on how many shares of Stock 1 are kept. In this table, just consider the possibility of selling whole blocks of 100 shares, so that the investor might keep 0 shares, 100 shares, 200 shares, etc., up to 600 shares of Stock 1.

**4.** (*Global Sports*, version 2.) Century Publications is considering the introduction of a new magazine. It is very likely to lose money in the first year, but in the long run it might be profitable. Managers at Century have focused on two unknown quantities: the net profit from the new magazine in its first year of publication, and the potential net present value of the profits from the magazine in all years after the first if it is continued in publication. For net profit in the first year (with negative numbers denoting a loss), they assessed four equally likely quartiles with boundary points \$3.7, –3.3, and –2.6 million. For the potential net present value of profits from the magazine in all subsequent years, they assessed quartile boundary points –\$9.0, –2.0, and 8.0 million. But if they learned that first-year profit would be –\$2.6 million, then they would believe that the potential net present value of profits in all subsequent years would be equally likely to be above or below \$7.0 million.

**(a)** Make a simulation model of these two unknown quantities, first-year profit and the potential net present value of profits in all subsequent years, that is consistent with the conditional median and quartile assessments of Century's managers.

**(b)** Suppose that Century plans to introduce this new magazine and continue publishing it forever, regardless of its profits or losses. Under this plan, the new magazine's overall net present value is just the sum of its first-year net profit and its net present value of profits in all subsequent years. Based on data from at least 1000 simulations of your model, estimate the expected value of this new magazine's overall net present value under this plan, and show its cumulative risk profile.

**(c)** Now suppose that Century plans to introduce this new magazine for one year and will continue publishing it forever if the first-year net profit is greater than –\$3.4 million, but will stop publishing it if the first-year net profit is worse than –\$3.4 million. So under this plan, if the first-year net profit is greater than –\$3.4 million, then the new magazine's overall net present value is the sum of its first-year net profit and its net present value of profits in all subsequent years, but if the first-year net profit is less than –\$3.4 million, then the new magazine's overall net present value is just its first-year net profit. Based on data from at least 1000 simulations of your model, estimate the expected value of this new magazine's overall net present value under this alternative plan, and show its cumulative risk profile.

**5.** Ed Coke is looking for partners to share the risks with him in a real estate development project that will yield its uncertain profits a year from now. The value of the total profits that will be earned from this project next year is a Normal random variable with mean \$100,000 and standard deviation \$25,000. So Coke has decided to issue 1000 shares

Myerson_C05.qxd 09-09-2004 10:34 Page 197

in the project, each of which will get 1/1000 of the total profit from the project. So the returns per share next year will be a Normal random variable with mean $100 and standard deviation $25.

Coke is hoping to sell shares in this project to many of his friends and relations, including his old colleague Frank Bacon. Coke knows that Bacon has about $45,000 in liquid assets, of which about half is in the bank and half is in a well-diversified mutual fund. The value returned next year per dollar invested now in this mutual fund is considered to be a Normal random variable with mean $1.10 and standard deviation $0.20. Bank accounts pay 4% interest, so each dollar invested now in the bank will yield $1.04 next year. Coke figures that Bacon has a constant risk tolerance of $15,000.

The returns from Coke's project and from the well-diversified mutual fund have a correlation equal to 0.7. Base your answers to the following questions on 2000 or more simulations of these random variables.

**(a)** Suppose first that Coke sells shares in his project now for $88 per share. How many shares of Coke's project should Bacon want to buy, as a part of an optimal portfolio with accounts in the bank and the mutual fund, to maximize Bacon's certainty equivalent?

**(b)** Suppose that Coke may sell the shares in his project now for some whole-integer number of dollars per share. What is the highest price per share such that Bacon would still want to buy some positive number of shares in Coke's project?

**(c)** How would your answers to parts (a) and (b) change if the returns from Coke's project and from the well-diversified mutual fund were independent random variables with correlation equal to 0?

Myerson_C05.qxd 09-09-2004 10:34 Page 198

Myerson_C06.qxd 07-09-2004 09:43 Page 199

# 6
# Conditional Expectation

The preceding chapter introduced correlation as a mathematical measure of relationships among random variables. In this chapter, we develop other ways to describe and analyze relationships among random variables in probability models. The emphasis in this chapter is on conditional expectation, which combines the ideas of conditional probability (introduced in Chapter 1) and expected value (introduced in Chapter 2). When we measure how the conditional expectation of one random variable may depend on the value of another random variable, we can find different kinds of relationships between these random variables that correlation alone could not distinguish.

Conditional expectations in any probability model must satisfy an important equation which says that the expected value of the conditional expectation of a random variable must equal the expected value of the random variable with current information. This expected-posterior law (also known as the law of iterated expectation) is explained and illustrated in the first part of this chapter. Later in this chapter, we introduce the basic ideas of regression analysis, the most widely used method of estimating conditional expectations from data.

## 6.1 Dependence Among Random Variables

We have emphasized the importance of building appropriate relationships among random variables in a simulation model. The concept of correlation was developed by statisticians as a way to measure relationships among random variables. Any simulation add-in like Simtools should include some device like CORAND to help us match the correlations of our random variables to the estimated correlations of the real-world unknown quantities that we are simulating. But correlation does not give a complete measure

Myerson_C06.qxd 07-09-2004 09:43 Page 200

of the relationship between any pair of random variables in general, and CORAND is not the only way that we should build relationships among random variables in our models. A more versatile way to build relationships among random variables in a spreadsheet model, which we emphasize in this chapter, is by making the formula for one random variable depend on the values of other random variables in the spreadsheet.

The word *depend* here deserves some careful discussion because it has two possible meanings. Two random variables are *statistically dependent* (or not independent) when learning more about one random variable could cause you to revise your beliefs about the other random variable. In this statistical sense, the dependence of two random variables does not imply that they have any cause-and-effect relationship in which one influences the other, because statistical dependence could also arise when two random variables are both influenced by some other random factor. But in spreadsheet models, we build cause-and-effect relationships among cells in their formulas, which gives us another concept of *formulaic dependence*. A cell is formulaically dependent on all the cells to which it refers in its formula, and on all the cells to which these cells refer, and so on.

The CORAND function that we studied in the preceding chapter gives us an easy way to build statistical dependence among random variables without either random variable formulaically depending on the other. But we should also use formulaic dependence, as a more general and versatile method for building appropriate statistical dependence among random variables.

We have learned to make random variables in spreadsheets by formulas in which an inverse cumulative function (like NORMINV) has a RAND( ) or CORAND as its first parameter. The inverse cumulative function then has other parameters (commonly the mean and standard deviation) that define the specific shape of the probability distribution. The simplest way to build formulaic dependence among random variables is to set these distributional parameters in one random variable by formulas that refer to the value of another random variable in the spreadsheet.

Let us begin by revisiting the example from Chapter 1 of the new salesperson whose skill level is unknown. Let **L** denote the unknown skill level, and let **X** denote the unknown number of sales that the salesperson will get in 20 calls. As in Chapter 1, let us assume that the successes or failures in the 20 calls would be conditionally independent given the person's skill level, and this skill level itself can be numerically identified as the conditional probability of success in each call.

In the simple discrete model of Chapter 1, the skill levels 1/3 and 2/3 each had probability 0.5, and so the salesperson's unknown skill level initially had mean 1/2 and standard deviation 1/6. Now suppose that these measures of the center and spread of the skill distribution are considered reasonable in this situation, but that the two-valued discreteness of the skill distribution is considered unrealistic. Because the skill level in this situation must be a number between 0 and 1, we want to use a continuous probability

Myerson_C06.qxd 07-09-2004 09:43 Page 201

distribution that admits any possible value between 0 and 1, but has 0 probability outside this interval, and that comes from a family of distributions that can be parameterized by the mean and standard deviation. A special family of probability distributions called *Beta distributions* has exactly these properties (as discussed in Section 4.8). So let us suppose now that, before observing the results of these 20 sales calls, our prior beliefs about the salesperson's unknown skill level can be described by a Beta distribution with mean 1/2 and standard deviation 1/6. (Using the Beta distribution here is a more sophisticated and versatile assumption than the simplistic Triangular distribution that was used in Section 1.5.)

Excel provides a function BETAINV for making Beta random variables, but its parameters are rather unintuitive. So we may use instead the Simtools add-in function BETINV for making Beta random variables, which takes the mean and the standard deviation as its second and third parameters. (To avoid confusion when you use one of these similarly named functions, you should use the Insert:Function dialog box to verify that BETINV is the function that has "mean" and "standard deviation" listed among its parameters.) BETINV's first "probability" parameter is where we put the RAND( ) to make a Beta random variable. Such a Beta random variable with mean (or expected value) 1/2 and standard deviation 1/6 has been entered in cell A2 of Figure 6.1 with the formula

=BETINV(RAND( ),F2,F3),

where cell F2 has the value 1/2 and cell F3 has the value 1/6.

The number of successes in $n$ trials, where each trial has independent probability $p$ of being a success, is a *Binomial random variable* with parameters $n$ and $p$ (as was discussed in Section 4.8). So once a skill level **L** has been given, the conditional probability distribution for the number of successful sales should be a Binomial distribution with parameters $n = 20$ and $p = \mathbf{L}$. As in Chapter 1, such a Binomial random variable could be simulated by the sum of 20 random cells containing IF statements, each of which takes the value 1 when an independent RAND( ) is less than the skill level **L**, and takes the value 0 otherwise. But such a Binomial random variable can also be simulated in one cell by the Simtools formula BINOMINV(RAND( ), 20, **L**). So with the skill level **L** simulated in cell A2, the number of successful sales can be simulated in cell A5 of Figure 6.1 by the formula

=BINOMINV(RAND( ),20,A2)

Thus, the cause-and-effect relationship between the salesperson's skill level and the number of successful sales in the real world is represented in this simple model by the formulaic dependence of the sales in cell A5 on the skill level in cell A2.

Does the formula =BINOMINV(RAND( ),20,A2) make cell A5 a Binomial random variable? The BINOMINV function is designed to return a Binomial random variable whenever its first parameter is a RAND( ) or CORAND and its other parameters are given constants. But because the third parameter of

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | Salesperson's Skill level (potential long-run sales rate) | | | | | | | | |
| 2 | 0.65486 | from Beta dist with | | | mean | 0.5 | | | |
| 3 | | | | | stdev | 0.16667 | | | Total |
| 4 | Total Sales in 20 calls | | | | | | Skill | Sales | 501 |
| 5 | 13 | | | | | Sales | E | Stdev | Freq |
| 6 | | Skill | | | | (,D10) | 0.46642 | 0.09719 | 56 |
| 7 | | 0.49423 | E | | 0 | | 0.08138 | 0.02915 | 2 |
| 8 | | 0.16349 | Stdev | | | 1 | 0.22364 | 0.03374 | 3 |
| 9 | | | | Sales= | | 2 | 0.17788 | 0.04023 | 7 |
| 10 | | Skill | Sales | 9 | | 3 | 0.2485 | 0.08809 | 15 |
| 11 | SimTabl | 0.65486 | 13 | Skill | Sales | 4 | 0.29003 | 0.07266 | 12 |
| 12 | 0 | 0.5741 | 9 | 0.5741 | | 5 | 0.31824 | 0.0616 | 25 |
| 13 | 0.002 | 0.58524 | 13 | .. | | 6 | 0.37102 | 0.07852 | 37 |
| 14 | 0.004 | 0.56375 | 10 | .. | | 7 | 0.37522 | 0.09889 | 29 |
| 15 | 0.006 | 0.4461 | 6 | .. | | 8 | 0.41398 | 0.07532 | 41 |
| 16 | 0.008 | 0.37011 | 5 | .. | | 9 | 0.46642 | 0.09719 | 56 |
| 17 | 0.01 | 0.46031 | 9 | 0.46031 | | 10 | 0.50592 | 0.09741 | 50 |
| 18 | 0.012 | 0.60782 | 11 | .. | | 11 | 0.52946 | 0.10524 | 51 |
| 19 | 0.014 | 0.81344 | 18 | .. | | 12 | 0.56456 | 0.10635 | 50 |
| 20 | 0.016 | 0.44203 | 12 | .. | | 13 | 0.57518 | 0.09002 | 38 |
| 21 | 0.018 | 0.64738 | 15 | .. | | 14 | 0.64211 | 0.07884 | 24 |
| 22 | 0.02 | 0.7647 | 17 | .. | | 15 | 0.68151 | 0.06761 | 21 |
| 23 | 0.022 | 0.50473 | 9 | 0.50473 | | 16 | 0.72132 | 0.08474 | 16 |
| 24 | 0.024 | 0.671 | 9 | 0.671 | | 17 | 0.79614 | 0.08313 | 11 |
| 25 | 0.026 | 0.43953 | 8 | .. | | 18 | 0.78445 | 0.04917 | 8 |
| 26 | 0.028 | 0.46632 | 9 | 0.46632 | | 19 | 0.82611 | 0.06462 | 3 |
| 27 | 0.03 | 0.17111 | 2 | .. | | 20 | 0.7538 | 0.16399 | 2 |
| 28 | 0.032 | 0.64037 | 14 | .. | | | | | |
| 29 | 0.034 | 0.36886 | 6 | .. | | E(E(Skill\|Sales)) | | E(Stdv(Sk\|Sa)) | |
| 30 | 0.036 | 0.23903 | 2 | .. | | | 0.49423 | 0.08773 | |
| 31 | 0.038 | 0.3202 | 9 | 0.3202 | FORMULAS | | | | |
| 32 | 0.04 | 0.28852 | 6 | .. | A2. =BETINV(RAND(),F2,F3) | | | | |
| 33 | 0.042 | 0.68211 | 11 | .. | A5. =BINOMINV(RAND(),20,A2) | | | | |
| 34 | 0.044 | 0.53242 | 10 | .. | B11. =A2 | | C11. =A5 | | |
| 35 | 0.046 | 0.61267 | 9 | 0.61267 | D12. =IF(C12=$D$10,B12,"..") | | | | |
| 36 | 0.048 | 0.57922 | 12 | .. | D12 copied to D12:D512 | | | | |
| 37 | 0.05 | 0.48706 | 13 | .. | B7. =AVERAGE(B12:B512) | | | | |
| 38 | 0.052 | 0.52173 | 12 | .. | B8. =STDEV(B12:B512) | | | | |
| 39 | 0.054 | 0.77622 | 15 | .. | I6. =COUNT(D12:D512) | | | | |
| 40 | 0.056 | 0.34727 | 10 | .. | G6. =IF(I6>0,AVERAGE(D12:D512),"NA") | | | | |
| 41 | 0.058 | 0.82216 | 17 | .. | H6. =IF(I6>1,STDEV(D12:D512),"NA") | | | | |
| 42 | 0.06 | 0.61336 | 14 | .. | G7:I27. {=TABLE(,D10)} | | | | |
| 43 | 0.062 | 0.60676 | 12 | .. | I4. =SUM(I7:I27) | | | | |
| 44 | 0.064 | 0.51366 | 13 | .. | G30. =SUMPRODUCT(G7:G27,$I$7:$I$27)/$I$4 | | | | |
| 45 | 0.066 | 0.52743 | 9 | 0.52743 | H30. =SUMPRODUCT(H7:H27,$I$7:$I$27)/$I$4 | | | | |

FIGURE **6.1**
**Conditional expected skills, given sales results, from a Beta prior**

BINOMINV is also a random variable here, the resulting value of cell A5 is a random variable with a more complicated probability distribution that may not fall into any of our well-known families. That is, the marginal probability distribution of cell A5 in Figure 6.1 is not Binomial. But this formula = BINOMINV(RAND( ),20,A2) tells us that the conditional probability distribution of cell A5 given any value of cell A2 is a Binomial probability distribution with parameters $n = 20$ and $p =$ A2. In effect, the random variable in cell A5 is being defined here by its conditional probability distribution given the skill level in cell A2, rather than by its marginal probability distribution.

Myerson_C06.qxd 07-09-2004 09:43 Page 203

## 6.2 Estimating Conditional Expectations and Standard Deviations

The spreadsheet shown in Figure 6.1 also includes a simulation table in A11:C512 that contains data from 501 simulations of the skill level and sales count from cells A2 and A5. We analyzed such simulation data in Chapter 1 to estimate a variety of conditional probabilities, but the concepts of expected value and standard deviation were not introduced until after Chapter 1. Let us now extend the analysis and show how the expected value and standard deviation of the skill level could be computed before or after the number of sales is observed.

The Beta distribution with mean 1/2 and standard deviation 1/6 is supposed to describe our prior beliefs about the salesperson before we observe how many sales are made in 20 calls. The simulated skill data in Figure 6.1 is contained in the range B12:B512, and the corresponding simulated sales counts are in the range C12:C512. So the expected value and standard deviation of this skill distribution can be estimated from our 501 sample points by the formulas

= AVERAGE(B12:B512) and = STDEV(B12:B512)

as shown, respectively, in cells B7 and B8 of Figure 6.1. Notice that this simulation data yields an estimated mean of 0.494 and an estimated standard deviation of 0.163, which are very close to the mean 1/2 and standard deviation 1/6 that we specified as parameters for BETINV in cell A2. So these estimates in B7 and B8 just tell us that BETINV has performed as it was supposed to.

But our beliefs about the salesperson's skill level would change after we observed, say, 9 successful sales in the first 20 calls. The conditional probabilities that would apply after such an observation may be called *posterior* beliefs, in contrast to the *prior* beliefs that applied before the observation. As in Chapter 1, we estimate conditional probabilities from simulation data by creating a column that filters out only the skill data from simulations where the sales match some criterion. In Figure 6.1, cell D10 contains the given number of sales (9) that the selected simulations should match. (In labels like D9 and D11, read "|" as "given.") The skill levels where this given number of sales occurred are selected in cells D12:D512 by entering into cell D12 the formula

=IF(C12= $D$10,B12,"..")

and then copying D12 to D12:D512. The formula

=COUNT(D12:D512)

in cell I6 tells us that the event of exactly 9 sales occurred 56 times in our 501 simulations, because the COUNT function only counts cells that contain numerical data.

Myerson_C06.qxd 07-09-2004 09:43 Page 204

The *conditional expected value* of skill given 9 sales is the expected value of the skill according to the posterior beliefs that we should have after we observe the salesperson making 9 sales. The skills in these selected simulations where 9 sales occurred are a random sample from the conditional probability distribution of skill given 9 sales, and so their average can be used as an estimate of this conditional expected value. The AVERAGE function in Excel ignores all cells that do not contain numerical data, and so the formula AVERAGE(D12:D512) in cell G6 of Figure 6.1 yields the average skill in these 56 simulations where 9 sales occurred.

The formula in cell G6 is actually a bit more complicated. If we chose another criterion in cell D10 that yielded no numerical data in cells D12:D512, then Excel would treat their AVERAGE as an error value, which could generate meaningless error values elsewhere (in cell G30). These error values can be avoided by enclosing the AVERAGE in an IF function that yields a text message instead of an error when there is insufficient numerical data in D12:D512. So cell G6 in Figure 6.1 computes the average skill among the simulations that match the D10 sales criterion by the formula

```
=IF(I6>0,AVERAGE(D12:D512),"NA")
```

Similarly, the formula

```
=IF(I6>1,STDEV(D12:D512),"NA")
```

in cell H6 returns the sample standard deviation of skills among the simulations that match the D10 sales criterion, and it returns a text message if there is insufficient data to compute this standard deviation.

Cells G6:I6 in Figure 6.1 tell us that among the 56 simulations where 9 sales occurred, the skill levels had mean 0.466 and standard deviation 0.097. So we may estimate that if we observed 9 sales, then our posterior probability distribution for the salesperson's skill level should have an expected value of about 0.466 and a standard deviation of about 0.097. The selected sample size in this case is rather small, but our 95% confidence interval for the true conditional expected value of skill given 9 sales is

$$0.466 \pm 1.96*0.097/(56\text{^}0.5) = 0.466 \pm 0.025 = [0.441, 0.491]$$

Thus, comparing cell G6 to the prior expected value (in F2 or B7), we find that the evidence of 9 sales should cause our expected value of the salesperson's skill to decrease by a small amount. The posterior standard deviation in cell H6 (0.097) is smaller than the prior standard deviation in cell B8 or F3 (0.16) because our observing 9 sales in the salesperson's 20 first calls would reduce our uncertainty about the salesperson's skill.

By substituting different numbers into cell D10, we can similarly estimate the conditional expected value and standard deviation of skill given any number of sales from 0 to 20. These conditional expected values and standard deviations are shown in cells G7:G27 and H7:H27 in Figure 6.1. Cells I7:I27 show the frequency of observing each number of sales from 0 to 20 in this simulation data. (All these results have been generated by a

Myerson_C06.qxd 07-09-2004 09:43 Page 205

column-input data table that fills the range G7:I27 with recalculated values of G6:I6 when the values in F7:F27 are substituted into cell D10.)

Let us compare these estimates of the posterior conditional expected skills in cell G7:G27 of Figure 6.1 to the prior expected skill, which is estimated in cell B7. The conditional expected skills here tend to increase as the number of sales increases, which corresponds to our basic intuition that getting more sales would be evidence of higher skill. When the number of sales is less than 10, we get posterior conditional expected skills that are less than the prior expected skill (as low as 0.081 according to these estimates). When the number of sales is greater than 10, we get posterior conditional expected skills that are greater than the prior expected skill (as high as 0.826 according to these estimates). So before we observe the number of sales, the conditional expected skill level that we will apply after this observation is an unknown quantity (somewhere between 0.081 and 0.826). The prior expected value of this posterior conditional expectation is estimated from our simulation data in cell G30 of Figure 6.1 by the formula

```
=SUMPRODUCT(G7:G27,$I$7:$I$27)/$I$4
```

That is, cell G30 multiplies each possible conditional expected skill (in G7:G27) by the frequency of the sales observation that would generate it (in I7:I27), sums these products, and then divides this sum by the total number of observations in our simulation data (501, counted in cell I4).

Notice that this prior expected value of the posterior conditional expected value of skill, as estimated in cell G30, is exactly equal to the prior expected value of skill, as estimated in cell B7 from the same simulation data. This equality is our first example of the important *expected-posterior law*. If you think about these calculations, you may see why G30 and B7 must be exactly equal. Each conditional expected value in cell G7:G27 is an average of the skills in a different subset of the data range B12:B512. When we multiply each of these averages by the size of the subset that generated it, the resulting product is the sum of all skills in that subset. So the sum of these products is the sum of the skills in the whole data range, and dividing this sum by the size of the whole data range yields the average skill in the whole data range.

Notice that the conditional standard deviations of skill in cells H7:H27 of Figure 6.1 are generally much smaller than the prior standard deviation of skill in cell B8. The expected value of the posterior conditional standard deviation, as computed in cell H30, is substantially less than the prior standard deviation of skill. That is, we expect that observing the number of successful sales in these 20 calls would substantially reduce our uncertainty about the salesperson's skill.

Figure 6.2 offers a fuller picture of how an observation of 9 successful sales in 20 calls should change our beliefs about the new salesperson's skill. The solid curve shows the inverse cumulative function for the salesperson's skill according to our prior probability distribution (before the number of successful sales is observed). The inverse cumulative values of this curve are

Myerson_C06.qxd 07-09-2004 09:43 Page 206

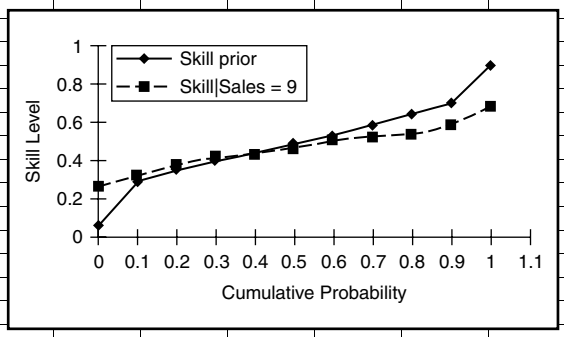


| | J | K | L | M | N | O | P | Q | R |
|---|---|---|---|---|---|---|---|---|---|
| 16 | | CumvPr | Skill | Skill\|Sales=9 | | | | | |
| 17 | | 0 | 0.06076 | 0.26348 | | | | | |
| 18 | | 0.1 | 0.28848 | 0.32756 | FORMULAS FROM RANGE L16:M27 | | | | |
| 19 | | 0.2 | 0.35304 | 0.37841 | L17. =PERCENTILE($B$12:$B$512,K17) | | | | |
| 20 | | 0.3 | 0.39833 | 0.41686 | M17. =PERCENTILE($D$12:$D$512,K17) | | | | |
| 21 | | 0.4 | 0.43777 | 0.43744 | L17:M17 copied to L17:M27 | | | | |
| 22 | | 0.5 | 0.49241 | 0.46332 | Chart series are (K17:K27,L17:L27) | | | | |
| 23 | | 0.6 | 0.53494 | 0.50473 | and (K17:K27,M17:M27). | | | | |
| 24 | | 0.7 | 0.58696 | 0.5241 | M16. ="Skill\|Sales="&D10 | | | | |
| 25 | | 0.8 | 0.64345 | 0.53695 | | | | | |
| 26 | | 0.9 | 0.70456 | 0.58713 | | | | | |
| 27 | | 1 | 0.90282 | 0.68397 | | | | | |

FIGURE **6.2**
**Prior and posterior inverse cumulative distributions for salesperson's skill**

PERCENTILE values (computed in cells L17:L27 of Figure 6.2) from our skills data (stored in cells B12:B512 of Figure 6.1) that correspond to various cumulative probabilities from 0 to 1 (listed in cells K17:K27 of Figure 6.2). The dashed curve in Figure 6.2 shows the inverse cumulative function for the salesperson's skill according to the posterior distribution after observing 9 sales, estimated as PERCENTILE values from the selected skills data for simulations with 9 sales (listed in cells D12:D512 of Figure 6.1). Comparing these two curves, you can see that the observation of 9 sales causes a substantial decrease in the higher levels of skill that we should consider likely, but it also causes a (somewhat smaller) increase in the lower levels of skill that we should consider likely.

## 6.3 The Expected-Posterior Law in a Discrete Example

A careful fundamental explanation of conditional expectation and the expected-posterior law should be developed in the general context of two discrete random variables. To illustrate these ideas as we develop them, let us consider again the discrete example of two stock prices, **X** and **Y,** that we discussed at the beginning of Chapter 5.

In this simple example, the possible values of **X** are 70, 75, and 80, and the possible values of **Y** are 30, 32, and 34. The joint probability distribution

| | A | B | C | D | E | F | G | H | I | J | K | L | M |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | JointProbys P(X=x∩Y=y) | | | | | | | CondlProbys P(Y=y\|X=x) | | | | | |
| 2 | | y= | | | | | | | y= | | | | |
| 3 | x= \ | 30 | 32 | 34 | | P(X=x) | | x= \ | 30 | 32 | 34 | | E(Y\|X=x) |
| 4 | 70 | 0.1 | 0.1 | 0 | | 0.2 | | 70 | 0.5 | 0.5 | 0 | | 31 |
| 5 | 75 | 0.2 | 0.3 | 0 | | 0.5 | | 75 | 0.4 | 0.6 | 0 | | 31.2 |
| 6 | 80 | 0 | 0.2 | 0.1 | | 0.3 | | 80 | 0 | 0.66667 | 0.33333 | | 32.67 |
| 7 | | P(Y=y) | | | | | | | | | | | |
| 8 | | 0.3 | 0.6 | 0.1 | | | | Let A denote the event that Y=32. | | | | | |
| 9 | | | | | | | | | | | | | |
| 10 | | E(X) | E(Y) | | | | | | | E(P(A\|X)) | | | E(E(Y\|X)) |
| 11 | | 76 | 32 | | | | | | | 0.6 | | | 31.6 |
| 12 | | | | | | | | | | | | | |
| 13 | SIMULATION MODEL | | | | | | | | | | | | |
| 14 | | X | Y | A | | | | y= | 30 | 32 | 34 | | |
| 15 | SimTa | 75 | 30 | 0 | | | P(Y=y\|X) | | 0.4 | 0.6 | 0 | | |
| 16 | 0 | 75 | 30 | 0 | | | | | | | | | |
| 17 | .002 | 75 | 32 | 1 | | | | | Estimates from simulation | | | | |
| 18 | .004 | 80 | 32 | 1 | | | | | E(Y) | P(A) | | | |
| 19 | .006 | 80 | 32 | 1 | | | | | 31.5768 | 0.62475 | | | |
| 20 | .008 | 75 | 30 | 0 | | | x | Freq | E(Y\|X=x) | P(A\|X=x) | P(X=x) | | |
| 21 | .010 | 75 | 30 | 0 | | | 70 | 107 | 30.9907 | 0.49533 | 0.21357 | | |
| 22 | .012 | 80 | 32 | 1 | | | 75 | 246 | 31.2439 | 0.62195 | 0.49102 | | |
| 23 | .014 | 80 | 32 | 1 | | | 80 | 148 | 32.5541 | 0.72297 | 0.29541 | | |
| 24 | .016 | 70 | 30 | 0 | | | | E(E(Y\|X)) | | E(P(A\|X)) | | | |
| 25 | .018 | 80 | 32 | 1 | | | | | 31.5768 | 0.62475 | | | |
| 26 | .020 | 75 | 30 | 0 | FORMULAS | | | | | | | | |
| 27 | .022 | 80 | 34 | 0 | F4. =SUM(B4:D4) | | | | | | copied to F4:F6 | | |
| 28 | .024 | 80 | 32 | 1 | I4. =B4/$F4 | | | | | | copied to I4:K6 | | |
| 29 | .026 | 70 | 32 | 1 | M4. =SUMPRODUCT(I4:K4,$B$3:$D$3) | | | | | | copied to M4:M6 | | |
| 30 | .028 | 75 | 32 | 1 | B8. =SUM(B4:B6) | | | | | | copied to B8:D8 | | |
| 31 | .030 | 70 | 30 | 0 | B11. =SUMPRODUCT(A4:A6,F4:F6) | | | | | | | | |
| 32 | .032 | 80 | 32 | 1 | C11. =SUMPRODUCT(B3:D3,B8:D8) | | | | | | | | |
| 33 | .034 | 75 | 32 | 1 | J11. =SUMPRODUCT(J4:J6,$F$4:$F$6) | | | | | | | | |
| 34 | .036 | 75 | 32 | 1 | M11. =SUMPRODUCT(M4:M6,$F$4:$F$6) | | | | | | | | |
| 35 | .038 | 75 | 30 | 0 | B15. =DISCRINV(RAND(),A4:A6,F4:F6) | | | | | | | | |
| 36 | .040 | 80 | 34 | 0 | I15:K15. {=INDEX(I4:K6,MATCH(B15,H4:H6),0)} | | | | | | | | |
| 37 | .042 | 80 | 32 | 1 | C15. =DISCRINV(RAND(),I14:K14,I15:K15) | | | | | | | | |
| 38 | .044 | 75 | 32 | 1 | D15. =IF(C15=32,1,0) | | | | | | | | |
| 39 | .046 | 70 | 30 | 0 | I19. =AVERAGE(C16:C516) | | | | | | copied to J19 | | |
| 40 | .048 | 80 | 34 | 0 | H21. =COUNTIF($B$16:$B$516,G21) | | | | | | copied to H21:H23 | | |
| 41 | .050 | 80 | 34 | 0 | I21. =SUMIF($B$16:$B$516,$G21,C$16:C$516)/$H21 | | | | | | | | |
| 42 | .052 | 70 | 32 | 1 | I21 copied to I21:J23 | | | | | | | | |
| 43 | .054 | 75 | 32 | 1 | K21. =H21/SUM($H$21:$H$23) | | | | | | copied to K21:K23 | | |
| 44 | .056 | 75 | 30 | 0 | I25. =SUMPRODUCT(I21:I23,$K$21:$K$23) | | | | | | | | |
| 45 | .058 | 70 | 30 | 0 | J25. =SUMPRODUCT(J21:J23,$K$21:$K$23) | | | | | | | | |

FIGURE **6.3**
**Conditional probability and conditional expectation in a discrete example**

of **X** and **Y** is shown in the range A3:D6 of Figure 6.3. For any number $x$ that is a possible value of **X**, the marginal probability $P(\mathbf{X} = x)$ is computed by summing the joint probabilities $P(\mathbf{X} = x \cap \mathbf{Y} = y)$ over all numbers $y$ that are possible values of **Y**, as shown in cells F4:F6 of Figure 6.3. The mathematical formula is

$$P(\mathbf{X} = x) = \sum_y P(\mathbf{X} = x \cap \mathbf{Y} = y)$$

where the sum includes all $y$ that are possible values of **Y**. Similarly, the marginal probabilities $P(\mathbf{Y} = y)$ are computed by summing down each column of the joint probability table, as shown in cells B8:D8.

Myerson_C06.qxd 07-09-2004 09:43 Page 208

For all numbers $x$ and $y$ that are possible values of **X** and **Y**, the conditional probability $P(\mathbf{Y} = y|\mathbf{X} = x)$ is computed by the basic formula

$$P(\mathbf{Y} = y|\mathbf{X} = x) = P(\mathbf{X} = x \cap \mathbf{Y} = y)/P(\mathbf{X} = x)$$

This formula is applied in the range I4:K6 in Figure 6.3. Comparing each row of conditional probabilities in I4:K6 with the corresponding row of joint probabilities in B4:D6, you can see that the conditional probabilities in the row are respectively proportional to the joint probabilities in the same row, where the constant of proportionality is what it takes to make the conditional probabilities sum to 1 across the row; for example, the top row of conditional probabilities given **X** = 70 in I4:K4,

$$P(\mathbf{Y} = 30|\mathbf{X} = 70) = 0.5, \qquad P(\mathbf{Y} = 32|\mathbf{X} = 70) = 0.5,$$
$$P(\mathbf{Y} = 34|\mathbf{X} = 70) = 0$$

which are proportional to the corresponding joint probabilities (0.1, 0.1, 0) in B4:D4.

If we learned that the value of **X** was some particular number $x$, then our *conditional expected value* of **Y** given that **X** = $x$ would be

$$E(\mathbf{Y}|\mathbf{X} = x) = \sum_y P(\mathbf{Y} = y|\mathbf{X} = x)*y$$

where the sum includes all $y$ that are possible values of **Y**. In our example, the conditional expected value of **Y** given **X** = 70 is

$$E(\mathbf{Y}|\mathbf{X} = 70) = 0.5*30 + 0.5*32 + 0*34 = 31.00$$

Similarly, the conditional expected value of **Y** given **X** = 75 in this example is

$$\mathrm{E}(\mathbf{Y}|\mathbf{X} = 75) = 0.4*30 + 0.6*32 + 0*34 = 31.20$$

The conditional expected value of **Y** given **X** = 80 is

$$E(\mathbf{Y}|\mathbf{X} = 80) = 0*30 + 0.667*32 + 0.333*34 = 32.67$$

These conditional expected values are computed by SUMPRODUCT formulas in cells M4:M6 in Figure 6.3.

If we are about to learn **X** but not **Y**, then our posterior expected value of **Y** after learning **X** will be either 31.00 or 31.20 or 32.67, depending on whether **X** turns out to be 70, 75, or 80. These three possible values of **X** have prior probability 0.2, 0.5, and 0.3, respectively. So our prior expected value of the posterior expected value of **Y** that we will apply after learning **X** is

$$E(E(\mathbf{Y}|\mathbf{X})) = 0.2*31.00 + 0.5*31.20 + 0.3*32.67 = 31.6$$

as shown in cell M9 of Figure 6.3. But before learning **X**, our expected value of **Y** is the same number, as shown in cell C11 of Figure 6.3:

$$E(\mathbf{Y}) = 0.3*30 + 0.6*32 + 0.1*34 = 31.6$$

More generally, the *expected-posterior law* says that before we observe a random variable **X**, our expected value of another random variable **Y** must be equal to our expected value of what will be the posterior conditional

expected value of **Y** after the observation of **X**. The formula for the expected-posterior law is

$$E(\mathbf{Y}) = \sum_x P(\mathbf{X} = x) * E(\mathbf{Y} | \mathbf{X} = x)$$

where the sum includes all $x$ that are possible values of **X**. This formula may be written as

$$E(\mathbf{Y}) = E(E(\mathbf{Y} | \mathbf{X}))$$

which says that the expected value of the conditional expected value of a random variable is its current expected value. If you would like a general derivation of this result, here it is:

$$\begin{aligned}\sum_x P(\mathbf{X} = x) * E(\mathbf{Y} | \mathbf{X} = x) &= \sum_x P(\mathbf{X} = x) * \sum_y P(\mathbf{Y} = \mathrm{y} | \mathbf{X} = y) * y \\ &= \sum_x \sum_y P(\mathbf{X} = x) * P(\mathbf{Y} = y | \mathbf{X} = x) * y \\ &= \sum_x \sum_y P(\mathbf{X} = x \cap \mathbf{Y} = y) * y \\ &= \sum_y P(\mathbf{Y} = y) * y = E(\mathbf{Y})\end{aligned}$$

(The basic formula for conditional probabilities is applied at the third step.)

As another application of this expected-posterior law, let $A$ denote any event, and define **W** to be a random variable such that **W** = 1 if the event $A$ is true, but **W** = 0 if the event $A$ is false. Letting $P(A)$ denote the probability of the event $A$, we get

$$E(\mathbf{W}) = P(A) * 1 + (1 - P(A)) * 0 = P(A)$$

That is, the expected value of this 0-or-1 random variable **W** is equal to the probability of the event $A$. (Indeed, we have been estimating probabilities by AVERAGEs of such 0-or-1 random variables in spreadsheets thoughout this book.) Now suppose that **X** is any other random variable such that observing the value of **X** might cause us to revise our beliefs about the probability of $A$. Then the conditional expected value of **W** given the value of **X** would be similarly equal to the conditional probability of $A$ given the value of **X**. That is, for any number $x$ that is a possible value of the random variable **X**,

$$E(\mathbf{W} | \mathbf{X} = x) = P(\mathrm{A} | \mathbf{X} = x)$$

Thus, the general equation

$$E(\mathbf{W}) = \sum_x P(\mathbf{X} = x) * E(\mathbf{W} | \mathbf{X} = x)$$

gives us the following probability equation:

$$P(A) = \sum_x P(\mathbf{X} = x) * P(A | \mathbf{X} = x)$$

(with the sum including all $x$ that are possible values of **X**). This equation may also be written as

$$P(A) = E(P(A | \mathbf{X}))$$

This equation says that the probability of $A$, as we assess it with our current information, must equal our current expected value of what we would

Myerson_C06.qxd 07-09-2004 09:43 Page 210

think is the probability of $A$ after learning the value of **X**. Learning **X** might cause us to revise our assessment of the probability of $A$ upwards or downwards, but the weighted average of these possible revisions, weighted by their likelihoods, must be equal to our currently assessed probability $P(A)$.

In our example, let $A$ denote the event that **Y** = 32. So before **X** is learned, the prior probability of $A$ is

$$P(A) = P(\mathbf{Y} = 32) = 0.6$$

as shown in cell C8 of Figure 6.3. Then the posterior conditional probability of $A$ given the value of **X** would be either 0.5 or 0.6 or 0.667, depending on whether **X** is 70 or 75 or 80, as shown in cells J4:J6. But the respective probabilities of these three values of **X** are 0.2, 0.5, and 0.3, as shown in cells F4:F6. So before we learn **X**, our prior expected value of the posterior conditional probability of $A$ given **X** is

$$E(P(A|\mathbf{X})) = 0.2*0.5 + 0.5*0.6 + 0.3*0.667 = 0.6$$

as shown in cell J11.

The idea of posterior conditional probabilities as random variables is applied to build a simulation model of **X** and **Y** for this example in row 15 of Figure 6.3. The random variable **X** is simulated straightforwardly in cell B15 by the formula

```
= DISCRINV(RAND( ),A4:A6,F4:F6)
```

using the list of possible values of **X** in A4:A6 and the corresponding marginal probabilities in F4:F6. But if we similarly simulated **Y** using its marginal probabilities and a different RAND( ), then we would have a pair of independent random variables. So the trick is to use the posterior conditional probability distribution for **Y** given **X** that would apply given that **X** is equal to its simulated value in cell B15. These conditional probabilities are listed in cells I15:K15 by the array formula

```
{= INDEX(I4:K6,MATCH(B15,H4:H6),0)}
```

(The MATCH function indicates the number of the cell in H4:H6 where a value is found that matches B15, and then the INDEX function here returns the values in the corresponding row of I4:K6.) So cells I15:K15 always contain the appropriate conditional probability distribution for **Y** given that **X** is equal to the value in cell B15. Thus, we can simulate **Y** in cell C15 by the formula

```
= DISCRINV(RAND( ),I14:K14,I15:K15)
```

(where I14:K14 contains the list of possible values of **Y**). So the formulaic dependence of C15 on B15 (indirectly via I15:K15) has been constructed to give us the required statistical dependence to jointly simulate **X** and **Y** in this example.

This model of (**X**, **Y**) in cells B15:C15 here may be contrasted with the simulation model that we constructed in cells I26:J26 in Figure 5.2, where

Myerson_C06.qxd 07-09-2004 09:43 Page 211

neither cell was formulaically dependent on the other and so statistical dependence had to be created by using the same RAND( ) (in I24) for both random variables. But the joint probability distribution of cells I26:J26 in Figure 5.2 in is exactly the same as the joint probability distribution of cells B15:C15 in Figure 6.3 here. No statistical test could ever distinguish one pair's simulation data from the other pair's.

The lower portion of Figure 6.3 contains an analysis of simulation data from the model in cells B15:C15 here to illustrate how the law of expected posteriors holds true also for estimates of conditional expectations and conditional probabilities that are computed from simulation data. In cells I21:I23, conditional expected values of **Y** given each possible value of **X** are estimated from the simulation data. In cells K21:K23, the corresponding prior probabilities of these possible values of **X** are also estimated from this data. Then cell I25 estimates the expected value of the posterior conditional expectation of **Y** given **X**. Notice that I25 has the same value as cell I19, which directly estimates $E(\mathbf{Y})$ from the simulation data, thus again illustrating the law of expected posteriors.

**NOTE** (These calculations from selected simulation data could have been done with a data table and a filtering range like D12:D512 in Figure 6.1. But instead, cells H21:K23 in Figure 6.3 illustrate an alternative way of doing these calculations, using Excel's COUNTIF and SUMIF functions. COUNTIF(*range*, *criterion*) counts the cells in the given range that match the given criterion. SUMIF(*range1*, *criterion*, *range2*) looks in its first range for cells that match the criterion and then sums the corresponding cells in its second range.)

## 6.4 Backwards Analysis of Conditional Expectations in Tree Diagrams

The expected-posterior law allows us to compute expectations, in tree diagrams by a kind of backwards analysis of expectations, as illustrated in Figure 6.4. This tree diagram is similar to the probability trees that were introduced in Section 1.6, but it has a few differences, and so we may call it an *expectation tree*. As in a probability tree, the diagram is read from left to right and begins with a single root node at the left. Going to the right from each circular node, there are branches that represent the possible values of some random variable. The three branches from the root in Figure 6.4 represent the three possible values of the random variable **X** (from the preceding section), which may equal 70, 75, or 80. Then each node after the root is followed here by three branches that represent the three possible values of the other random variable **Y**, which may equal 30, 32, or 34. Above each branch, we list the conditional probability of the event that is represented

Myerson_C06.qxd 07-09-2004 09:43 Page 212

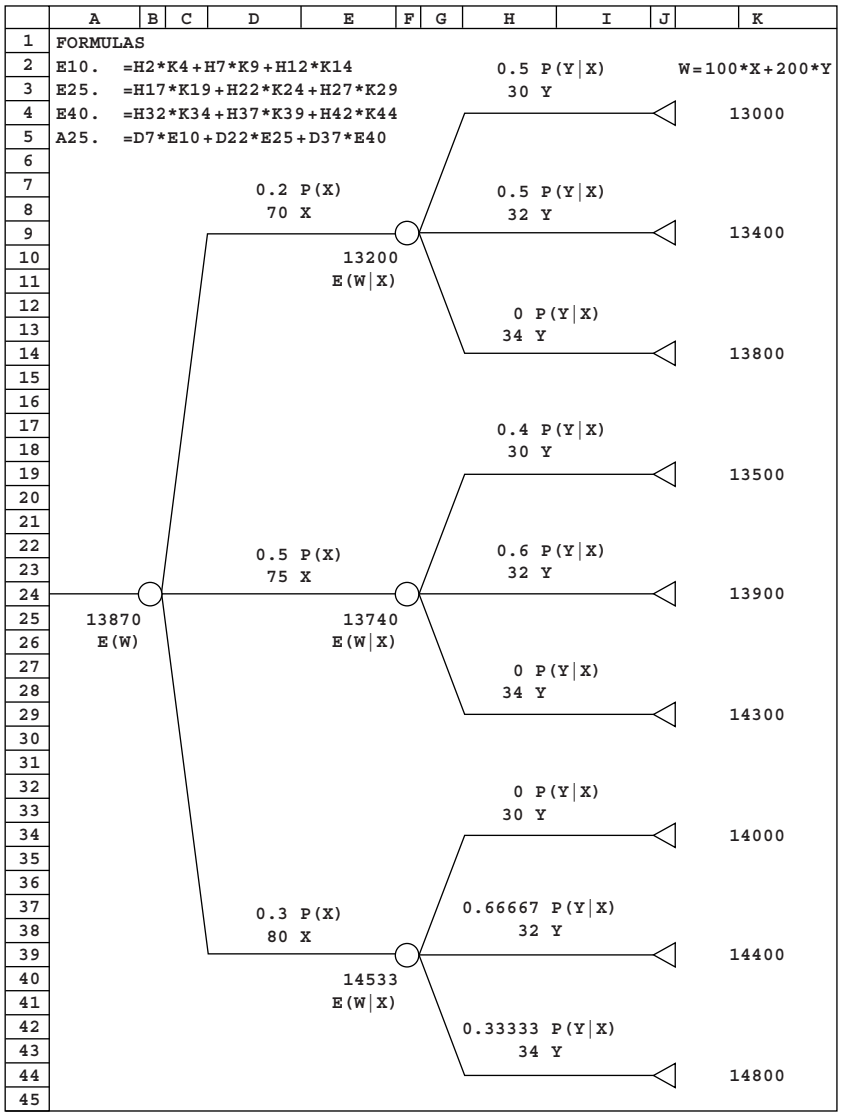


FIGURE **6.4**
**Backwards analysis of conditional expectations in a tree diagram**

by this branch given all other events on the path from the root to this branch. For example, the top branch on the right side of this diagram represents the event that $\mathbf{Y} = 30$, and it follows a branch representing the event $\mathbf{X} = 70$, and so the probability on this branch is $P(\mathbf{Y} = 30|\mathbf{X} = 70) = 0.5$. Similarly, each of the conditional probabilities for values of $\mathbf{Y}$ given $\mathbf{X}$ in cells I4:K6 of Figure 6.3 appears in the H column of Figure 6.4 over a corresponding branch of this tree diagram. For the three branches that follow the root, we list the marginal (unconditional) probability of the corresponding value of $\mathbf{X}$ (because there are no previous events to condition), and so

Myerson_C06.qxd 07-09-2004 09:43 Page 213

the branch probabilities in the D column of Figure 6.4 are the same as in cells F4:F6 of Figure 6.3.

Such expectation trees differ from the probability trees of Section 1.6 in that after the endpoints of the tree, instead of listing the intersection probabilities of the paths, we list the values of some quantity that depends on the random variables in the tree. In Figure 6.4, the values of

$$\mathbf{W} = 100*\mathbf{X} + 200*\mathbf{Y}$$

are listed at each endpoint, given the values of **X** and **Y** listed on the path to this endpoint. As in Sections 5.1 and 5.3, we may interpret **X** and **Y** as next year's price per share of two stocks, and then **W** is the value next year of a portfolio that contains 100 shares of the **X** stock and 200 shares of the **Y** stock.

Next to each node (just below it and to its left) in the expectation tree, we compute the conditional expectation of the quantity that is listed at the endpoints, given the events (if any) on the path from the root to this node. The conditional expectation $E(\mathbf{W}|\mathbf{X} = 70)$ is computed in cell E10 of Figure 6.4, next to the upper node in cell F9, because this node follows a branch representing the event $\mathbf{X} = 70$. Similarly, the conditional expectations $E(\mathbf{W}|\mathbf{X} = 75)$ and $E(\mathbf{W}|\mathbf{X} = 80)$ appear in cells E25 and E40, next to the nodes that follow the $\mathbf{X}= 75$ and $\mathbf{X} = 80$ branches. The unconditional expectation $E(\mathbf{W})$ is computed in cell A25, next to the root of the tree, because there are no "given" events on the path from the root to itself.

The conditional expectation at each node of an expectation tree is computed by multiplying the value at the end of each branch by the probability of the branch and summing these products for all branches that go out to the right from the node. For example, the conditional expectation in cell E10 of Figure 6.4 is

$$E(\mathbf{W}|\mathbf{X} = 70) = 0.5*13{,}000 + 0.5*13{,}400 + 0*13{,}800 = 13{,}200$$

which is computed in this spreadsheet by the formula

```
=H2*K4+ H7*K9+ H12*K14
```

Copying this formula from E10 to cells E25 and E40, we can similarly compute

$$E(\mathbf{W}|\mathbf{X} = 75) = 0.4*13{,}500 + 0.6*13{,}900 + 0*14{,}300 = 13{,}740$$

$$E(\mathbf{W}|\mathbf{X} = 80) = 0*14{,}000 + 0.66667*14{,}400 + 0.33333*14{,}800 = 14{,}533$$

Finally, the unconditional expectation $E(\mathbf{W})$ is computed at the root by multiplying the conditional expectation after each branch from the root times the probability of that branch and summing these products for all branches that go out to the right from the root. That is, $E(\mathbf{W})$ is computed in cell A25 of Figure 6.4 by the formula

```
=D7*E10+ D22*E25+ D37*E40
```

Myerson_C06.qxd 07-09-2004 09:43 Page 214

This formula is an application of the expected-posterior law, because it says

$$E(\mathbf{W}) = P(\mathbf{X} = 70)*E(\mathbf{W}|\mathbf{X} = 70) + P(\mathbf{X} = 75)*E(\mathbf{W}|\mathbf{X} = 75) + P(\mathbf{X} = 80)*E(\mathbf{W}|\mathbf{X} = 80)$$
$$= 0.2*13{,}200 + 0.5*13{,}740 + 0.3*14{,}533 = 13{,}870$$

(Recall that the same value of $E(\mathbf{W})$ was computed directly from the probability distribution of **W** in cell I16 of Figure 5.2.) Thus, the expected-posterior law enables us to compute expected values by backwards analysis in such tree diagrams, where the expected value at each node is computed by looking only at the probabilities and expected values on the branches to the right of this node.

## 6.5 Conditional Expectation Relationships and Correlation

We often think of correlation as a way of measuring how strong the relationship is between two random variables. But Figure 6.5 contains an example to illustrate the limitations of correlation as a way of measuring relationships between random variables. Cells B8 and C8 here contain two random variables, called **X** and **Y**, that have a correlation equal to 0 but that nevertheless have a very strong relationship derived from the formulaic dependence of **Y** on **X**. The random variable **X** is constructed in cell B8 by the formula

```
=NORMINV(RAND( ),0,2)
```

So **X** is a Normal random variable with mean 0 and standard deviation 2. The random variable **Y** is constructed in cell C8 by the formula

```
=NORMINV(RAND( ),B8^2,1)
```

So given **X**, the random variable **Y** has a conditional probability distribution that is Normal with mean **X**^2 and standard deviation 1. That is, when **X** is known, the conditional expected value of **Y** is **X**^2, the conditional standard deviation of **Y** is 1, and the remaining uncertainty about **Y** can be described by a Normal probability distribution. The marginal probability distribution of **Y** (unconditional, without information about **X**) is harder for us to describe here, but our simulation data indicate that $E(\mathbf{Y})$ is approximately 4 and *Stdev*(**Y**) is about 6. Notice that the standard deviation of **Y**, as estimated in cell C2 (6.020) is much larger than the conditional standard deviation of **Y** given **X** (1, by the formula in cell C8), because knowing **X** would reduce our uncertainty about **Y**.

Could we discover this relationship between **X** and **Y** just by looking at the simulation data? The answer is Yes, but not if we looked only at the

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | E(X) | 0.062 | 4.474 | E(Y) | | | | | |
| 2 | Stdev(X) | 2.114 | 6.020 | Stdev(Y) | | | | | |
| 3 | | | -0.018 | Correl(X,Y) | | | Event A=(x-1<=X<x+1) | | |
| 4 | | | | | | x | E(Y\|A) | P(A) | |
| 5 | | | | \|X near | X>= | (,D6) | 0.294 | 0.345 | |
| 6 | | | | 0 | -1 | -6 | 34.196 | 0.006 | |
| 7 | | X | Y | | X< | -4 | 14.367 | 0.072 | |
| 8 | SimTable | -1.159 | 0.4878 | Y\|X | 1 | -2 | 3.939 | 0.232 | |
| 9 | 0 | 0.8493 | 1.783 | 1.78303 | | 0 | 0.294 | 0.345 | |
| 10 | 0.002 | -1.377 | 1.8327 | .. | | 2 | 3.580 | 0.259 | |
| 11 | 0.004 | -1.155 | 2.2549 | .. | | 4 | 13.909 | 0.080 | |
| 12 | 0.006 | 2.6579 | 6.673 | .. | | 6 | 30.777 | 0.006 | |
| 13 | 0.008 | -1.429 | 2.3493 | .. | | | | | |
| 14 | 0.01 | 4.5172 | 21.261 | .. | | | E(E(Y\|X)) | sum(probys) | |
| 15 | 0.012 | -1.654 | 1.426 | .. | | | 4.474 | 1 | |
| 16 | 0.014 | -0.108 | 0.0484 | 0.04838 | FORMULAS | | | | |
| 17 | 0.016 | 2.0767 | 4.2103 | .. | B8. =NORMINV(RAND(),0,2) | | | | |
| 18 | 0.018 | 1.1512 | 0.7283 | .. | C8. =NORMINV(RAND(),B8^2,1) | | | | |
| 19 | 0.02 | -2.113 | 4.2104 | .. | B1. =AVERAGE(B9:B509) | | | | |
| 20 | 0.022 | 2.2322 | 6.8581 | .. | B2. =STDEV(B9:B509) | | | | |
| 21 | 0.024 | 1.2106 | -0.53 | .. | B1:B2 copied to C1:C2 | | | | |
| 22 | 0.026 | 0.9715 | 1.4704 | 1.47041 | C3. =CORREL(B9:B509,C9:C509) | | | | |
| 23 | 0.028 | -1.164 | 0.1279 | .. | E6. =D6-1 | | | | |
| 24 | 0.03 | 0.0813 | 1.0392 | 1.03924 | E8. =D6+1 | | | | |
| 25 | 0.032 | 1.0612 | -0.409 | .. | D9. =IF(AND($E$6<=B9,B9<$E$8),C9,"..") | | | | |
| 26 | 0.034 | -0.135 | 1.2979 | 1.29789 | D9 copied to D9:D509 | | | | |
| 27 | 0.036 | -2.346 | 3.3985 | .. | G5. =AVERAGE(D9:D509) | | | | |
| 28 | 0.038 | -1.874 | 4.677 | .. | H5. =COUNT(D9:D509)/COUNT(C9:C509) | | | | |
| 29 | 0.04 | -2.714 | 8.053 | .. | G6:H12. {=TABLE(,D6)} | | | | |
| 30 | 0.042 | -3.342 | 11.355 | .. | G15. =SUMPRODUCT(G6:G12,H6:H12) | | | | |
| 31 | 0.044 | 0.7178 | -0.204 | -0.2038 | H15. =SUM(H6:H12) | | | | |
| 32 | 0.046 | -0.136 | 1.1187 | 1.11874 | | | | | |
| 33 | 0.048 | -0.278 | -1.473 | -1.4734 | | | | | |
| 34 | 0.05 | -2.774 | 7.5497 | .. | | | | | |
| 35 | 0.052 | 1.4922 | 3.8439 | .. | | | | | |
| 36 | 0.054 | -4.495 | 22.321 | .. | | | | | |
| 37 | 0.056 | -2.381 | 5.5311 | .. | | | | | |
| 38 | 0.058 | 0.9376 | 0.401 | 0.40096 | | | | | |
| 39 | 0.06 | -0.111 | -1.163 | -1.163 | | | | | |
| 40 | 0.062 | -1.705 | 4.6293 | .. | | | | | |
| 41 | 0.064 | 0.364 | 1.7449 | 1.7449 | | | | | |
| 42 | 0.066 | -3.539 | 12.021 | .. | | | | | |
| 43 | 0.068 | 0.9495 | 0.1189 | 0.11893 | | | | | |

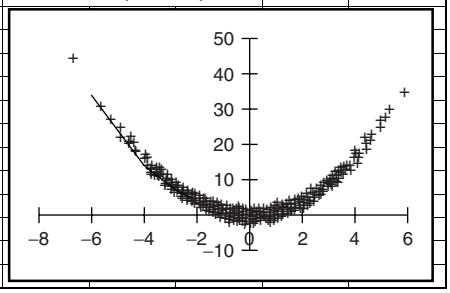


FIGURE **6.5**
**An example of dependence among random variables with correlation 0**

correlation. The correlation between **X** and **Y** in this spreadsheet's simulation data is –0.018, as computed by CORREL in cell C3 of Figure 6.5, and this value is statistically indistinguishable from 0. (You can check that the two columns of 501 independent RAND( )s have a correlation that differs from 0 by more than 0.018 about 70% of the time.) The strong relationship between **X** and **Y** can be found from this simulation data when we see how the conditional expected value of **Y** depends on the value of **X**.

No single value of **X** appears more than once in our data because **X** is a continuous random variable, and so we may have difficulties estimating the conditional expected value of **Y** given any single value of **X** from the data. But our simulation data can generate good estimates of the conditional expected values of **Y** given that **X** lies in any interval that occurred

Myerson_C06.qxd 07-09-2004 09:43 Page 216

many times in this data. Based on this simulation data, for each even integer $x$ from –6 to 6, cells G6:G12 in Figure 6.5 show the conditional expected values of **Y** given the event that the random variable **X** is between $x - 1$ and $x + 1$. That is, this table in Figure 6.5 contains our estimates of the conditional expected value of **Y** given the value of **X** rounded to the nearest even integer. The strong quadratic relationship between **X** and the conditional expected value of **Y** is readily apparent in the table. It is also clear why this relationship is not manifested in the correlation, because the tendency of **Y** to increase when **X** is increased among the positive numbers is counterbalanced by an equal tendency of **Y** to decrease when **X** is increased among the negative numbers. The equality between cells G15 and C1 in Figure 6.5 is another illustration of the expected-posterior law.

## 6.6 Uncertainty About a Probability

We have argued that a decision maker should always be able to describe uncertainty by probabilities. But now imagine a situation where we are working as decision analysis consultants for a decision maker who expresses great uncertainty about the probability of some event $A$ that is important in the decision problem. In such a case, you might worry about whether we should treat the subjective probability of this event as an unknown quantity $\mathbf{P}(\mathbf{A})$ in our model. But if we did so, then we would have to assess its probability distribution, and so we would have to ask new questions about more probabilities (such as $P(\mathbf{P}(\mathbf{A}) < 1/2)$), about which the decision maker may express even more uncertainty.

Saying that the decision maker's subjective probability of the event is some number $P(A)$ means that, in the decision maker's judgment, the event $A$ is as likely to occur as a RAND( ) is to be between 0 and $P(A)$ after its next recalculation. That is, if there were no personal consequences depending on the event $A$, then the decision maker would be indifferent between a promise to pay \$1000 if $A$ occurs and a promise to pay \$1000 if a particular RAND( ) is less than $P(A)$. Thus, the subjective probability of an event is a measure of how likely the event seems to the decision maker and as such is an aspect of feelings about the world that this person should be able to assess by pure introspection. Of course, introspection is often difficult, but in principle there is nothing to prevent the person from assessing the probability of the event.

When people say that they do not know the probability of some event, it may be because they are not speaking about their own personal subjective probability of this event, but rather the probability of this event that would be assessed by someone who has more expertise in this area. In such cases, decision makers may assume that their own inferior judgment is less

Myerson_C06.qxd 07-09-2004 09:43 Page 217

important than what the experts think. And if the experts are not available to say what they think, then decision makers may have very real uncertainty about it. In such situations, decision makers may imagine questioning an expert on this subject and getting some advice $\mathbf{X}$ that would cause them to adjust their beliefs about the probability of $A$ to some number $P(A|\mathbf{X})$. But until they actually learn this expert's advice, this posterior probability $P(A|\mathbf{X})$ that they would assess with such advice would be an unknown quantity, and their beliefs about $P(A|\mathbf{X})$ should be described by some probability distribution.

The expected-posterior law says that the decision maker's current assessment of $P(A)$ must be equal to the expected value of the unknown posterior $P(A|\mathbf{X})$, that is,

$$P(A) = E(P(A|\mathbf{X}))$$

If the decision maker actually has no possibility of any communication with experts, then this probability $P(A)$ is all that we need in our decision-analysis work, and there is no need to be concerned about the distribution of the unknown quantity $P(A|\mathbf{X})$, because it has no further relevance to the decision.

For example, suppose that our decision maker is a venture capitalist who is considering a risky investment project that would cost \$5 million and that would subsequently generate either \$10 million in revenues if the project is a success or no revenues (\$0) if the project is not successful. Suppose that the decision maker is risk-neutral but is worried about what an expert in this field would say about the chances of success for this project. Let $A$ denote the event that the project will be successful, and let $P(A|\mathbf{X})$ denote the probability of $A$ that the decision maker would assess after getting advice $\mathbf{X}$ from an expert. Since the expert's advice $\mathbf{X}$ is not actually available to the decision maker, $P(A|\mathbf{X})$ is an unknown quantity. But suppose that the venture capitalist has dutifully answered our questions about the probability of $P(A|\mathbf{X})$ being above or below various numbers, and on the basis of these answers we have assessed that beliefs about the unknown probability $P(A|\mathbf{X})$ can be described by a Beta probability distribution with mean 0.4 and standard deviation 0.2. (A probability density for such a Beta distribution is shown in the inset chart at the bottom-right of Figure 6.6. This chart was made in a spreadsheet similar to Figure 4.1, but with BETINV instead of NORMINV.)

Saying that $P(A|\mathbf{X})$ is a random variable with mean 0.4 means that 0.4 is the decision maker's current expected value of the posterior probability of success after getting expert advice. That is,

$$E(P(A|\mathbf{X})) = 0.4$$

The decision maker's subjective probability of $A$, given current information only, must be equal to this expected posterior, and so

$$P(A) = 0.4$$

Myerson_C06.qxd 07-09-2004 09:43 Page 218

| | A | B | C | D | E | F | G | H | I | J |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | We do not know what an expert would say is project's P(Success). | | | | | | | | | |
| 2 | 0.4 | Our expected value of expert's P(Success) | | | | | | | | |
| 3 | 0.2 | Our standard deviation of expert's P(Success) | | | | | | | | |
| 4 | | | | | | | | | | |
| 5 | We get payoff 0 if we do not do the project. | | | | | | | | | |
| 6 | From doing the project, we get | | | | | | | | | |
| 7 | 5 | payoff if Success | | | | | | | | |
| 8 | -5 | payoff if Not Success | | | | | | | | |
| 9 | | | | | | | | | | |
| 10 | E | 0.42 | 0.396914 | 0.393 | | -1.074 | 0.4296 | 0.2912 | E | |
| 11 | Stdev | 0.494 | 0.200364 | 0.489 | | 4.8857 | 2.7047 | 0.124 | Stdev | |
| 12 | MODEL 1 | | MODEL 2 | | | | | | | |
| 13 | | Succ? | Exprt'sP | Succ? | | $ from project | | | | |
| 14 | SimTab | 0 | 0.639231 | 0 | | always | ifP>.5 | P | P<.5 | |
| 15 | 0 | 1 | 0.692971 | 1 | | 5 | 5 | .. | | |
| 16 | 0.001 | 1 | 0.464805 | 1 | | 5 | 0 | 0.4648 | | |
| 17 | 0.002 | 0 | 0.688664 | 1 | | 5 | 5 | .. | | |
| 18 | 0.003 | 1 | 0.255294 | 0 | | -5 | 0 | 0.2553 | | |
| 1012 | 0.997 | 1 | 0.047693 | 0 | | -5 | 0 | 0.0477 | | |
| 1013 | 0.998 | 0 | 0.27682 | 1 | | 5 | 0 | 0.2768 | | |
| 1014 | 0.999 | 0 | 0.482012 | 0 | | -5 | 0 | 0.482 | | |
| 1015 | 1 | 0 | 0.465665 | 0 | | -5 | 0 | 0.4657 | | |
| 1016 | | | | | | | | | | |
| 1017 | FORMULAS | | | | | | | | | |
| 1018 | B14. =IF(RAND()<A2,1,0) | | | | | | | | | |
| 1019 | C14. =BETINV(RAND(),A2,A3) | | | | | | | | | |
| 1020 | D14. =IF(RAND()<C14,1,0) | | | | | | | | | |
| 1021 | F15. =IF(D15=1,$A$7,$A$8) | | | | | | | | | |
| 1022 | G15. =IF(C15>0.5,F15,0) | | | | | | | | | |
| 1023 | H15. =IF(C15<0.5,C15,"..") | | | | | | | | | |
| 1024 | F15:H15 copied to F15:H1015 | | | | | | | | | |
| 1025 | B10. =AVERAGE(B15:B1015) | | | | | | | | | |
| 1026 | B11. =STDEV(B15:B1015) | | | | | | | | | |
| 1027 | B10:B11 copied to C10:D11 | | | | | | | | | |
| 1028 | B10:B11 copied to F10:H11 | | | | | | | | | |

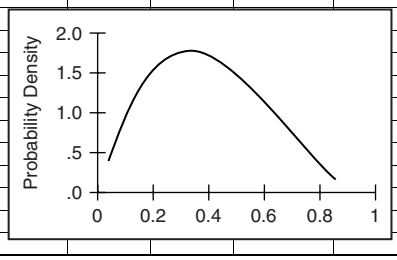


FIGURE **6.6**
**Modeling uncertainty about a probability**

Figure 6.6 compares two models of this situation. In Model 1, we use this assessment that the decision maker's current subjective probability of success is 0.4, but we ignore all the uncertainty about what an expert might tell us about the probability of success. With this probability 0.4 entered into cell A2, the event of a successful product can be simulated in cell B14 by the formula

=IF(RAND( )<A2,1,0)

In Model 2, we take account of the uncertainty about what an expert might tell us about the probability of success. In cell C14, the unknown posterior probability of $A$ that we would assess after getting expert advice is simulated as a Beta random variable with mean 0.4 (in cell A2) and standard deviation 0.2 (in cell A3) by the formula

=BETINV(RAND( ),A2,A3)

Then in cell D14, the event of a success can be simulated by the formula

=IF(RAND( )<C14,1,0)

Myerson_C06.qxd 07-09-2004 09:43 Page 219

Notice that this formula makes the conditional probability of a simulated success in cell D14 given C14 equal to the simulated posterior probability $P(A|\mathbf{X})$ in cell C14.

Model 2 is more complicated than Model 1, but it is not necessarily better. If there really is no possibility of any communication with such an expert, then the additional complexity of Model 2 contributes nothing to our decision analysis. In fact, if we were given a column of simulation data that came from either cell B14 (Model 1) or cell D14 (Model 2) in Figure 6.6, there would be no statistical test that could indicate which cell had generated the data. Independent simulations of either cell would be expected to generate about 40% 1s and about 60% 0s, shuffled together randomly. The fact that cell D14 depends on another random probability does not make it any riskier, and the simulation data from cells B14 and D14 both have standard deviations of about 0.49. (*Note:* Rows 19 to 1011 of the spreadsheet have been hidden in Figure 6.6 by Excel's Data:Group command.)

In this example, the decision maker's expected net profit from the project would be \$5 million if there is a success and –\$5 million if there is not a success. Because the probability of success is 0.4, the decision maker's expected payoff is $0.4*5 + (1 - 0.4)*(-5) = -1$. So both simulation models should generate expected payoffs of –\$1 million (with a standard deviation of about \$4.9 million), as illustrated for Model 2 in the F column of Figure 6.6. Since Model 1 is simpler and would yield the same results, it may be considered the better model. Thus, we may simply advise the decision maker to assess the subjective probability of success on the basis of available information and to not be concerned about unavailable expert opinions.

But if there is some possibility of communication with such experts, then the uncertainty about what they would say may be important, in which case Model 2 should be used. If an expert could be paid for advice about the probability of success, then the decision maker would want to do this project when the expert suggested a probability of success greater than 1/2. The results of this strategy are computed for the simulation data from Model 2 in the G column of Figure 6.6, and the expected payoff from this strategy is found to be about 0.43 in cell G10. So this analysis suggests that the decision maker should be willing to pay up to about \$0.43 million for advice from an expert who can determine $P(A|\mathbf{X})$ before deciding about this project. We will see more about such *value of information* calculations in Section 7.2.

There is another important case where the unavailable expert opinions may be very important to take into account. Suppose that experts who know $P(A|\mathbf{X})$ may be working for our decision maker's competitors, and these competitors would have already taken advantage of this investment opportunity if they knew that $P(A|\mathbf{X})$ were greater than 1/2. Then the fact that the project has become available to our decision maker may itself be evidence that $P(A|\mathbf{X})$ is less than 1/2. If so, then the decision maker should take this fact into account and recompute $P(A)$ as the conditional expectation of $P(A|\mathbf{X})$ given the information that $P(A|\mathbf{X}) < 0.5$. Cells H15:H10105

Myerson_C06.qxd 07-09-2004 09:43 Page 220

in Figure 6.6 echo the simulated expert probabilities from Model 2 for only those cases where the probability is less than 1/2, and nonnumerical text is displayed otherwise. Then cell H10 computes the average of these cells to be about 0.29. So this analysis suggests that the conditional expectation of $P(A|\mathbf{X})$ given that it is less than 1/2 is approximately 0.29. Thus (by another application of the expected-posterior law), if the availability of the project tells our decision maker that $P(A|\mathbf{X})$ is less than 1/2, then the conditional probability of success given the project's availability should be

$$P(A|P(A|\mathbf{X}) < 1/2) = E(P(A|\mathbf{X})|P(A|\mathbf{X}) < 1/2) \approx 0.29$$

So if the availability of this investment opportunity may be an indication that experts are not optimistic about it, then the decision maker may need to significantly reduce his subjective probability of its success. This important effect is called the *winner's curse* and is considered further in Section 7.5.

## 6.7 Linear Regression Models

We now consider a special form of formulaic dependence among random variables that is probably assumed by statisticians more often than any other: the *linear regression model*. In a linear regression model, a random variable $\mathbf{Y}$ is made dependent on other variables $(\mathbf{X}_1, \ldots, \mathbf{X}_K)$, which are called the *explanatory variables* for $\mathbf{Y}$. The conditional expected value of $\mathbf{Y}$ given these explanatory variables is assumed to be linear function that can be written

$$E(\mathbf{Y}|\mathbf{X}_1, \ldots, \mathbf{X}_K) = \alpha_0 + \alpha_1 * \mathbf{X}_1 + \cdots + \alpha_K * \mathbf{X}_K$$

for some nonrandom numbers $\alpha_0, \alpha_1, \ldots, \alpha_K$. The conditional standard deviation of $\mathbf{Y}$ given the explanatory variables in the regression model is assumed to be a constant $\sigma$ that is called the *standard error* of the regression. It is also assumed that the conditional distribution of $\mathbf{Y}$ given these explanatory variables is a Normal probability distribution. That is, the linear regression model assumes that the conditional distribution of $\mathbf{Y}$ given the explanatory variables is a Normal probability distribution with a mean that is a linear function of the explanatory variables and a standard deviation that is a constant (independent of the explanatory variables). To define such a linear regression model where $\mathbf{Y}$ depends on $K$ other random variables, we must specify $K + 2$ parameters: the coefficients for each of the $K$ explanatory random variables in the linear function $(\alpha_1, \ldots, \alpha_K)$, the $\mathbf{Y}$ axis intercept of the linear function $(\alpha_0)$, and the standard error of the regression $(\sigma)$. Once these parameters are specified, the dependent random variable can be simulated in a spreadsheet by the formula

$$\text{NORMINV(RAND( ), } \alpha_0 + \alpha_1 * \mathbf{X}_1 + \cdots + \alpha_K * \mathbf{X}_K, \sigma)$$

An example of a linear regression model is shown in Figure 6.7, where a random variable $\mathbf{Y}$ in cell B12 is made to depend on a random variable $\mathbf{X}$

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | Parameters: | | | | Computed from parameters: | | | |
| 2 | for X | for Y given X | | | E(Y) | Covar(X,Y) | | |
| 3 | E(X) | X coefficient | | | 20 | 12 | | |
| 4 | 10 | 3 | | | Stdev(Y) | Correl(X,Y) | Corands | |
| 5 | Stdev(X) | Intercept | | | 10 | 0.6 | 0.3644 | 0.3810 |
| 6 | 2 | -10 | | | | | X | Y |
| 7 | | Std error | | | | | 9.306551 | 16.97084 |
| 8 | | 8 | | | | | | |
| 9 | | | | | | | | |
| 10 | Data: | | | | Computed from data: | | | |
| 11 | X | Y | | Estimates for X | | Estimates for Y given X | | |
| 12 | 12.05 | 16.53 | | | E(X) | Regressn Array ... 7 Rows x 1 | | |
| 13 | 11.12 | 18.26 | | | 10.048 | X(1) Coefficient | | |
| 14 | 10.15 | 26.93 | | | Stdev(X) | 2.889 | | |
| 15 | 11.54 | 22.85 | | | 1.770 | Intercept | | |
| 16 | 11.57 | 24.48 | | | | -8.372 | | |
| 17 | 10.07 | 23.57 | | | | Std Error | | |
| 18 | 5.87 | -6.71 | FORMULAS | | | 7.131 | | |
| 19 | 10.33 | 31.02 | A12. =NORMINV(RAND(),$A$4,$A$6) | | | | | |
| 20 | 6.00 | 22.79 | B12. =NORMINV(RAND(), $B$6+$B$4*A12, $B$8) | | | | | |
| 21 | 12.35 | 30.34 | A12:B12 copied to A12:B31 | | | | | |
| 22 | 8.38 | 11.64 | E13. =AVERAGE(A12:A31) | | | | | |
| 23 | 10.70 | 22.64 | E15. =STDEV(A12:A31) | | | | | |
| 24 | 11.23 | 21.85 | F12:F18. {=REGRESSN(A12:A31,B12:B31)} | | | | | |
| 25 | 10.80 | 33.75 | E3. =B6+B4*A4 | | | | | |
| 26 | 9.70 | 23.16 | F3. =B4*A6^2 | | | | | |
| 27 | 8.37 | 13.79 | E5. =((A6*B4)^2+B8^2)^0.5 | | | | | |
| 28 | 11.44 | 22.08 | F5. =F3/(A6*E5) | | | | | |
| 29 | 9.94 | 13.49 | G5:H5. {=CORAND(F5)} | | | | | |
| 30 | 10.01 | 21.55 | G7. =NORMINV(G5,A4,A6) | | | | | |
| 31 | 9.31 | 19.22 | H7. =NORMINV(H5,E3,E5) | | | | | |

FIGURE **6.7**
**Estimating a linear regression relationship among random variables**

in cell A12. The value 3 in cell B4 is the coefficient of **X** in the linear function, the value –10 in cell B6 is the **Y** axis intercept of the linear function, and the value 8 in cell B8 is the standard error of the regression. So with these parameters in cells B4, B6, and B8, the formula for **Y** in cell B12 is

=NORMINV(RAND( ), $B$6+$B$4∗A12, $B$8)

Entering the value of the given parameters would make this formula

=NORMINV(RAND( ), –10+3∗A12, 8)

Adding a constant to a Normal random variable yields a Normal random variable with mean increased by the amount of this constant. So the formula for **Y** in cell B12 could have been written equivalently as

=$B$6+$B$4∗A12 + NORMINV(RAND( ),0,$B$8)

or

=–10+3∗A12 + NORMINV(RAND( ),0,8)

That is, the dependence of **Y** on **X** could be described by saying that **Y** is a linear function of **X** plus a Normal random variable with mean 0 that may be called a *Normal error term*.

Myerson_C06.qxd 07-09-2004 09:43 Page 222

The value of the **X** in cell A12 of Figure 6.7 is generated by the formula

=NORMINV(RAND( ),$A$4,$A$6)

where cell A4 contains the value 10 and cell A6 contains the value 2. That is, the explanatory variable **X** here is a Normal random variable with mean 10 and standard deviation 2.

The formulas from A12:B12 in Figure 6.7 have been copied to fill the range A12:B31 with 20 independent pairs that have the same joint distribution. That is, in each row from 13 to 31, the cells in the A and B columns are drawn from the same joint distribution as **X** and **Y** in A12:B12. In each of these rows, the value of the B cell depends on the value of the A cell according to the linear regression relationship specified by the parameters in cells B4, B6, and B8, but the values in each row are independent of the values in all other rows.

The fundamental problem of statistical regression analysis is to look at statistical data like the 20 pairs in A12:B31 and guess the linear regression relationship that underlies the data. In such statistical analysis, we can look only at the values of these random variables because statisticians cannot see underlying formulas that generated their data. Excel offers several different ways of doing statistical regression analysis. In particular, with the Excel's Analysis ToolPak added in, you can do regression analysis by the menu commands Tools:DataAnalysis:Regression. But for a first illustration of regression analysis, Figure 6.7 uses a simple regression analysis function called REGRESSN that is added by Simtools.

The REGRESSN function has two parameters, where the first is the XDataRange and second is the YDataRange. The YDataRange should be a range of cells in one column that contain the sample data for the dependent random variable **Y**. When **Y** is dependent on *K* explanatory variables in the regression, the XDataRange should be a range with *K* columns (one for each explanatory variable) and the same number of rows as the YDataRange. With *K* explanatory variables, the REGRESSN function should be entered as an array function filling 7 rows by *K* columns, where it will return its best estimate of the regression relationship between the dependent variable and the *K* explanatory variables. In this example, our random variable **Y** is dependent on just one explanatory variable, **X**, and so we have $K = 1$. The XDataRange is A12:A31, the YDataRange is B12:B31, and the array formula

{= REGRESSN(A12:A31,B12:B31)}

has been entered into the seven cells F12:F18 to display the estimated regression parameters: the coefficient of **X** in the regression, the **Y** axis intercept, and the standard error of the regression.

If you try recalculating the spreadsheet in Figure 6.7 so that the simulated **X** and **Y** data are regenerated in cells A12:B31, you will see the regression estimates in F12:F18 varying around the true underlying parameters that we can see in B3:B8. More sophisticated regression analysis software

Myerson_C06.qxd 07-09-2004 09:43 Page 223

(such as that in Excel's Analysis ToolPak) will generate 95% confidence intervals for the regression parameters around these estimates.

Regression analysis in F12:F18 here only assesses the parameters of the conditional probability distribution of **Y** given **X**. So to complete our estimation of the joint distribution of **X** and **Y**, we just need to estimate the parameters of the marginal probability distribution of **X**, as we have done in earlier chapters. In particular, if we assume that **X** is drawn from a Normal probability distribution but we do not know its mean and standard deviation, then we can use the sample average and the sample standard deviation of our **X** data, as shown in cells E13 and E15 of Figure 6.7, to estimate the mean and standard deviation parameters of the Normal distribution that generated our **X** data.

In a linear regression model for **Y**, if the explanatory variables are Multivariate Normal, then the joint distribution of **Y** with these explanatory variables is also Multivariate Normal. In the special case of a linear regression model with just one explanatory variable, if the explanatory variable is Normal, then its joint distribution with **Y** is Multivariate Normal. Thus, because **X** has been given as a Normal random variable in cell A12, the random variables **X** and **Y** in cells A12 and B12 of Figure 6.7 are Multivariate Normal random variables,

The difference between the treatment of Multivariate Normals here and in Chapter 5 is the parametric form that we use to characterize these Multivariate Normal random variables. Here the parameters are the mean $E(\mathbf{X})$ and standard deviation $Stdev(\mathbf{X})$, which characterize the marginal distribution of **X** and the intercept, coefficient, and standard error parameters of the linear regression model, which characterize the conditional probability distribution of **Y** given **X**. But in Chapter 5 we parameterized the Multivariate Normal random variables by their individual means and standard deviations and their pairwise correlations. In cases like this, where we have only one explanatory variable **X** in the regression model, the other parameters of the Multivariate Normal distribution of **X** and **Y** can be computed from the regression parameters by some equations that can be mathematically derived from equations (3)–(6) of Chapter 5. To be specific, suppose that **Y** depends on **X** in a linear regression model of the form

$$E(\mathbf{Y}|\mathbf{X}) = \alpha_0 + \alpha_1 * \mathbf{X} \quad \text{and} \quad Stdev(\mathbf{Y}|\mathbf{X}) = \sigma$$

where $Stdev(\mathbf{Y}|\mathbf{X})$ denotes the conditional standard deviation of **Y** given **X**. Then the marginal probability distribution of **Y** satisfies

$$E(\mathbf{Y}) = \alpha_0 + \alpha_1 * E(\mathbf{X})$$
$$Stdev(\mathbf{Y}) = ((\alpha_1 * Stdev(\mathbf{X}))\text{^}2 + \sigma\text{^}2)\text{^}0.5$$
$$Correl(\mathbf{X}, \mathbf{Y}) = \alpha_1 * Stdev(\mathbf{X})/Stdev(\mathbf{Y})$$

These equations are applied to compute these parameters in E2:F5 of Figure 6.7. These computed parameters have been used with CORAND (in G5:H5) to

Myerson_C06.qxd 07-09-2004 09:43 Page 224

make a simulation of **X** and **Y** (in cells G7:H7) that is completely equivalent to the simulation that we made with regression in cells A12:B12. That is, no statistical test could distinguish between data taken from cells A12:B12 and data taken from cells G7:H7.

Notice in Figure 6.7 that the standard deviation of **Y** (10, in cell E5) is greater than the conditional standard deviation of **Y** given **X** (8, the standard error in cell B8). This inequality holds because the standard deviation is a measure of uncertainty, and learning the value of the explanatory variable **X** would reduce our uncertainty about the dependent variable **Y**.

## 6.8 Regression Analysis and Least-Squared Errors

Regression analysis is a powerful statistical method that you should study in detail in another course. But somewhat more needs to be said here about how the estimates of a statistical regression analysis are determined from the data. So consider again the annual growth ratio data for Funds 1 and 2 from Chapter 5, shown here in cells B2:C11 of Figure 6.8. Let us try now to estimate a regression model in which the growth ratio of Fund 2 depends on the growth ratio of Fund 1 in each year.

We start by searching for a linear function that can be used to predict the growth ratio of Fund 2 as a function of the growth ratio of Fund 1. In Figure 6.8, the values of cells A14 and B14 will be interpreted as the intercept and coefficient of this linear function. Then, based on the 1980 growth ratio of Fund 1 in cell B2, our linear estimate of estimate of Fund 2's growth ratio in 1980 can be computed by the formula

```
=$A$14+$B$14*B2
```

which has been entered into cell D2. Then copying cell D2 to D2:D11 gives us a column of linear estimates of Fund 2's growth based only on the growth of Fund 1 in the same year.

The numerical value of these estimates will depend, of course, on the intercept and coefficient that we have entered into cells A14 and B14. For example, if we enter the value 0 into cell A14 and the value 1 into cell B14, then our linear estimate of Fund 2's growth will be just the same as Fund 1's growth for each year. Unfortunately, no matter what intercept and coefficient we may try, our linear estimates in D2:D11 will be wrong (different from C2:C11) in most or all the years for which we have data. But we can try to find an intercept and coefficient that will generate linear estimates which have the smallest overall errors, in some sense.

The overall measure of errors that statisticians use in regression analysis is an adjusted average of the squared errors of our linear estimates. The

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | | Fund 1 | Fund 2 | Linear estimate of 2 | Squared error | | |
| 2 | 1980 | 1.22 | 1.28 | 1.2427 | 0.00141 | | |
| 3 | 1981 | 1.09 | 1.20 | 1.0962 | 0.01036 | | |
| 4 | 1982 | 0.92 | 0.83 | 0.9033 | 0.00509 | | |
| 5 | 1983 | 1.37 | 1.48 | 1.4226 | 0.00321 | | |
| 6 | 1984 | 1.01 | 0.96 | 1.0053 | 0.00221 | | |
| 7 | 1985 | 1.15 | 1.22 | 1.1645 | 0.00263 | | |
| 8 | 1986 | 1.26 | 1.15 | 1.2921 | 0.01982 | | |
| 9 | 1987 | 1.25 | 1.17 | 1.2856 | 0.01309 | | |
| 10 | 1988 | 0.93 | 0.96 | 0.9163 | 0.00153 | | |
| 11 | 1989 | 1.19 | 1.31 | 1.2198 | 0.00759 | | |
| 12 | | | | | | | |
| 13 | Intrcep | Coefficient | | Adjusted avg sq err | | DegreesOfFreedom | |
| 14 | -0.131 | 1.1305 | | | 0.00837 | 8 | |
| 15 | | | | | Std err | | |
| 16 | SOLVER: minimize E14 by changing A14:B14 | | | | 0.0915 | | |
| 17 | | | | | FORMULAS | | |
| 18 | | | FUND2 | FUND1 | D2. =$A$14+$B$14*B2 | | |
| 19 | | | Regressn Array ... 7 Rows x | | E2. =(D2-C2)^2 | | |
| 20 | | | X(1) Coefficient | | E2:D2 copied to E2:D11 | | |
| 21 | | | 1.1305 | F14. =COUNT(E2:E11)-COUNT(A14:B14) | | | |
| 22 | FUND1 Mean | | Intercept | | E14. =SUM(E2:E11)/F14 | | |
| 23 | 1.1375 | | -0.131 | | E16. =E14^0.5 | | |
| 24 | FUND1 Stdev | | Std Error | | B23. =AVERAGE(B2:B11) | | |
| 25 | 0.1520 | | 0.0915 | | B25. =STDEV(B2:B11) | | |
| 26 | | | | C19:C25. {=REGRESSN(B2:B11,C2:C11)} | | | |
| 27 | Simulation model: | | | B29. =NORMINV(RAND(),B23,B25) | | | |
| 28 | | Fund 1 | Fund 2 | C29. =NORMINV(RAND(),C23+C21*B29,C25) | | | |
| 29 | SimTabl | 1.213 | 1.246 | | | | |
| 30 | 0 | 0.9737 | 1.0909 | | | | |
| 31 | 0.005 | 1.0162 | 0.9851 | | | | |
| 32 | 0.01 | 0.9525 | 0.9901 | | | | |
| 33 | 0.015 | 0.9306 | 0.9358 | | | | |
| 34 | 0.02 | 1.3578 | 1.4118 | | | | |
| 35 | 0.025 | 1.2084 | 1.1936 | | | | |
| 36 | 0.03 | 1.015 | 1.117 | | | | |
| 37 | 0.035 | 1.172 | 1.1917 | | | | |
| 38 | 0.04 | 1.05 | 0.9767 | | | | |
| 39 | 0.045 | 1.2618 | 1.4721 | | | | |
| 40 | 0.05 | 1.1637 | 1.0842 | | | | |
| 41 | 0.055 | 0.914 | 0.9924 | | | | |
| 42 | 0.06 | 0.9896 | 0.9103 | | | | |
| 43 | 0.065 | 1.3599 | 1.5734 | | | | |
| 44 | 0.07 | 1.2124 | 1.3562 | | | | |
| 45 | 0.075 | 1.2794 | 1.2302 | | | | |

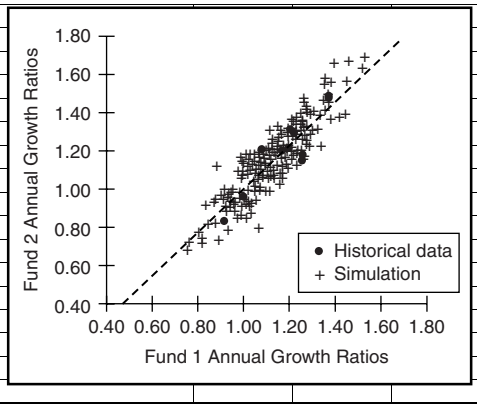


FIGURE **6.8**
**Minimizing squared errors of linear estimates in regression analysis**

*squared error* of the linear estimate in 1980 is computed in cell E2 of Figure 6.8 by the formula

= (D2–C2)^2

and cell E2 has then been copied to cells E2:E11. The adjusted average squared error is computed in cell E14 by the formula

=SUM(E2:E11)/(COUNT(E2:E11)–COUNT(A14:B14))

That is, the *adjusted average* differs from the true average in that we divide, not by the number of data points, but by the number of data

Myerson_C06.qxd 07-09-2004 09:43 Page 226

points minus the number of linear parameters that we have to find in the linear function. This denominator (number of data points minus number of parameters in the linear function) is called the *degrees of freedom* in the regression.

Now let us ask Solver to minimize this adjusted average squared error in cell E14 by changing the linear estimation parameters in cells A14:B14. The resulting optimal solution that Solver returns for this optimization problem is shown in cells A14:B14 of Figure 6.8 (an intercept of –0.1311 and a coefficient of 1.1305). This optimal solution is exactly the same as the intercept and coefficient that the REGRESSN function returns in cells C21 and C19 because statistical regression analysis is designed to find the linear function that minimizes the adjusted average squared error in this sense.

Cell E16 in Figure 6.8 computes that square root of the adjusted average squared error, with the formula

```
=E14^0.5
```

Notice that the result of this formula in cell E16 is exactly the same as the standard error of the regression that REGRESSN returns in cell C25 (0.0915). That is, the estimated standard error of the regression is the square root of the minimal adjusted average squared error.

If we had used the simple average in cell E14 (dividing by the number of data points, instead of dividing by the number of degrees of freedom), then Solver would have still given us the same optimal intercept and coefficient in A14:B14 when we asked it to minimize E14, but the square root of the average squared error would have been slightly smaller. Regression analysis uses the adjusted average squared error (dividing by the number of degrees of freedom) because the unadjusted average would bias downwards our estimate of the variance of the Normal errors in the regression.

In cells B29:C29 of Figure 6.8, this statistically estimated regression model is applied to make a simulation model for forecasting the growth ratios of Funds 1 and 2 in future years. The growth ratio of Fund 1 is simulated in cell B29 by the formula

```
=NORMINV(RAND( ),B23,B25)
```

where cells B23 and B25 contain the mean and standard deviation of our Fund 1 data in B2:B11. So the formula in cell B29 is equivalent to

```
=NORMINV(RAND( ),AVERAGE(B2:B11),STDEV(B2:B11))
```

Then the growth ratio of Fund 2 is simulated in cell C29 by the formula

```
=NORMINV(RAND( ),C23+ C21*B29,C25)
```

using the intercept C23, the coefficient C21, and standard error C25 from the REGRESSN array formula. A simulation table containing 201 independent recalculations of cells B29:C29 has been made in this spreadsheet and the results are shown in the scatter chart in the lower right. The optimal

Myerson_C06.qxd 07-09-2004 09:43 Page 227

linear function for estimating Fund 2 from Fund 1 is also shown by a dashed line in that chart.

The scatterplot generated by our regression model in Figure 6.8 here should look very similar to the right-hand scatterplot in Figure 5.6, which shows 200 simulations that were generated from a CORAND model that we fit to the same historical data. There is actually a very small difference between the Multivariate Normal probability distributions that generated these two scatterplots. The difference is caused by the practice of dividing by the number of degrees of freedom to estimate the standard error in a regression analysis, instead of dividing by the sample size minus 1 as we do in calculating sample standard deviations. (The problem is that the regression analysis formulas have been designed to yield unbiased estimates of the conditional variance of the dependent variable given the explanatory variables, whereas the sample standard deviation formula has been designed to yield unbiased estimates of the unconditional variance.) The result is that Fund 2's growth ratios have a slightly higher standard deviation in the regression model here than in the CORAND model of Chapter 5. But this difference is small, and it vanishes when the number of data points is large.

## 6.9 Summary

This chapter considered a variety of spreadsheet models in which the value of one random variable may depend on the value of other random variables. We use such models because it is often more convenient to characterize a random variable in terms of its conditional probability distribution given the values of these other random variables (rather than in terms of its marginal probability distribution). In particular, we emphasized models where a random variable **Y** is drawn from a distribution whose mean depends on the value of another random variable **X**. In this case, this mean is the conditional expected value of **Y** given **X**.

The expected-posterior law says that before we observe a random variable **X**, our expected value of another random variable **Y** must be equal to our expected value of what would be the posterior conditional expected value of **Y** after observing **X**. The formula for the expected-posterior law is

$$E(\mathbf{Y}) = \sum_x P(\mathbf{X} = x) * E(\mathbf{Y} | \mathbf{X} = x)$$

This expected-posterior law was illustrated in a more sophisticated version of our old skill-and-sales example, using a Beta distribution for the unknown skill. The logic of the expected-posterior law was then examined more closely in a discrete example where **X** and **Y** each have just three possible values. We also showed how the expected-posterior law allows

Myerson_C06.qxd 07-09-2004 09:43 Page 228

decision analysts to compute expected values in tree diagrams by a kind of backwards analysis.

An example was considered where the conditional expected value of **Y** depends substantially on the value of **X**, but the correlation of **Y** and **X** is 0. Such relationships obviously cannot be detected if we look only at correlations and cannot be modeled with CORANDs, but require the conditional expectation methods of this chapter.

The expected-posterior law also implies that before we observe a random variable **X**, the probability of an event $A$ must be equal to our expected value of what will be the posterior conditional probability of $A$ after the observation of **X**. That is,

$$P(A) = E(P(A|\mathbf{X})) = \sum_x P(\mathbf{X} = x) * P(A|\mathbf{X} = x)$$

With this formula, we considered the problem of assessing a decision maker's subjective probability of an event when uncertain about what better-informed experts might think about the probability of this event. Uncertainty about the experts' probability can be ignored in our decision analysis when there is no possibility of any communication with such experts. But we also discussed two cases where this uncertain expert opinion should be taken into account in our analysis: in situations where the decision maker can buy expert advice (so that we need to compute the expected value of getting such information) and in situations where the decision maker's options include opportunities that are available only because better-informed competitors have not already taken them (so that we need to compute the winner's curse effect).

In a linear regression model, a random variable **Y** can be characterized as a linear function of some other explanatory random variables plus an error term that is an independent Normal random variable. In this model, the conditional probability distribution of **Y** given the explanatory variables is a Normal distribution with a conditional mean that is a linear function of the explanatory variables and a conditional standard deviation (or standard error) that is a constant. Statistical regression software is designed to estimate the parameters of this relationship (the intercept and coefficients of the linear function and the standard error) from observations of **Y** and its explanatory variables. We saw how such regression estimates are computed by finding the linear formula that minimizes the sum of squared errors in the observed data.

In a regression model for **Y**, if the explanatory variables are Multivariate Normals, then the joint distribution of **Y** with these explanatory variables is also Multivariate Normal. So CORAND could be used to make a model that is equivalent to any regression model where the explanatory variables are themselves Multivariate Normals. In such situations, we can use whichever method seems more convenient.

This chapter introduced the Simtools functions BETINV, BINOMINV, and REGRESSN, along with the Excel functions INDEX, MATCH, COUNTIF, and SUMIF.

Myerson_C06.qxd 07-09-2004 09:43 Page 229

## Exercises

**1.** A construction project has two stages, Let **X** denote the cost of the first stage, and let **Y** denote the cost of the second stage (both in thousands of dollars). Suppose that the joint distribution of these costs is as follows:

| | *Y = 200* | *Y = 250* | *Y = 300* |
|---|---|---|---|
| **X** = 90 | 0 | 0.15 | 0.15 |
| **X** = 70 | 0.10 | 0.20 | 0.16 |
| **X** = 50 | 0.09 | 0.09 | 0.06 |

**(a)** Compute the expected value of **Y**.

**(b)** Compute the conditional expected value of **Y** given each of the possible values of **X**.

**(c)** Show that the formula $E(\mathbf{Y}) = \sum_m P(\mathbf{X} = m) * E(\mathbf{Y} | \mathbf{X} = m)$ is satisfied for this example.

**(d)** If you learned that $\mathbf{X} > 60$ (that is, **X** may equal 70 or 90, but not 50), then what would be your conditional expected value of **Y** given this new information? (*Hint*: if you have difficulty with (d), try first to compute $P(\mathbf{Y} = 200 | \mathbf{X} > 60)$, $P(\mathbf{Y} = 250 | \mathbf{X} > 60)$, and $P(\mathbf{Y} = 300 | \mathbf{X} > 60)$.)

**2.** **A** and **B** are unknown quantities. Given my current information only, I think that $E(\mathbf{A}) = 3000$. But if I knew the value of **B**, I would revise my beliefs about **A** to $E(\mathbf{A} | \mathbf{B} = 0) = 1000$ or $E(\mathbf{A} | \mathbf{B} = 1) = 9000$. If **B** can only equal 0 or 1, what is $P(\mathbf{B} = 0)$?

**3.** Dick Gatling has offered John Morgan an opportunity to finance Gatling's new high-technology company. But Morgan recognizes that other specialist bankers may have much more expertise in this technology, and he does not know what they would think about Gatling's chances of making this company succeed. From Morgan's perspective, the probability of Gatling's success that would be assessed by a specialist banker is an unknown quantity. Morgan first thought that if he could assume Gatling has not already approached any other banker, then this unknown quantity should have a Beta distribution with mean 0.45 and standard deviation 0.25. But Morgan had to recognize the possibility that Gatling might have already approached such a specialist banker. If he did, then the specialist banker must have thought the probability of Gatling's success to be below 0.4, because otherwise he would have already financed the company and so Gatling would not have approached Morgan. Morgan figured that a man in Gatling's position would be equally likely to bring such an investment opportunity to Morgan first, or to a specialist banker first.

**(a)** What would be the conditional probability of Gatling's success, if we were given that Gatling had already approached a specialist banker who then refused to finance Gatling's company?

**(b)** Given only that Gatling has approached Morgan for financing, estimate the probability that Gatling previously approached a specialist banker and failed to get financing.

**(c)** In this situation, knowing only that Gatling has approached Morgan for financing, what should Morgan think is the probability of Gatling's success?

**4.** Esther Bahn, a senior partner at Leverone Consulting, has just given Sam Mermes authorization to form a new team that will offer consulting services to small, privately held manufacturing firms whose owners want to make them more attractive targets for acquisition by larger companies. The team will work with each client over a period of months to reorganize finanicial accounting systems and help the client identify potential acquirers. The team will operate with considerable autonomy, and Mermes will have full independent responsibility for getting the word out about his new team and finding clients who will pay for his consulting services. But after two years, if Mermes' team has not generated enough business, then Bahn will shut down the team.

To be profitable, Mermes' team will need to have a long-run average of at least 6 new clients per year. In the first year, the new team will surely get only a couple of clients. But in the second year and every year thereafter, the number of new clients for Mermes' team will be a random variable drawn from a Poisson random variable with some fixed mean. The mean of this Poisson distribution can be called the *long-run potential rate* of Mermes' new consulting team.

The problem is that Bahn is not sure what this long-run potential rate will be. To describe her current beliefs about this number, we can say that Bahn thinks of the team's long-run potential rate as a Gamma random variable that has mean 6 and standard deviation 4.

At the end of the second year, Bahn will look at the number of new clients that Mermes' team got during that year. If this number of clients is less than some threshold value *K*, then Bahn will dissolve the team. Bahn has not yet decided what this threshold number *K* should be.

Myerson_C06.qxd 07-09-2004 09:43 Page 230

**(a)** Build a simulation model of the relationship between the new team's long-run potential rate and its number of clients in its second through fifth years. Make a table of at least 1000 simulations of this model. Using this simulation table, answer the following questions.

**(b)** For each number $K$ from 0 to 10, estimate the probability that the new team will get $K$ new clients in the second year, and estimate the conditionally expected value of the new team's long-run potential rate given that it gets $K$ new clients in the second year.

**(c)** What is the probability that the new team will get more than 10 new clients in the second year? In the event that the team gets more than 10 new clients in the second year, what would be the conditionally expected value of the team's long-run potential rate?

**(d)** Suppose that (after the first year) Bahn will dissolve Mermes' team at the end of any year in which it gets 5 or fewer new clients. Estimate the probability that Mermes' team will survive past the end of its fifth year. Conditional on surviving past the end of the fifth year, what would be the conditionally expected value of the team's long-run potential rate?

Myerson_C07.qxd 09-09-2004 10:35 Page 231

# 7

# Optimization of Decision Variables

Decision variables are the quantities that the decision maker controls. The goal of any decision analysis is to recommend optimal values for these decision variables. This chapter is concerned with fundamental techniques for using simulations to determine the optimal values of decision variables in decisions under uncertainty.

We focus here on two basic types of examples. First, we consider an example where the main decision variable is a manufacturer's production quantity of a commodity for which the future demand is uncertain. (Such examples have been called "newsboy problems" in other texts.) In the context of this example, we study the strategic value of information. Then we consider bidding problems, where the main decision variable is a price to offer in a competitive auction or bargaining situation. In the context of these bidding examples, we study the winner's curse effect and develop models to take it properly into account.

## 7.1 General Techniques for Using Simulation in Decision Analysis

A *decision variable* is any quantity that we have to choose in a decision problem. In the Scotia case, the principal decision variable is the quantity of snowboards that Scotia will produce for next winter. In a decision under uncertainty, we are interested in maximizing some *payoff* value, such as profit, that depends on our decision variables and on various unknown quantities. A rule for choosing all decision variables may be called a *strategy*. In the Scotia case, the payoff to be maximized is Scotia's profit from its snowboard operations, and this profit depends both on the production quantity and on the unknown demand for Scotia's snowboards. The demand

Myerson_C07.qxd 09-09-2004 10:35 Page 232

## CASE ■ Scotia Snowboards

The Scotia Corporation will sell snowboards next winter. Scotia's cost of manufacturing is $20 per snowboard. Scotia will get $48 in revenue for every snowboard that it sells next winter. If demand is less than supply, then the snowboards that Scotia still has unsold in its inventory at the end of the winter will have a value of $8 each to Scotia. If demand is greater than Scotia's supply, then excess demand will be lost to other competitors.

Scotia's factory and shipping schedules require that all snowboards that it sells next winter must be manufactured by September of this year, three months before the winter season begins. As this year's production begins, Scotia does not have any old snowboards in inventory.

The predictions about demand for snowboards next winter depend on the general weather patterns, which may be *normal* or *cold*. If next winter's weather is normal, then demand for Scotia's snowboards will have probability 1/4 of being below 60,000, probability 1/2 of being below 75,000, and probability 3/4 of being below 90,000. If next winter's weather is cold, then demand for Scotia's snowboards have probability 1/4 of being below 80,000, probability 1/2 of being below 100,000, and probability 3/4 of being below 125,000. (In either case, we may assume a Generalized Lognormal probability distribution.) It is currently estimated that the probability of cold weather next winter is 1/3, and the probability of normal weather is 2/3.

The staff engineer assigned to analyze this production planning decision has generally assumed that Scotia's objective in such situations should be to maximize the expected value of its profits. To verify this assumption, he asked Scotia's chief executive officer to think about a simpler gamble where Scotia could earn either $2 million or $1 million, each with probability 1/2, and asked whether Scotia would exchange this gamble for anything less than $1.5 million. The CEO replied, "Well, not for anything less than $1.4 million."

In a production planning session, a marketing manager remarked that good forecasts of the coming winter's general weather would be available in November. But the manufacturing director replied that a delay of snowboard production until November could substantially increase Scotia's total productions costs, perhaps by $100,000 or more.

for Scotia's snowboards next winter is the major unknown quantity or random variable in this case.

People sometimes confuse decision variables with random variables, because in some intuitive sense they are both "unknown" when we start a decision analysis. But of course the decision maker controls a decision variable and so can stop being uncertain about it any time by making the decision. Random variables are used in decision analysis models only to represent quantities that the decision maker does not know and cannot control.

To analyze a decision problem like the Scotia case, we begin by constructing a model that describes how the payoff variables depend on the decision variables that we control and on the random variables that we do not control. In this model, we can try different values for the decision variables, and we can simulate the random variables to see how the probability distribution of payoff outcomes would depend on our decision variables. Then we can try to find values of the decision variables that yield the best possible outcome distribution, according to some optimality criterion such

Myerson_C07.qxd 09-09-2004 10:35 Page 233

as expected-profit maximization or certainty-equivalent maximization with some appropriate risk tolerance.

There are three general methods for doing simulation analyses that should be distinguished.

**Method 1: For each alternative strategy, make a separate simulation table of payoffs.** We can set up our simulation model assuming one strategy for the decision maker and generate a large table of simulated profits (or other payoffs, suitably measured) for this strategy. Then we can make revised versions of the simulation model that assume other feasible strategies, one at a time, and similarly generate a separate table of simulated payoffs for each strategy that we want to consider. The difficulty with this first method is that our alternative strategies are being evaluated with different simulated values of the random variables. So a strategy that is always worse than some other strategies might appear better in such analysis if the random variables that were used to generate its simulated payoffs happened to be more favorable (that is, included more high-demand outcomes and fewer low-demand outcomes than the random data that were used to evaluate the other strategies). Using larger data sets can reduce the probability of this kind of error, but it may be better to avoid it altogether by using the same simulated random variables to evaluate all of our alternative strategies.

**Method 2: Simultaneously simulate payoffs from several alternative strategies.** We can compute payoffs under several alternative strategies simultaneously in our simulation model using the same values of the simulated random variables in the computations for all alternatives. The payoffs from these alternative strategies can be listed in the model-output row of the simulation table. Thus we can generate one simulation table that lists a payoff series for each strategy in a separate column, where all these payoff series were generated with the same series of simulated random variables. The difficulty with this second method is that it requires us to specify in advance the strategies that we want to evaluate, and the simulation table can become big and unwieldy if we want to evaluate many strategies.

**Method 3: Make a table listing simulated values of all payoff-relevant random variables.** If the number of random variables that influence payoffs under all strategies is not too large, then we can make a table of simulated values of these random variables. With this fixed simulation table of all payoff-relevant random variables, we can then compute payoffs for as many different strategies as we like. With this method, we can ask Excel Solver to help us search for strategies that maximize the expected payoff or any other target formula. Even when Solver has difficulties searching, a data table can be used to show how expected payoffs and other summary statistics would depend on different alternative values of any decision variable. This method can become unwieldy and infeasible, however, if the number of payoff-relevant random variables is very large, in which case we should try the second method instead.

Myerson_C07.qxd 09-09-2004 10:35 Page 234

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | SCOTIA PARAMETERS: | | | | FORMULAS | | | |
| 2 | Demand distribution given weather | | | | B15. =IF(RAND()<B8,1,0) | | | |
| 3 | Quartiles | Normal | Cold | | E16. =IF($B$15=1,C4,B4) | | | |
| 4 | q1 (.25) | 60000 | 80000 | | E16 copied to E16:E18 | | | |
| 5 | q2 (.50) | 75000 | 100000 | | B16. =GENLINV(RAND(),E16,E17,E18) | | | |
| 6 | q2 (.75) | 90000 | 125000 | | B18. =MIN(B16,C17) | | | |
| 7 | | | | | B19. =C17-B18 | | | |
| 8 | P(Cold winter) | 0.33333 | | | B20. =B11*B18+B12*B19-B10*C17 | | | |
| 9 | | | | | B28. =B20 | | | |
| 10 | $Cost/unit | 20 | | | B23. =AVERAGE(B29:B429) | | | |
| 11 | $SellingPrice | 48 | | | B24. =STDEV(B29:B429) | | | |
| 12 | RemainderValue | 8 | | | B25. =PERCENTILE(B29:B429,0.1) | | | |
| 13 | | | | | B26. =PERCENTILE(B29:B429,0.9) | | | |
| 14 | SIMULATION MODEL | | | | | | | |
| 15 | Weather cold? | 0 | Demand quartiles given weather | | | | | |
| 16 | Demand | 71698 | | | 60000 | | | |
| 17 | Production Quantity | | 85000 | | 75000 | | | |
| 18 | Sales | 71698 | | | 90000 | | | |
| 19 | Remainder | 13302 | | | | | | |
| 20 | Profit ($) | 1847914 | | ASSESSING RISK TOLERANCE | | | | |
| 21 | | | | RiskTol | from | Hi $ | Low $ | CE |
| 22 | PROFIT STATISTICS | | | 1216303 | | 2000000 | 1000000 | 1400000 |
| 23 | E(Profit) | 1935006 | | | | | | |
| 24 | Stdev | 606928 | | 1746769 | Certainty equivalent | | | |
| 25 | 10%-level | 1096978 | | | | | | |
| 26 | 90%-level | 2380000 | | | FORMULAS FROM RANGE D20:H24 | | | |
| 27 | Simulated Profit | | | | D22. =RISKTOL(F22,G22,H22) | | | |
| 28 | SimTable | 1847914 | | | D24. =CE(B29:B429,D22) | | | |
| 29 | 0 | -319970 | Data in B29:B429 is sorted for chart. | | | | | |
| 30 | 0.0025 | -312721 | Chart plots (A29:A429,B29:B429). | | | | | |
| 31 | 0.005 | -146251 | | | | | | |
| 32 | 0.0075 | -36333 | | | | | | |
| 33 | 0.01 | 9491.48 | | | | | | |
| 34 | 0.0125 | 171755 | | | | | | |
| 35 | 0.015 | 188559 | | | | | | |
| 36 | 0.0175 | 257077 | | | | | | |
| 37 | 0.02 | 270572 | | | | | | |
| 38 | 0.0225 | 275419 | | | | | | |
| 39 | 0.025 | 278428 | | | | | | |
| 40 | 0.0275 | 373929 | | | | | | |
| 41 | 0.03 | 471828 | | | | | | |
| 42 | 0.0325 | 489793 | | | | | | |
| 43 | 0.035 | 498560 | | | | | | |
| 44 | 0.0375 | 525783 | | | | | | |
| 45 | 0.04 | 544715 | | | | | | |
| 46 | 0.0425 | 618400 | | | | | | |

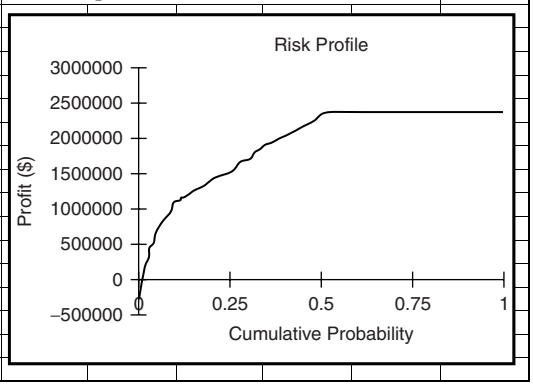


FIGURE **7.1**
**Simulation analysis of one production quantity in the Scotia case**

Let us illustrate these three methods in the context of the Scotia Snowboards case. To begin, consider Figure 7.1, which applies the first method described, that of generating a separate simulation table of payoffs for each possible decision option.

The basic parameters of the Scotia case (demand quartiles in each weather state, probability of cold weather, prices and costs) are summarized in the range A1:C12 of Figure 7.1. The random weather pattern (with 0 for "normal" and 1 for "cold") is simulated in cell B15. Cells E16:E19 contain IF formulas that return the quartile boundaries for the conditional distribution

Myerson_C07.qxd 09-09-2004 10:35 Page 235

of demand given the simulated weather pattern in B15. Then demand for Scotia snowboards is simulated in cell B16 with a Generalized Lognormal random variable that has the quartiles listed in E16:E19. So the probability distribution of simulated demand in cell B16 depends on the simulated weather pattern in cell B15 just as Scotia's demand is supposed to depend on the actual weather pattern.

Cell C17 in Figure 7.1 is used to represent the quantity of snowboards that Scotia will produce, which is the principal decision variable in this model. So the current value of cell C17 in Figure 7.1, which is 85,000, should be interpreted as just one tentative proposal to be evaluated in the model. Scotia's actual sales of snowboards will be either the demand quantity or the production quantity, whichever is smaller. So cell B18, with the formula =MIN(B16,C17) contains the sales quantity (in snowboards) that would result from the production quantity in C17 and the demand in B16. Then the remaining inventory of snowboards at the end of the season (which have a remainder value of $8 per snowboard) is computed in cell B19 by the formula =C17–B18, and the final profit (in dollars) is computed in cell B20 by the formula

=B11*B18+B12*B19–B10*C17

(Cell B10 contains the $20 cost per snowboard, B11 contains the $48 selling price, and B12 contains the $8 value of snowboards in inventory at the end of the winter.)

Cell B28 in Figure 7.1 echoes the simulated profit in cell B20 (by the formula =B20), and 401 simulated values of this simulated profit have been listed below in cells B29:B429. This profit data has been sorted to make the cumulative risk profile shown in the Figure 7.1, which plots the percentile index in cells A29:A429 on the horizontal axis and the simulated profit values in cells B29:B429 on the vertical axis. (Because of this sorting, Figure 7.1 shows only the worst profit outcomes at the top of the SimTable.) The expected profit for Scotia with the production quantity of 85,000 snowboards is estimated to be $1,935,006 in cell B23 by the formula =AVERAGE(B29:B429).

The case suggests that the decision makers at Scotia do not feel any strong discomfort with the criterion of maximizing expected profits, and so it seems reasonable to do most of our decision analysis here under the assumption that Scotia is risk-neutral, which corresponds to a risk tolerance of $+\infty$. But assessing a decision maker's risk tolerance is always difficult. So we should probe a bit further to look for evidence that some lower level of risk tolerance might be appropriate here, and if we find such evidence, then we should see how such lower risk tolerance might affect our conclusions. This process, of assessing a reasonable range for a parameter that is hard to measure exactly, and then seeing how much our conclusions could change if this parameter took a different value within this range, is called *sensitivity analysis*.

In response to such probing, Scotia's CEO has indicated here that if there is any risk aversion among the company's decision makers, the company's risk tolerance should be considered very high, because a gamble paying

Myerson_C07.qxd 09-09-2004 10:35 Page 236

$2,000,000 or $1,000,000, each with probability 1/2, would be worth at least $1,400,000 for sure to the company. This statement indicates that Scotia's risk tolerance should not be less than

RISKTOL(2000000, 1000000, 1400000) = $1,216,303

as computed in cell D22 of Figure 7.1. Such risk tolerance would decrease Scotia's certainty equivalent for the profit gamble in Figure 7.1 to

*CE*(B29:B429,D22) = $1,746,769

as shown in cell D24. This calculation gives us a lower bound on how much risk aversion might decrease the value of this gamble to the decision makers at Scotia.

Figure 7.1 only shows us an analysis of one possible decision: that of producing 85,000 snowboards. But it is easy for us to enter other possible values of this decision variable into cell C17 and regenerate the simulation table. By entering the production quantities 80,000 through 85,000 successively into cell C17 and regenerating the simulation table in this spreadsheet once for each production quantity, we get the following estimates of expected profit in cell B23:

| *Production Quantity* | *Average Simulated Profit* |
|---|---|
| 80,000 | $1,930,656 |
| 81,000 | $1,891,684 |
| 82,000 | $1,909,517 |
| 83,000 | $1,851,352 |
| 84,000 | $1,953,704 |
| 85,000 | $1,935,006 |

As we move down this table, we find that small increases in the production quantity are estimated to have rather large changes in expected profit, which alternates strangely between increases and decreases. Most striking is the result of increasing production from 83,000 to 84,000, which is estimated here to increase expected profit by more than $100,000, even though we cannot earn more than $28 per snowboard! This absurd conclusion is a result of our using different simulated demands to evaluate different production quantities. In the simulations, large demands may have occurred more frequently in the simulations when evaluating the production quantity 84,000 in cell C17, and less frequently in the simulations when evaluating the production quantity 83,000 in cell C17.

To say which production quantity is better under the expected value criterion, we need to estimate how expected profit changes when the production quantity changes. That is, what we really need to estimate accurately are the *differences* between expected profits under different production quantities. These differences can be estimated more accurately when the different production quantities are evaluated with the same series of simulated demands, as in Figure 7.2.

| | A | B | C | D | E | F | G | H | I | J | K | L |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | SCOTIA SNOWBOARDS PARAMETERS: | | | FORMULAS | | | | | | | | |
| 2 | Demand distribution \| weather | | | B16. =IF(RAND()<B8,1,0) | | | | | | | | |
| 3 | Quartiles | Normal | Cold | C16. =IF(B16=1,GENLINV(RAND(),C4,C5,C6),GENLINV(RAND(),B4,B5,B6)) | | | | | | | | |
| 4 | q1 (.25) | 60000 | 80000 | B29. =($B$12-$B$10)*B20+($B$11-$B$12)*MIN(B20,$C$16) | | | | | | | | |
| 5 | q2 (.50) | 75000 | 100000 | B29 copied to B29:E29 | | | | | | | | |
| 6 | q3 (.75) | 90000 | 125000 | B22. =AVERAGE(B30:B430) | | | | | | | | |
| 7 | | | | B23. =STDEV(B30:B430) | | | | | | | | |
| 8 | P(cold winter) | 0.33333 | | B24. =PERCENTILE(B30:B430,0.1) | | | | | | | | |
| 9 | | | | B25. =PERCENTILE(B30:B430,0.9) | | | | | | | | |
| 10 | $Cost/unit | 20 | | B22:B25 copied to B22:E26 | | | | | | | | |
| 11 | $SellingPrice | 48 | | B27. =RISKTOL(2000000,1000000,1400000) | | | | | | | | |
| 12 | RemainderValue | 8 | | B26. =CE(B30:B430,$B$27) | | | | | | | | |
| 13 | | | | B26 copied to B26:E26 | | | | | | | | |
| 14 | | SIMULATION MODEL | | | | | | | | | | |
| 15 | | Cold? | Demand | | | | | | | | | |
| 16 | | 0 | 73334 | | | | | | | | | |
| 17 | | | | | | | | | | | | |
| 18 | ANALYSIS OF FOUR ALTERNATIVE DECISIONS | | | | | | | | | | | |
| 19 | | With Production Quantity = | | | | | | | | | | |
| 20 | | 85000 | 95000 | 105000 | 115000 | | | | | | | |
| 21 | Profit Statistics: | | | | | | | | | | | |
| 22 | E(Profit) | 1937071 | 1961756 | 1927204 | 1865501 | | | | | | | |
| 23 | Stdev | 587543 | 710961 | 805984 | 886796 | | | | | | | |
| 24 | 10%-level | 1006952 | 886952 | 766952 | 646952 | | | | | | | |
| 25 | 90%-level | 2380000 | 2660000 | 2940000 | 3220000 | | | | | | | |
| 26 | CertainEquiv | 1766507 | 1724523 | 1639151 | 1536509 | | | | | | | |
| 27 | with RiskTol | 1216303 | | | | | | | | | | |
| 28 | Sim'd profits with ProdnQ as above | | | | | | | | | | | |
| 29 | SimTable | 1913375 | 1793375 | 1673375 | 1553375 | | | | | | | |
| 30 | 0 | 2380000 | 2660000 | 2645188 | 2525188 | | | | | | | |
| 31 | 0.0025 | 2338534 | 2218534 | 2098534 | 1978534 | | | | | | | |
| 32 | 0.005 | 2220079 | 2100079 | 1980079 | 1860079 | | | | | | | |

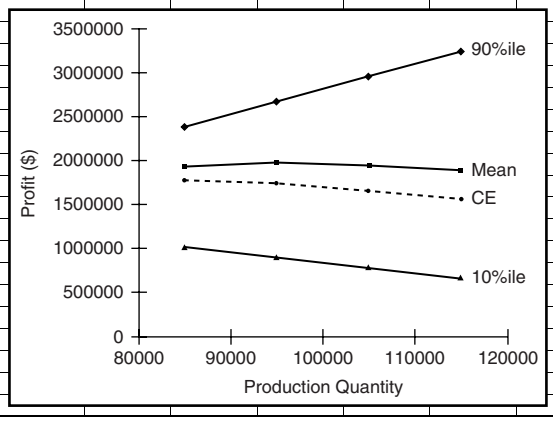


FIGURE **7.2** **Simulation analysis of four alternative production quantities in the Scotia case**

Our second method, of simultaneously simulating payoffs from selected alternatives with the same simulated demands, is applied in Figure 7.2 for the selected production quantities of 85,000, 95,000, 105,000, and 115,000, which have been entered into cells B20:E20. The simulated demands to compare these alternative decisions are generated by the random variable in cells B16:C16. Cell B16 simulates the weather pattern by the formula

=IF(RAND( )<B8,1,0)

where B8 contains the value 1/3 = *P*(Cold). The simulated demand in cell C16 is generated by the formula

=IF(B16=1,GENLINV(RAND( ),C4,C5,C6),GENLINV(RAND( ),B4,B5,B6)

where cells C4:C6 contain the quartile boundaries for demand in cold weather and cells B4:B6 contain the quartile boundaries for demand in normal weather.

The profit that would result from the demand in cell C16 and the production quantity in cell B20 is computed in cell B29 of Figure 7.2 by the formula

=($B$12−$B$10)∗B20+($B$11−$B$12)∗MIN(B20,$C$16)

Myerson_C07.qxd 09-09-2004 10:35 Page 238

To interpret this formula, notice that B12 contains the $8 remainder value and B10 contains the $20 cost per unit, so (B12–B10)*B20 is the (negative) profit that would result from producing the quantity B20 and leaving all production in our final inventory. But then taking each unit out of inventory and selling it at the $48 price in cell B11 increases profit by B11 – B12, and the number of units sold is the minimum of production and demand MIN(B20,C16). Copying this formula from B29 to B29:E29 gives us the alternative profits that would result from the alternative production quantities in B20:E20, all with the same simulated demand in C16.

The range B30:E430 in Figure 7.2 contains a table of profit outcomes for these four production quantities with 401 independent simulations of the demand random variable B16. The sample averages of these simulated profits are computed in cells B22:E22 to estimate the expected profit for each production quantity. By the criterion of expected profit maximization, the best of these four production quantities is apparently 95,000, which yields an estimated expected profit of $1,961,756. These estimated expected values (or means) of profit are plotted above their corresponding production quantities by the thick curve in the chart in Figure 7.2.

The top and bottom curves in the chart in Figure 7.2 show the levels of profit that have cumulative probabilities 0.90 and 0.10 in the profit distributions that result from these four alternative production quantities. These 0.9 and 0.1 cumulative probability levels are computed by entering the formulas

=PERCENTILE(B30:B430,0.1) and =PERCENTILE(B30:B430,0.9)

into cells B24 and B25, and then copying B24:B25 to B24:E25. The increasing spread between these 0.10 percentile and 0.90 percentile values shows that higher production quantities may be associated with more risk.

With the lower bound on Scotia's risk tolerance ($1,216,303) entered into cell B27 of Figure 7.2, the corresponding certainty equivalents of profit for the four selected production quantities are estimated in cells B26:E26 by entering the formula

=CE(B30:B430,$B$27)

into cell B26, and then copying cell B26 to cells B26:E26. The certainty equivalent curve in the chart, which plots these certainty equivalent values from B26:E26, is always below the expected monetary value (mean) curve, but the certainty equivalent is farther below the expected monetary value for the higher production quantities. The greatest of these four certainty equivalents is in cell B29, generated by the production quantity 85,000. Thus, a decision analysis that assumes risk aversion may lead us to recommend a lower optimal production quantity than a decision analysis that assumes risk neutrality.

Figure 7.3 applies the third method to the Scotia case, in which we first make a table of simulated values of the random variables, and then we compute profits from this fixed simulation data by formulas that depend actively on a cell that represents our decision variable. The random variables

| | A | B | C | D | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | SCOTIA SNOWBOARDS PARAMETERS: | | | FORMULAS FROM RANGE A1:E1023 | | | | | | | |
| 2 | Demand distribution | weather | | B22. =IF(RAND()<B8,1,0) | | | | | | | |
| 3 | Quartiles | Normal | Cold | C22. =IF(B22=1,GENLINV(RAND(),C4,C5,C6),GENLINV(RAND(),B4,B5,B6)) | | | | | | | |
| 4 | q1 (.25) | 60000 | 80000 | C18. =AVERAGE(C$23:C$1023) | | | | | | | |
| 5 | q2 (.50) | 75000 | 100000 | C19. =STDEV(C$23:C$1023) | | | | | | | |
| 6 | q3 (.75) | 90000 | 125000 | C18:C19 copied to E18:E19 | | | | | | | |
| 7 | | | | E23. =($B$12-$B$10)*$D$18+($B$11-$B$12)*MIN($D$18,C23) | | | | | | | |
| 8 | P(Cold) | 0.33333 | | E23 copied to E23:E1023 | | | | | | | |
| 9 | | | | B15. =RISKTOL(2000000,1000000,1400000) | | | | | | | |
| 10 | $Cost/unit | 20 | | E20. =CE(E23:E1023,$B$15) | | | | | | | |
| 11 | $SellingPrice | 48 | | SOLVER: maximize E18 (or E20) by changing D18. | | | | | | | |
| 12 | RemainderVal | 8 | | | | | | | | | |
| 13 | | | | | | | | | | | |
| 14 | For CEs: | | | | | | | | | | |
| 15 | RiskTolerance | 1216303 | | DECISION ANALYSIS | | | ProdnQ | E(Proft) | 10%ile | 90%ile | CE |
| 16 | | | | Without forecast | | | (,D18) | 1968188 | 943976 | 2689017 | 1716953 |
| 17 | | | Demand | ProdnQ | ProfitStatistics | | 70000 | 1775992 | 1256412 | 1960000 | 1698493 |
| 18 | | Mean | 84574 | 96036 | 1968188 | E(Proft) | 80000 | 1901571 | 1136412 | 2240000 | 1761550 |
| 19 | | Stdev | 28556 | | 725049 | Stdev | 90000 | 1959554 | 1016412 | 2520000 | 1749579 |
| 20 | SIMULATION MODEL | | | | 1716953 | CE | 100000 | 1964204 | 896412 | 2800000 | 1687932 |
| 21 | | Cold? | Demand | | | | 110000 | 1924151 | 776412 | 3080000 | 1595937 |
| 22 | SimTable | 1 | 86010 | | Profit | | 120000 | 1854014 | 656412 | 3332364 | 1488696 |
| 23 | 0.000 | 0 | 49923 | | 844490 | | | | | | |
| 24 | 0.001 | 0 | 94891 | | 2643193 | | FORMULAS FROM RANGE G15:K22 | | | | |
| 25 | 0.002 | 1 | 60225 | | 1256550 | | H16. =E18 | | K16. =E20 | | |
| 26 | 0.003 | 1 | 146634 | | 2689017 | | I16. =PERCENTILE(E23:E1023,0.1) | | | | |
| 27 | 0.004 | 0 | 42152 | | 533661 | | J16. =PERCENTILE(E23:E1023,0.9) | | | | |
| 28 | 0.005 | 1 | 156649 | | 2689017 | | H17:K22. {=TABLE(,D18)} | | | | |

FIGURE **7.3** **A spreadsheet for finding the optimum among all possible production quantities for Scotia**

for weather and demand are simulated in cells B22 and C22 of Figure 7.3 by the formulas

=IF(RAND( ) <B8,1,0)

to simulate the weather pattern (cold = 1, normal = 0) in cell B22, and

=IF(B22=1,GENLINV(RAND( ),C4,C5,C6),GENLINV(RAND( ),B4,B4,B6)

to simulate Scotia's snowboard demand in cell C22. The results of 1001 simulations of these random variables have been stored below in the range B23:C1023. So the expected value of demand can be estimated to be 84,574, according to the average of sampled demands that is computed in cell C18.

A possible production quantity has been entered into cell D18 of Figure 7.3, and the profits that would result from this production quantity for each of the simulated demand values are computed in E23:E1023 by entering the formula

=($B$12–$B$10)∗$D$18+($B$11–$B$12)∗MIN($D$18,C23)

into cell E23, and then copying E23 to cells E23:E1023. (As in cell B29 of Figure 7.2, this formula deducts B10 – B12 = $20 – $8 from profit for each unit when it is produced, and then credits B11 – B12 = $48 – $8 to profit for

Myerson_C07.qxd 09-09-2004 10:35 Page 240

each unit when it is sold.) Then the expected profit with the production quantity D18 is estimated in cell E18 of Figure 7.3 by the formula

=AVERAGE(E23:E1023)

The big advantage of this method of analysis is that this estimated expected profit in cell E18 here depends actively on the decision variable in cell D18. We can try different decisions in D18 and watch how the expected profits change in E18. A data table can be used to show how expected profit and other statistics of the simulated profit distribution would vary with any group of selected production quantities, as shown by the data table in cells H17:K22.

We can even ask Solver to search for the value of our decision variable in cell D18 that would maximize the expected profit in cell E18. The value 96,036 that is shown in cell D18 in Figure 7.3 is the result of Solver's optimization when directed to maximize E18 by changing D18. The global optimality of this result was confirmed by repeated reapplications of Solver with different initial values entered into cell D18, some much higher and others much lower, all of which terminated with Solver reporting this production quantity as the optimal solution. So according to our simulation analysis, Scotia's optimal production quantity is 96,036 snowboards, which yields an expected profit of $1,968,188.

Notice that this optimal production quantity for maximizing the expected value of profit is quite different from the expected value of demand (84,574) that we estimated in cell C18 of Figure 7.3. If demand were perfectly predictable, then the profit-maximizing production quantity would be equal to demand. But when there is uncertainty about demand, the optimal production quantity is not necessarily equal to expected demand, even when our goal is to maximize expected profit. That is, producing to expected demand would not maximize profit here, and believing that it would is an example of the common error known as the "fallacy of averages."

To do sensitivity analysis in Scotia's risk tolerance, we should see how the optimal production quantity would change with the lower bound on Scotia's risk tolerance, which has been entered into cell B15 of Figure 7.3. With this risk tolerance in cell B15, the formula

=CE(E23:E1023,B15)

in cell E20 computes the certainty equivalent of the profit distribution generated in E23:E1023 by the production quantity that we enter in cell D18. Then we can ask Solver to maximize this certainty equivalent in cell E20 by changing the production quantity in cell D18. With the given simulated demand data in this spreadsheet, Solver found the *CE*-maximizing production quantity to be 83,311 snowboards, which yielded a certainty equivalent of $1,765,006. Using this optimal production quantity, the mean and standard deviation of Scotia's profits, as computed in cells C18 and C19 of Figure 7.3, were found to be $1,928,144 and $561,925, respectively. Thus, compared to the risk-neutral optimization shown in Figure 7.3 (where the production quantity is 96,036), we again find that risk aversion makes

Myerson_C07.qxd 09-09-2004 10:35 Page 241

Scotia reduce its production somewhat, which decreases expected profit by about 2% but decreases the standard deviation of profit by about 20%.

## 7.2 Strategic Use of Information

The last paragraph of the Scotia case raises the question of whether Scotia should be willing to pay some additional costs of about $100,000 to be able to base its production decision on better information about the future weather pattern that will prevail next winter. To answer this question, we must estimate how much information about the weather pattern would be worth to Scotia. To compute the expected value of information in such cases, we must first compute how the information might affect our optimal decision. This is a subtle question, because we do not yet know the information, and so we must consider all possibilities of what the information might turn out to be. In this case, we are considering information about a predicted weather pattern that could be either "normal" or "cold." So we must ask what Scotia's optimal production quantity would be if we learned that the forecast was "normal," and we must also ask what Scotia's optimal production quantity would be if we learned that the forecast was "cold." (For simplicity, let us assume here that the $100,000 cost would enable us to get a perfect prediction of the normal or cold weather patterns.) Before we pay for this weather forecast information, however, we only know that the "normal" forecast has probability 2/3 and the "cold" forecast has probability 1/3. So Scotia's expected payoffs from these two possible forecasts must be averaged with these probabilities 2/3 and 1/3. Then the resulting expected value can be compared to the expected value that Scotia could get without the forecast (as computed in the previous section), and the difference will be the expected value of this weather forecast information for Scotia.

Figure 7.4 shows a spreadsheet to calculate this expected value of this information for Scotia, using the criterion of expected profit maximization (risk neutrality). Figure 7.4 was constructed from a copy of Figure 7.3. The simulation model and the simulation data in columns B and C of Figure 7.4 are the same as in Figure 7.3, and the profits in cells E23:E1023 of Figure 7.4 depend on the simulation data and on the production quantity in cell D18 in the same way as in Figure 7.3. With the simulation data copied from Figure 7.3, maximizing the mean profit in cell E18 by changing the production quantity in cell D18 in Figure 7.4 yields the same optimal solution (96,036 snowboards) as in Figure 7.3.

To investigate the question of what would be optimal if Scotia knew that the weather pattern would be cold, the value 1 has been entered into cell G21 of Figure 7.4, the formula

```
=IF(B23=$G$21,E23,"..")
```

| | A | B | C | D | E | F | G | H | I | J | K | L | M |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | SCOTIA PARAMETERS | | | FORMULAS FROM RANGE A1:L1023 | | | | | | | | | |
| 2 | Demand distrib'n | weather | | B22. =IF(RAND()<B8,1,0) | | | | | | | | | |
| 3 | Quartiles | Normal | Cold | C22. =IF(B22=1,GENLINV(RAND(),C4,C5,C6),GENLINV(RAND(),B4,B5,B6)) | | | | | | | | | |
| 4 | q1 (.25) | 60000 | 80000 | E23. =($B$12-$B$10)*$D$18+($B$11-$B$12)*MIN($D$18,C23) | | | | | | | | | |
| 5 | q2 (.50) | 75000 | 100000 | G23. =IF(B23=$G$21,E23,"..") | | | | | | | | | |
| 6 | q3 (.75) | 90000 | 125000 | J23. =VLOOKUP(B23,$I$18:$J$19,2,0) | | | | | | | | | |
| 7 | | | | K23. =($B$12-$B$10)*J23+($B$11-$B$12)*MIN(J23,C23) | | | | | | | | | |
| 8 | P(Cold) | 0.3333 | | E23:K23 copied to E23:K1023. | | | | | | | | | |
| 9 | | | | E18. =AVERAGE(E23:E1023) | | | | E18 copied to B18,C18. | | | | | |
| 10 | $Cost/unit | 20 | | E19. =STDEV(E23:E1023) | | | | | | | | | |
| 11 | SellingPrice | 48 | | E18:E19 copied to G18:G19,K18:K19. | | | | | | | | | |
| 12 | RemainderVal | 8 | | K15. =K18-E18 | | | | | | | | | |
| 13 | | | | SOLVER: maximize K18 by changing J18:J19. | | | | | | | | | |
| 14 | | | | | | | | | | | E(ValueOfInfo) | | |
| 15 | | | | DECISION ANALYSIS | | | | | | | 42333 | | |
| 16 | | From SimTable | | Without forecast | | Conditional | | | Strategy with forecast | | | | |
| 17 | P(Cold) | E(D) | | ProdnQ | | on weather | | | Cold? | ProdnQ | Profit | | |
| 18 | 0.3357 | 84574 | | 96036 | 1968188 | E(Profit) | 2340521 | | 0 | 86255 | 2010521 | E(Profit) | |
| 19 | | | | | 725049 | Stdev | 486394 | | 1 | 114839 | 746333 | Stdev | |
| 20 | SIMULATION MODEL | | | | | with (Cold?)= | | | | | | | |
| 21 | Cold? | Demand | | | | | 1 | | | | | | |
| 22 | SimTable | 1 | 117564 | | Profit | Profit | (Cold?) | | | ProdnQ | Profit | | |
| 23 | 0.000 | 0 | 49923 | | 844490 | | .. | | | 86255 | 961862 | | |
| 24 | 0.001 | 0 | 94891 | | 2643193 | | .. | | | 86255 | 2415151 | | |
| 25 | 0.002 | 1 | 60225 | | 1256550 | | 1256550 | | | 114839 | 1030916 | | |
| 26 | 0.003 | 1 | 146634 | | 2689017 | | 2689017 | | | 114839 | 3215496 | | |
| 27 | 0.004 | 0 | 42152 | | 533661 | | .. | | | 86255 | 651032 | | |
| 28 | 0.005 | 1 | 156649 | | 2689017 | | 2689017 | | | 114839 | 3215496 | | |
| 29 | 0.006 | 1 | 78325 | | 1980555 | | 1980555 | | | 114839 | 1754922 | | |
| 30 | 0.007 | 1 | 122266 | | 2689017 | | 2689017 | | | 114839 | 3215496 | | |
| 31 | 0.008 | 1 | 66303 | | 1499674 | | 1499674 | | | 114839 | 1274041 | | |
| 32 | 0.009 | 0 | 72144 | | 1733342 | | .. | | | 86255 | 1850713 | | |

FIGURE **7.4** **Strategic value of information about weather in the Scotia case**

has been entered into cell G23, and G23 has been copied to cells G23:G1023. Thus, cells G23:G1023 select the profits from the E column in rows where the weather pattern is cold (as denoted by a 1 in the B column), but the profits are omitted and replaced by nonnumerical text ("..") in rows where the weather pattern is normal. Nonnumerical data is ignored by the statistical functions AVERAGE, STDEV, and CE. So the formula

=AVERAGE(G23:G1023)

in cell G18 returns our estimate of Scotia's conditional expected profit given a cold weather pattern when the production quantity is as specified in cell D18. With the production quantity of 96,036 (which maximizes Scotia's expected profit when we do not know the weather pattern), Figure 7.4 shows the resulting conditional expected profit given cold weather would be $2,340,521. But if we asked Solver to maximize this conditional expected profit in cell G18 by changing the production quantity in cell D18, then Solver would return a production quantity of 114,839 in D18, yielding an increased conditional expected profit of $2,421,032 in G18.

Myerson_C07.qxd 09-09-2004 10:35 Page 243

Next suppose that we changed the value of cell G21 to 0. Then cells G23:G1023 would instead select all the profits from the E column in rows where the weather pattern is normal (not cold), and the profits would be omitted in rows where the weather pattern is cold. So with G21 equal to 0, the average reported in cell G18 would become an estimate of the conditional expected profit given normal weather when the production quantity is as in cell D18. If we kept the production in D18 equal to 96,036, then the resulting conditional expected profit given normal weather would be $1,780,062. But if we asked Solver to maximize this conditional expected profit in cell G18 by changing the production quantity in cell D18, then Solver would return a production quantity of 86,255 in D18, yielding a conditional expected profit of $1,803,105 in G18.

The I, J, K columns of Figure 7.4 show how to analyze the optimal production decision given weather information for both possible states of the weather at the same time by using strategic analysis. Decision analysts use the word *strategy* to refer to any plan that specifies what decisions would be made as a function of any information that the decision maker may get now and in the future. So consider the situation when Scotia's managers have decided to anticipate getting a perfect forecast of the winter weather pattern before they choose the production quantity, but they do not yet have the forecast in hand. In this situation, a strategy for Scotia would be any rule that specifies how much to produce if the weather forecast is normal and how much to produce if the forecast is cold.

The range I18:J19 in Figure 7.4 is used to display the strategy being considered. Cells I18 and I19 contain, respectively, the numbers 0 and 1 to represent the cases of normal and cold forecasts. The intended production quantities in these two cases are listed, respectively, in cells J18 and J19. For example, entering the value 90,000 into cell J18 and 100,000 into cell J19 would represent the strategy of planning to order 90,000 snowboards if the forecast is normal but 100,000 snowboards if the forecast is cold.

The formula

```
=VLOOKUP(B23,$I$18:$J$19,2,0)
```

has been entered into cell J23 in Figure 7.4. This formula tells Excel to look for the value of B23 in the leftmost (I) column of the range I18:J19, and then return the value found in column 2 within this range (the J column). Including 0 as the optional fourth parameter of the VLOOKUP function tells it to look for an exact match for the B23 value in the I18:I19 range. Then this formula in cell J23 has been copied to J23:J1023, and so each cell in J23:J1023 displays the production quantity that would be implemented under our strategy for the weather pattern given in the corresponding cell in B23:B1023. The profits that would result from implementing this strategy are computed in the K column by entering the formula

```
=($B$12–$B$10)*J23+($B$11–$B$12)*MIN(J23,C23)
```

into cell K23 and then copying cell K23 to K23:K1023.

Myerson_C07.qxd 09-09-2004 10:35 Page 244

Scotia's expected profit under this I18:J19 strategy is computed in cell K18 by the formula

=AVERAGE(K23:K1023)

Now we can ask Solver to maximize this expected profit in cell K18 by changing the strategic plan's production quantities in cells J18:J19. The result is that Solver will find the optimal strategy for using this forecast information to maximize Scotia's expected profit. The resulting values returned by Solver are shown in cells J18:J19 of Figure 7.4. That is, the optimal strategy found by Solver is to produce 87,226 snowboards if the forecast is normal and to produce 118,443 snowboards if the forecast is cold. In fact, these two conditional production quantities are the same as the quantities that would maximize the conditionally expected profit in G18 given each of these two weather patterns in G21. This optimal strategy yields the expected value $2,010,521, shown in cell K18.

Figure 7.4 also shows in cell D18 the optimal production quantity for maximizing the expected profit in cell E18 when Scotia does not get any weather forecast, which is the production quantity 96,036 that we found in Figure 7.3 with the same simulation data. So the difference between cell E18 and cell K18, which is computed in cell K15, is the amount that Scotia's expected profit would be increased by allowing Scotia to use such a perfect weather forecast in its production planning. This increase in expected profit is called the *expected value of information*. So according to our simulation analysis, the expected value of weather information to Scotia would be

$2,010,521 – $1,968,188 = $42,333

This expected value of information is the most that Scotia should be willing to pay to get a perfect forecast of the winter weather pattern before choosing its production quantity. Thus, the estimated costs of $100,000 to delay production until forecast information is available are greater than the expected benefits of getting the information.

Figure 7.5 summarizes the conditional expected profits that we have computed from this simulation data for Scotia's optimal strategies, both for the case where Scotia can base its production decision on information about the weather pattern and for the case where Scotia must make its production decision without such information. The optimal production values in each case have been copied to cells Q3 and R3:S3, and then the conditional expectations in cells Q4:S5 are computed by a two-way data table. (To make this data table, first enter the formula =G18 into cell P3, so that it echoes the conditional expected profit that is computed in cell G18 of Figure 7.4. Then select the range P3:S5, enter the command sequence Data:Table, and indicate the row input cell D18 and the column input cell G21. The result is that cells Q4:S5 are filled with recalculated values of the data table's top-left corner, cell P3, as the top-row Q3:S3 values are entered into the production quantity cell D18 of Figure 7.4, and the left-column P4:P5 values are entered into the given weather state cell G21 of Figure 7.4.)

| | O | P | Q | R | S | T | U | V | W |
|---|---|---|---|---|---|---|---|---|---|
| 1 | A TWO-WAY DATA TABLE(D18,G21) | | | | | | | | |
| 2 | | E($\|State) | | ProdnQ | | | | | |
| 3 | | 2340521 | 96036 | 86255 | 114839 | | Weather states | | |
| 4 | State | 0 | 1780062 | 1803105 | 1626008 | | 0=normal | | |
| 5 | | 1 | 2340521 | 2227280 | 2421032 | | 1=cold | | |
| 6 | | | | | | | | | |
| 7 | CONDITIONAL OPTIMIZATION RESULTS WITH AND WITHOUT FORECAST | | | | | | | | |
| 8 | | | Deciding without info | | | With weather info | | | |
| 9 | State | Proby | ProdnQ | E($\|State) | | ProdnQ | E($\|State) | | Gain |
| 10 | 0 | 0.664336 | 96036 | 1780062 | | 86255 | 1803105 | | 23044 |
| 11 | 1 | 0.335664 | 96036 | 2340521 | | 114839 | 2421032 | | 80511 |
| 12 | | | | E($) | | | E($) | | |
| 13 | | | | 1968188 | | | 2010521 | | |
| 14 | | | Expected value of weather information | | | | | | |
| 15 | | | | | 42333 | | | | 42333 |
| 16 | | FORMULAS | | | | | | | |
| 17 | | P3. =G18 | | Q4:S5. {=TABLE(D18,G21)} | | | | | |
| 18 | | P11. =B18 | | P10. =1-P11 | | | | | |
| 19 | | Q10. =Q3 | | R10. =Q4 | | | | | |
| 20 | | Q11. =Q3 | | R11. =Q5 | | | | | |
| 21 | | T10. =R3 | | U10. =R4 | | W10. =U10-R10 | | | |
| 22 | | T11. =S3 | | U11. =S5 | | W11. =U11-R11 | | | |
| 23 | | R13. =SUMPRODUCT(R10:R11,$P$10:$P$11) | | | | | | | |
| 24 | | U13. =SUMPRODUCT(U10:U11,$P$10:$P$11) | | | | | | | |
| 25 | | S15. =U13-R13 | | | | | | | |
| 26 | | W15. =SUMPRODUCT(W10:W11,$P$10:$P$11) | | | | | | | |

FIGURE **7.5**
**Conditional expected profits given weather, for Scotia's optimal strategies**

The unconditional expected values for each strategy are computed from these conditional expected values in cells R13 and U13 of Figure 7.5 by an application of the expected-posterior law that we studied in Chapter 6.

The analysis in Figure 7.4 could be easily adapted to use the criterion of certainty equivalent maximization instead of expected profit maximization. To apply the risk tolerance constant $1,216,303, we could simply change the estimated expected profit in K18 by the certainty equivalent formula

=CE(K23:K1023, 1216303)

and ask Solver to maximize this certainty equivalent by changing the production quantities in J18:J19. With the simulation data in this spreadsheet, Solver found an optimal strategy of producing 77,467 snowboards if the forecast is normal and producing 100,104 snowboards if the forecast is cold. The estimated certainty equivalent of the profit distribution with this risk-averse optimal strategy was found to be $1,796,508. From the discussion at the end of Section 7.1, recall that the highest certainty equivalent that could be achieved without forecast information (using the same simulation data) was found to be $1,765,006, achieved with the production quantity 83,311. The value of weather forecast information for Scotia with this constant risk tolerance can then be estimated by the difference between these two *CE* values:

$1,796,508 –$1,765,006 = $31,502

So when risk aversion is taken into account, we still find that the production delay is not worthwhile for Scotia.

Myerson_C07.qxd 09-09-2004 10:35 Page 246

## 7.3 Decision Trees

Decision problems like the Scotia case can be illustrated using tree diagrams, similar to those that we studied in Section 1.6 and Section 6.4. Figure 7.6 illustrates the Scotia case when it has to choose its production quantity without weather information. Starting at the left (from cell A21), each path in this tree describes a possible sequence of events. First the decision maker at Scotia must choose a production quantity. Then, as our model is constructed, the next event to be determined is the winter's general weather pattern (normal or cold), and finally the actual demand for Scotia's snowboards is determined according to a probability distribution that depends

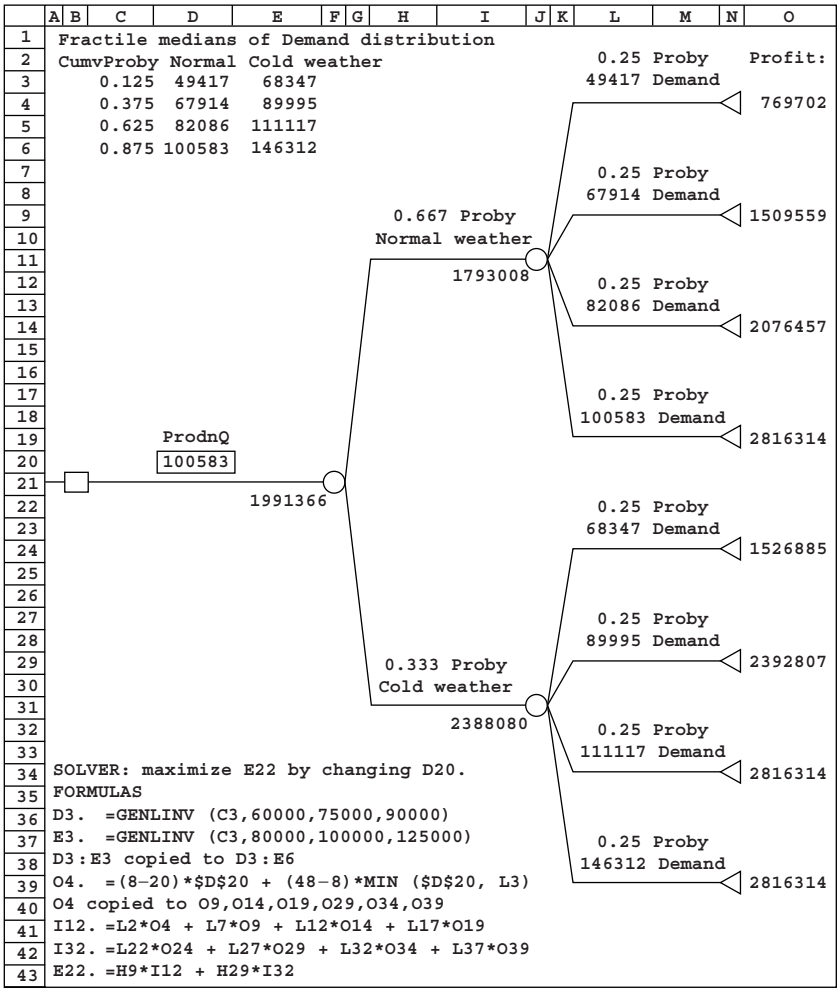


FIGURE **7.6**
**Tree diagram for Scotia's decision without forecast information**

Myerson_C07.qxd 09-09-2004 10:35 Page 247

on the weather pattern. The resulting profit for Scotia from its snowboard operations is listed after each endpoint of the tree.

We cannot draw a tree diagram that shows all possible demand quantities in this problem. Indeed, it is hard to work with a tree diagram that has more than a few dozen endpoints. So decision analysts who work with tree diagrams often approximate continuous random variables by discrete random variables that have a small number of possible values, with the hope that a few carefully chosen values might still represent a fair sample of the wide range of possible outcomes. In Figure 7.6, a technique called *quartile median selection* is applied to generate probability distributions with four equally likely values, as simple discrete approximations of the continuous distributions that were originally assessed for demand in each state of the weather.

For any of the four quartiles, the *quartile median* is the value such that if the random variable were in this quartile, then it would be equally likely to be above or below the quartile median. For a RAND( ) that is uniformly distributed on the interval from 0 to 1, the four quartile medians are 0.125, 0.375, 0.625, and 0.875. That is, if we were told that a RAND( ) value was less than 1/4 then we would say that it was equally likely to be above or below 1/8 = 0.125; but if we were told that the RAND( ) value was between 1/4 and 1/2, then we would say that it was equally likely to be above or below 3/8 = 0.375; and so on. Demand in the normal weather state is simulated by the random variable

GENLINV(RAND( ), 60000, 75000, 90000)

which is an increasing function of its RAND( ). If we were told that this random variable was in its lowest quartile of possible values for the normal weather state, then we would know that its RAND( ) was between 0 and 1/4, which would make the RAND( ) equally likely to be above or below 0.125, and so the demand random variable would be equally likely to be above or below

GENLINV(0.125, 60000, 75000, 90000) = 49,417

which is thus the first quartile median for the demand distribution in the normal weather state. More generally, the four quartile medians for the demand distribution with normal weather can be computed by the formula

GENLINV($x$, 60000, 75000, 90000) with $x$ = 0.125, 0.375, 0.625, or 0.875

which yields values 49,417 and 67,914 and 82,086 and 100,583 in cells D3:D6 of Figure 7.6. Similarly, the four quartile medians for the demand distribution with a cold weather pattern are computed by the formula

GENLINV($x$, 80000, 100000, 125000) with $x$ = 0.125, 0.375, 0.625, or 0.875

Myerson_C07.qxd 09-09-2004 10:35 Page 248

which yields the values 68,347 and 89,995 and 111,117 and 146,312 in cells E3:E6.

The decision tree shown in Figure 7.6 simplifies the demand distribution in each state of the weather by assuming that in each state of the general weather pattern, the demand for Scotia's snowboards will be one of the four quartile medians for the demand distribution in this state, each with probability 1/4. In effect, we have simplified the problem by collapsing all probability in each quartile onto the quartile median.

The conditional expected value of profit at each node in this tree is computed by backwards analysis, as discussed in Section 6.4. That is, at each node in the J column, we multiply the profit after each branch times the probability of the branch, and we sum these products for all the branches that go to the right from the node, and the resulting conditionally expected profit is listed just below and to the left of the node in Figure 7.6. Then, working backwards, we repeat this expected-profit calculation at the node in cell F21, but using the conditionally expected profits from the subsequent J-column nodes.

The tree diagram in Figure 7.6 may be called a *decision tree* because the profits at the endpoints depend on our decision about the production quantity. A standard convention in decision-tree diagrams is to represent alternative decisions by alternative branches at a square *decision node*. But in a spreadsheet where profit depends on our decision by a formula, we do not need to actually draw different branches for different alternatives, but instead we can represent the decision by a number in a cell and make the profits depend on this number by the appropriate formula. In Figure 7.6, the decision about production quantity is represented by the number in cell D20, and the final profits are computed in the O column with reference to this cell (as well as the cell that represents the demand in this outcome). For example, the profit after the top branch (with normal weather and lowest-quartile demand) is computed in cell O4 by the formula

```
=(8–20)*$D$20+(48–8)*MIN($D$20,L3)
```

With such functional dependence of profits on the production quantity in cell D20, we can ask Solver to maximize the expected profit in cell E22 by changing the production quantity in cell D20. The result, as shown in Figure 7.6, is the profit-maximizing plan to produce 100,583 snowboards.

Figure 7.7 shows how this decision tree would change if Scotia delayed the production decision until information about the winter's general weather pattern would be available. The big difference is that getting this weather information before the production decision allows us to have two different production quantities in the tree diagram: one in cell H10 for the case of normal weather, and one in cell H30 for the case of cold weather. The profit formulas in O4:O19 refer to the normal-weather production quantity in cell H10, and the profit formulas in O24:O39 refer to the cold-weather production quantity in cell H30. But it is also assumed here that

Myerson_C07.qxd 09-09-2004 10:35 Page 249

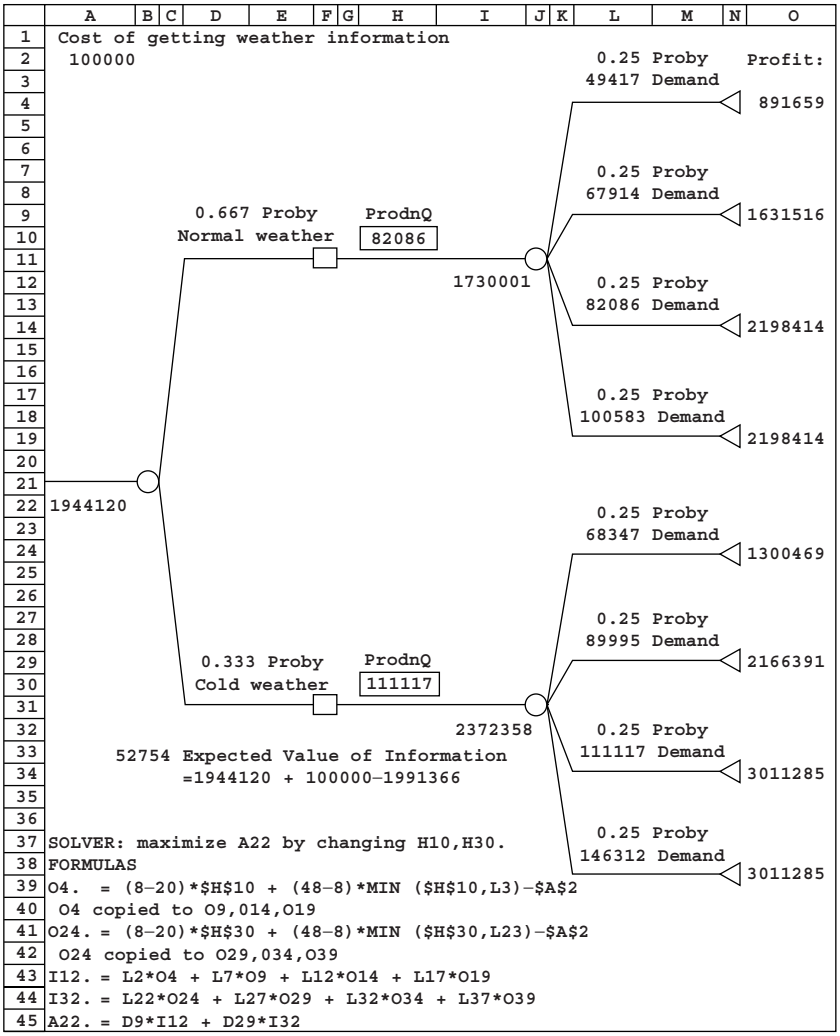


FIGURE **7.7**
**Tree diagram for Scotia's decision with forecast information**

the delay of production to get more information would increase Scotia's production costs by $100,000, as indicated in cell A2, and this quantity is subtracted from the profits at all endpoints of the tree. Thus, for example, the profit after the top branch (with normal weather and lowest-quartile demand) is computed in cell O4 by the formula

=(8–20)∗$H$10+(48–8)∗MIN($D$20,L3)–$A$2

Similarly, cell O24 computes profit with cold weather and lowest-quartile demand by the formula

=(8–20)∗$H$30+(48–8)∗MIN($H$30,L23)–$A$2

Myerson_C07.qxd 09-09-2004 10:35 Page 250

The conditional expected value of profit at each node in Figure 7.7 is computed by backwards analysis, just as in the previous tree. In particular, cell A22 shows the expected value of profit, as assessed before we learn the weather information but based on the assumption that we pay the cost in A2 and use the weather information according to cells H10 and H30. The conditional production quantities in cells H10 and H30 have been chosen by Solver to maximize this expected value in A22. But notice that this expected profit from this optimal production plan with (costly) weather information is $1,944,120, which is less than the expected profit ($1,991,366) that we got in Figure 7.6 without weather information. So the comparison of these two decision trees indicates that Scotia should not pay $100,000 for the privilege of basing its production decision on a forecast of the winter's general weather pattern.

If the information cost in cell A2 of Figure 7.7 were changed to 0, then all the profits in the O column would increase by $100,000, and the expected profit in cell E22 would increase to

$$\$1,944,120 + \$100,000 = \$2,044,120$$

(Changing the cost of information in cell A2 would not change the optimal production decisions in cells H10 and H30.) This increased expected profit represents our estimate of Scotia's expected profit if it could get the weather information for free, using the simplified four-value demand distributions that are assumed in these tree diagrams. Similarly, if we changed the cost in cell A2 to any other number $c$, then the expected profit in cell A22 would change to $\$2,044,120 - c$, and this expected value would be greater than the expected profit without weather information ($1,991,366) from Figure 7.6 only if

$$c \le \$2,044,120 - \$1,991,366 = \$52,754$$

That is, Scotia should not pay more than $52,754 for the privilege of basing its production decision on a perfect forecast of the winter's general weather pattern. Thus, the expected value of perfect information about next winter's weather pattern is $52,754 for Scotia, according to this decision-tree analysis. Obviously, an imperfect forecast would be worth even less to Scotia.

This conclusion differs somewhat from our estimated value of information in Section 7.2 ($42,333), which was based on many simulations sampled from the full demand distribution. There are of course different advantages to each approach. Decision trees require a great simplification of our probability distributions, and simulation involves an element of randomness in our computational results. But decision trees become more unwieldy as the complexity of the decision problem increases, and the randomness of simulation results can always be mitigated by increasing the number of simulations used, and so this book emphasizes the simulation approach as a general method of analysis.

Myerson_C07.qxd 09-09-2004 10:35 Page 251

## CASE ■ The Bates Restoration (A)

Russ Gehrig, a construction general contractor, has decided to bid for the contract to do an extensive restoration of the old Bates mansion. The contract will be awarded to the bidder who submits the lowest bid, and the winner will be paid the amount of his bid at the completion of the job. Gehrig estimates that his cost to do the work described in the contract is equally likely to be above or below $100,000, has probability 1/4 of being below $86,000, and has probability 1/4 of being above $120,000.

Gehrig knows that there will be at least one other bidder for the job, but he does not know how many. He has assessed the following probability distribution for the number of other bidders (not including himself):

| *# Other Bidders* | *Probability* |
|---|---|
| 1 | 0.2 |
| 2 | 0.3 |
| 3 | 0.3 |
| 4 | 0.1 |
| 5 | 0.1 |

He also figures that each opposing bidder's bid is equally likely to be above or below $140,000, has probability 1/4 of being below $120,000, and has probability 1/4 of being above $180,000.

To assess his risk tolerance, Gehrig thought about a simple hypothetical gamble that would either generate a loss of $25,000 or generate some profit of $$X$ for him, each with probability 1/2. He tried to ask himself, what was the smallest number $X$ such that he would be willing to accept such a gamble? If the potential profit $$X$ was less than $25,000, then the expected monetary value of the gamble would be less than $0, and he would not accept such a gamble. Indeed, he might still walk away from the gamble even if $$X$ was slightly larger than $25,000. But if the potential profit $$X$ was $40,000 or more, then he would feel willing to accept such a gamble.

# 7.4 A Simple Bidding Problem

Selecting a bid for an auction is a decision under uncertainty. A bidder might think about submitting an aggressive bid that would have a high probability of winning the auction, but such an aggressive bid might yield the bidder little or no profit in the event that it wins. On the other hand, a less aggressive bid that would allow a larger profit margin for the bidder typically would have a smaller probability of winning. So our criterion for choosing a bid must take account of both the probability of winning and the potential profit that the bid would yield if it wins.

One such criterion for determining optimal bids is expected-profit maximization. In an auction where losing bidders pay nothing and get nothing, a bidder's expected profit from submitting some bid $\beta$ is the probability of the bid $\beta$ winning the auction multiplied by the conditionally expected profit that the bidder would get if he won the auction with this bid $\beta$.

To be more precise, let us consider the auction described in part A of the Bates Restoration case. The auction in this case is an example of a first-price

Myerson_C07.qxd 09-09-2004 10:35 Page 252

sealed-bid auction in which the bidders are bidding to sell. In an auction where the bidders are *bidding to sell*, the winner of the auction will be the one with the lowest bid. In a *sealed-bid* auction, the bids are submitted independently, and there is no opportunity to revise one's bids after hearing others' bids. In a *first-price* auction where the bidders are sellers, the winner will be paid the amount of his bid to provide the services specified in the contract. Of course, the bidders who do not win are paid nothing and have no obligations to provide any services.

So suppose that we are consultants for the bidder (Gehrig) in the Bates case. Let β denote his bid, let **C** denote his cost of fulfilling the contract, and let **A** denote the lowest opposing bid submitted by other bidders. The bid β is a decision variable for Gehrig, but the cost **C** and lowest opposing bid **A** must be treated as random variables (as conventionally indicated here by boldface) because our decision maker does not know them. With this notation, the profit from submitting the bid β in this auction is

$$\text{Profit} = \beta - \mathbf{C} \text{ if } \beta < \mathbf{A}, \quad \text{and} \quad \text{Profit} = 0 \text{ if } \beta > \mathbf{A}$$

For a risk-neutral bidder, the optimal bid is one that maximizes the bidder's expected profit from submitting the bid β, which is

$$E(\text{Profit from submitting bid } \beta) = P(\mathbf{A} > \beta) * (\beta - E(\mathbf{C} | \mathbf{A} > \beta))$$

This formula expresses a natural trade-off between the goals of having a high probability of winning and of having a large profit margin when the bid wins. Choosing a lower bid β could increase the probability of winning $P(\mathbf{A} > \beta)$, but it also would decrease the profit $(\beta - E(\mathbf{C} | \mathbf{A} > \beta))$ that he could expect when he wins.

Figure 7.8 shows an analysis of this bidding problem. The cost of fulfilling the contract is simulated in cell A4 by a Generalized Lognormal random variable with Gehrig's assessed quartiles for the cost (86, 100, and 120 thousand dollars), and this random variable is echoed in cell B18. The uncertain number of other bidders is simulated in cell A6 using the discrete probability distribution shown in cells C8:D12. Opposing bids from these other bidders are simulated in cells A8:A12 by entering the formula

```
=IF(C8<=$A$6, GENLINV(RAND( ),$A$2,$B$2,$C$2), "..")
```

into cell A8, and then copying A8 to A8:A12. Notice that cells C8:C12 contain the numbers 1 to 5. So given any number from 1 to 5 in cell A6, the cells A10:A14 will contain as many simulated bids as this number in cell A6, with each bid being an independent Generalized Lognormal random variable with the assessed quartile points 120, 140, 180; and the (5–A5) extra cells in A10:A14 will contain nonnumerical text (".."). Cell A13 contains the buyer's reservation price of $250 thousand (which may also be counted as another "opposing bid," because Gehrig cannot win the auction if he bids above this value). So Gehrig can win the auction, in this simulation, only if his bid

Myerson_C07.qxd 09-09-2004 10:35 Page 253

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | Quartiles for each other bid | | | | Quartiles for Gehrig's actual cost | | | | |
| 2 | 120 | 140 | 180 | | 86 | 100 | 120 | | |
| 3 | | | | | | | | | |
| 4 | 131.936 | Gerhig's Cost to do the Bates job ($1000s) | | | | | | | |
| 5 | | | | | | | | | |
| 6 | 3 | #Other bidders | | | | | | | |
| 7 | Other bids | | Bidders | Proby | | | | | |
| 8 | 108.273 | | 1 | 0.2 | | | | | |
| 9 | 202.129 | | 2 | 0.3 | | | | | |
| 10 | 109.638 | | 3 | 0.3 | | | | | |
| 11 | .. | | 4 | 0.1 | | | | | |
| 12 | .. | | 5 | 0.1 | | | | | |
| 13 | 250 | Buyer's reserve price | | | | | | RiskTol | |
| 14 | | | | | | | | 67.88 | |
| 15 | | | | G's bid | P(Win) | E(Profit) | | CE from bid | |
| 16 | | | | 130 | 0.34283 | 7.72131 | | 4.10974 | |
| 17 | G's cost | Bid to beat | | | | | | | |
| 18 | SimTabl | 131.936 | 108.273 | | G wins? | G's profit | | | |
| 19 | 0 | 79.0744 | 110.079 | | 0 | 0 | | | |
| 20 | 0.0005 | 105.235 | 122.482 | | 0 | 0 | | | |
| 21 | 0.001 | 83.7517 | 108.935 | | 0 | 0 | | | |
| 22 | 0.0015 | 99.3089 | 105.418 | | 0 | 0 | | | |
| 23 | 0.002 | 116.732 | 147.865 | | 1 | 13.268 | | | |
| 24 | 0.0025 | 87.1279 | 102.47 | | 0 | 0 | | | |
| 25 | 0.003 | 84.941 | 110.737 | | 0 | 0 | | | |
| 2015 | 0.998 | 98.2963 | 113.021 | | 0 | 0 | | | |
| 2016 | 0.9985 | 87.1836 | 140.505 | | 1 | 42.8164 | | | |
| 2017 | 0.999 | 83.9332 | 132.979 | | 1 | 46.0668 | | | |
| 2018 | 0.9995 | 75.1348 | 111.675 | | 0 | 0 | | | |
| 2019 | 1 | 92.4269 | 119.384 | | 0 | 0 | | | |
| 2020 | | | | | | | | | |
| 2021 | FORMULAS | | | | | | | | |
| 2022 | A4. =GENLINV(RAND(),E2,F2,G2) | | | | | | | | |
| 2023 | A6. =DISCRINV(RAND(),C8:C12,D8:D12) | | | | | | | | |
| 2024 | A8. =IF(C8<=$A$6,GENLINV(RAND(),$A$2,$B$2,$C$2),"..") | | | | | | | | |
| 2025 | A8 copied to A8:A12 | | | | | | | | |
| 2026 | B18. =A4 | | | | | | | | |
| 2027 | C18. =MIN(A8:A13) | | | | | | | | |
| 2028 | E19. =IF(C19>$D$16,1,0) | | | | | | | | |
| 2029 | F19. =E19*($D$16-B19) | | | | | | | | |
| 2030 | E16. =AVERAGE(E19:E2019) | | | | | | | | |
| 2031 | F16. =AVERAGE(F19:F2019) | | | | | | | | |
| 2032 | H14. =RISKTOL(40,-25,0) | | | | | | | | |
| 2033 | H16. =CE(F19:F2019,H14) | | | | | | | | |

FIGURE **7.8**
**Model of bidding to sell in the Bates case, part A (no winner's curse)**

is less than the lowest opposing bid, which is computed in cell C18 by the formula

=MIN(A8:A13)

(The MIN function ignores any nonnumerical cells in A8:A13.)

The simulation table in cells A18:C2019 of Figure 7.8 contains 2001 simulated values of the potential cost (in B19:B2019) and the lowest opposing bid (in C19:C2019). It is good to use a large simulation table in the analysis of bidding problems, because we often want to consider bids that have a relatively small chance of winning, and so the expected profit may

Myerson_C07.qxd 09-09-2004 10:35 Page 254

depend on only a small subset of our simulation data. A proposed bid has been entered into cell D16. For each row of simulation data, the question of whether the D16 bid would win the auction in this simulation is answered (1 = yes, 0 = no) in column E by entering

```
=IF(C19>$D$16,1,0)
```

in cell E19, and then copying E19 to E19:E2019. The corresponding profit for Gehrig in each simulation row is computed in column F by entering the formula

```
=E19*($D$16–B19)
```

into cell F19, and then copying F19 to F19:F2019. This formula in F19 returns $0 if Gehrig's D16 bid would lose the auction (if E19 = 0), but it returns Gehrig's net profit from the contract (D16 – B19) if his bid would win (if E19 = 1). So Gehrig's expected profit from the bid in cell D16 can be estimated in cell F16 by the formula

```
=AVERAGE(F19:F2019)
```

This formula in F16 yields the expected value $7.72 thousand for the $130 thousand bid in Figure 7.8.

Gehrig has indicated that a gamble that would either generate a profit of about $40 thousand or a loss of $25 thousand, each with probability 1/2, might not be substantially better than a sure alternative of $0. So we can estimate that his risk tolerance is approximately $67.88 thousand, as computed by the formula =RISKTOL(40,–25,0) in cell H14 of Figure 7.8. Then Gehrig's certainty equivalent for the gamble that he gets by submitting the bid in cell D16 can be computed in cell H16 by the formula

```
=CE(F19:F2019,H14)
```

which yields the certainty-equivalent value $4.11 thousand for the $130 thousand bid in Figure 7.8.

Unfortunately, Solver may be unable to maximize F16 (or H16) by changing the bid in D16 in Figure 7.8. The reason is that the profits depend very discontinuously on the bid. Solver tries to find improvements in the target cell by trying small changes in the changing cells. Making a very small increase in the bid almost always increases Gehrig's profit here because it increases his income in all simulations where he wins. But when the increased bid goes past the bid-to-beat value from some simulation in the C19:C2019 data, then Gehrig's profit from that simulation suddenly goes to $0. So the expected profit depends very erratically on small changes in the bid, and Solver can find itself at a local maximum almost anywhere. Instead of using Solver, we should try to find the optimal bid in this example by recalculating the expected profit or certainty equivalent many times in a data table.

The data table and chart in Figure 7.9 show the results of recalculating the expected profit (in cell F16 in Figure 7.8) and the certainty equivalent (in cell H16 in Figure 7.8) when the bid in cell D16 of Figure 7.8 is changed

| | K | L | M |
|---|---|---|---|
| 1 | Bid | E($) | CE |
| 2 | (,D16) | 7.72131 | 4.10974 |
| 3 | 100 | -6.8456 | -17.888 |
| 4 | 105 | -1.5578 | -12.191 |
| 5 | 110 | 2.3415 | -7.2104 |
| 6 | 115 | 5.10983 | -3.3534 |
| 7 | 120 | 6.52351 | 0.8383 |
| 8 | 125 | 7.26213 | 3.17185 |
| 9 | 130 | 7.72131 | 4.10974 |
| 10 | 135 | 7.62366 | 4.38832 |
| 11 | 140 | 7.49047 | 4.46952 |
| 12 | 145 | 7.13795 | 4.28855 |
| 13 | 150 | 7.15569 | 4.58916 |
| 14 | 155 | 6.92797 | 4.36377 |
| 15 | 160 | 6.9038 | 4.42034 |
| 16 | 165 | 6.52826 | 4.06806 |
| 17 | 170 | 6.38742 | 3.91026 |
| 18 | 175 | 6.19858 | 3.73579 |
| 19 | 180 | 5.71093 | 3.36252 |
| 20 | 200 | 4.94106 | 2.60721 |
| 21 | 225 | 3.92994 | 1.81518 |
| 22 | 250 | 0 | 0 |
| 23 | | | |
| 24 | max | 7.72131 | 4.58916 |
| 25 | | | |

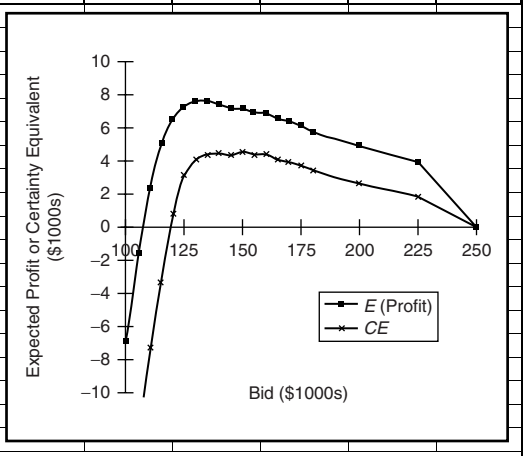


Chart plots (K3:K22,L3:L22),(K3:K22,M3:M22)

FORMULAS

L2. =F16 M2. =H16

L3:M22. {=TABLE(,D16)}

L24. =MAX(L3:L22)

M24. =MAX(M3:M22)

FIGURE **7.9**
**Finding optimal bids for the Bates case, part A (no winner's curse)**

to a range of values from $100 thousand to $250 thousand. According to this analysis, if Gehrig's goal is to maximize expected profit, then the optimal bid is $130 thousand, which is estimated to yield an expected profit of about $7.72 thousand. On the other hand, if Gehrig's limited risk tolerance is taken into account, then we find that a somewhat higher bid may be optimal. With this simulation, we find a maximal *CE* of about $4.59 thousand, when the bid is $150 thousand.

## 7.5 The Winner's Curse

In the previous section we considered an auction where bidders were bidding to sell. Now let us consider the problem of someone who is trying to buy something in a bargaining situation, taking a well-known example from an article on "The Market for Lemons" by George Akerlof (1970). In this example, we are trying to help someone who has been negotiating to buy a used car. The problem is that the buyer has substantial uncertainty about the condition of this used car and recognizes that the seller has much better information about the car's condition, because the seller has owned it for a long time. (A used car of poor quality is often called a "lemon" in colloquial American English.)

Bargaining processes can be quite complicated, so let us simplify the situation by supposing that our buyer can make a take-it-or-leave-it offer to buy at some price. Then the buyer's only question is what price to offer the

Myerson_C07.qxd 09-09-2004 10:35 Page 256

seller, and this price is like a bid in an auction. If the seller understands that this price is the buyer's final offer, then the seller should accept it if the offered price is more than the seller's own value (or *reserve price*) for keeping the car. If this were an auction, then our buyer would win the car if the bid were higher than the highest opposing bid. (In a situation of *bidding to buy*, it is the highest bid that wins!) So we can view this bargaining situation as a simple auction in which this final offer price is the buyer's bid, and the seller's value for keeping the car is the opposing bid that our buyer's bid must exceed to get the car. If the buyer's bid is greater than the seller's value, then the buyer's profit is the buyer's value of the car minus the buyer's bid. But if the buyer's bid is less than the seller's value then the buyer gets zero profit because there is no transaction.

As long as the car is worth more to the buyer than the seller, the buyer's ideal bid would be one cent more than the seller's value for keeping the car. But of course the buyer does not know the seller's reserve price. So this bargaining problem is a decision under uncertainty for the buyer. To find an optimal bid that maximizes the buyer's expected profit (assuming risk neutrality), we must ask the buyer to assess probability distributions for the seller's reserve price and the buyer's value of getting the car.

Suppose that the (female) buyer tells us that the value of the car to her will depend on the quality of its condition (which determines how long it will run before breaking down). Because she does not know the actual condition of the car, she thinks that it could be worth anywhere from \$0 to \$1500 to her. The buyer also believes that the (male) seller's value of keeping the car could be anywhere from \$0 to \$1000, also depending on the condition of the car (which the seller knows). In each case, the buyer is willing to assess a Uniform distribution for these two unknowns. Let **V** denote the value that the car could yield for the buyer, and let **A** denote the seller's reserve value for keeping the car. With this notation, the buyer's profit from a bid $\beta$ is

$$\text{Profit} = \mathbf{V} - \beta \text{ if } \beta > \mathbf{A}, \quad \text{and} \quad \text{Profit} = 0 \text{ if } \beta < \mathbf{A}$$

Because the buyer's value **V** is drawn from a Uniform distribution from \$0 to \$1500, the expected value of the car to the buyer is

$$E(\mathbf{V}) = 750$$

For any bid $\beta$ that is between \$0 and \$1000, the probability of $\beta$ being greater than the seller's reserve value **A** is

$$P(\mathbf{A} < \beta) = \beta/1000$$

because **A** is drawn from a Uniform distribution from \$0 to \$1000.

Consider now the consequences of bidding $\beta = 375$. This bid $\beta$ happens to be the value that maximizes the quantity $P(\mathbf{A} < \beta)*(E(\mathbf{V}) - \beta)$, and it would be accepted by the seller if $\mathbf{A} < 375$, which has probability 375/1000. So the buyer has probability 0.375 of being able to buy the car for \$375. Since the expected value of the car for her is \$750, we might initially guess that this bid should give her a positive expected profit.

Myerson_C07.qxd 09-09-2004 10:35 Page 257

| | A | B | C | D | E | F |
|---|---|---|---|---|---|---|
| 1 | AKERLOF'S LEMON PROBLEM | | | | | |
| 2 | | | E(ValueToBuyer) | | E(ValueToBuyer\|BidWins) | |
| 3 | | | 761.50 | | | 278.90 |
| 4 | | | | | | |
| 5 | | | | | E(Buyer'sProfit) | P(BidWins) |
| 6 | Correl | (corands) | | | -35.23 | 0.3666 |
| 7 | 1 | 0.17269 | 0.17269 | | -35.23 | |
| 8 | | Seller's value | Buyer's value | Buyer's bid | | |
| 9 | SimTabl | 172.69 | 259.03 | 375 | B's Profit | ValToB\|Win |
| 10 | 0 | 129.82 | 194.73 | | -180.27 | 194.73 |
| 11 | 0.0025 | 44.71 | 67.06 | | -307.94 | 67.06 |
| 12 | 0.005 | 511.11 | 766.67 | | 0.00 | .. |
| 13 | 0.0075 | 478.56 | 717.84 | | 0.00 | .. |
| 14 | 0.01 | 108.49 | 162.73 | | -212.27 | 162.73 |
| 15 | 0.0125 | 727.37 | 1091.06 | | 0.00 | .. |
| 16 | 0.015 | 29.86 | 44.78 | | -330.22 | 44.78 |
| 17 | 0.0175 | 531.87 | 797.80 | | 0.00 | .. |
| 18 | 0.02 | 273.75 | 410.62 | | 35.62 | 410.62 |
| 406 | 0.99 | 628.00 | 942.00 | | 0.00 | .. |
| 407 | 0.9925 | 753.02 | 1129.53 | | 0.00 | .. |
| 408 | 0.995 | 677.36 | 1016.04 | | 0.00 | .. |
| 409 | 0.9975 | 174.76 | 262.14 | | -112.86 | 262.14 |
| 410 | 1 | 168.38 | 252.57 | | -122.43 | 252.57 |
| 411 | | | | | | |
| 412 | FORMULAS | | | | | |
| 413 | B7:C7. {=CORAND(A7)} | | | | | |
| 414 | B9. =1000*B7 | | | | | |
| 415 | C9. =1500*C7 | | | | | |
| 416 | E10. =IF($D$9>B10,C10-$D$9,0) | | | | | |
| 417 | F10. =IF($D$9>B10,C10,"..") | | | | | |
| 418 | E10:F10 copied to E10:F410 | | | | | |
| 419 | C3. =AVERAGE(C10:C410) | | | | | |
| 420 | F3. =AVERAGE(F10:F410) | | | | | |
| 421 | E6. =AVERAGE(E10:E410) | | | | | |
| 422 | F6. =COUNT(F10:F410)/COUNT(C10:C410) | | | | | |
| 423 | E7. =F6*(F3-D9) | | | | | |

FIGURE **7.10** **Bidding to buy a "lemon"**

There is a serious problem with this analysis, however, because we have not assessed the relationship between the seller's value **A** and the buyer's value **V**. Both are unknown quantities in the buyer's model because they depend on the condition of the car, which is known to the seller but not to the buyer. So there is good reason to suspect that these two random variables should be highly correlated. To make things simple, suppose that the buyer tells us that **A** and **V** have a perfect correlation of 1, because the value of the car to her would always be 50% more than the value of the car to the seller. That is, suppose that **V** depends on **A** by the simple linear relation

$$\mathbf{V} = 1.5 * \mathbf{A}$$

With this equation, having **A** drawn from a Uniform distribution over \$0 to \$1000 does indeed make **V** a random variable with a Uniform distribution between \$0 and \$1500. Notice also that this equation implies that the car would always have more value to the buyer than to the seller.

But a simulation analysis of this bargaining situation, shown in Figure 7.10, indicates that the buyer's expected profit from bidding \$375 is actually

Myerson_C07.qxd 09-09-2004 10:35 Page 258

negative. An average of the results of 401 simulations, computed in cell E6 of Figure 7.10, indicates that bidding $375 would actually yield an expected net loss of about $35!

What has gone wrong? The problem is that our bid is accepted only when the seller's value **A** for the car is relatively low (less than $375), in which case the buyer's value **V** must also be relatively low (less than $562.50 = 1.5∗375), and the mean of these relatively low values is substantially less than $750. The conditional expectation $E(\mathbf{V}|\mathbf{A} < 375)$ has been estimated from simulation data in cell F3 of Figure 7.10, and these results show that when the bid $375 is accepted, the conditionally expected value of the car for the buyer is about $280. So learning that the bid of $375 is accepted should make the buyer expect to lose almost $100 from the transaction!

The general formula for the buyer's profit from submitting a bid β is

$$E(\text{Profit}) = P(\mathbf{A} < \beta)*(\text{E}(\mathbf{V}|\mathbf{A} < \beta) - \beta)$$

This formula is applied in cell E7 of Figure 7.10, using the simulation estimates of $P(\mathbf{A} < \beta)$ and $E(\mathbf{V}|\mathbf{A} < \beta)$ from cells F6 and F3, and the result in cell E7 is indeed the same as the expected profit that we estimated in cell E6 by averaging the profits from this same simulation data. In fact, when these calculations are repeated with any positive bid β entered into cell D9 of Figure 7.10, the expected buyer's profit in cells E6 and E7 is always negative. So the optimal bid for the buyer in this example is $0.

The problem with our initial guess was that we used $E(\mathbf{V})$ instead of $E(\mathbf{V}|\mathbf{A} < \beta)$ in the formula for expected profit, and $E(\mathbf{V}|\mathbf{A} < \beta)$ is typically less than $E(\mathbf{V})$. That is, learning that a bid β has been accepted by the seller may make the buyer revise downward her expected value of the object that she is buying. This inequality

$$E(\mathbf{V}|\mathbf{A} < \beta) < E(\mathbf{V})$$

is called the *winner's curse*. Such a winner's curse can arise in bidding problems whenever the value of winning to the bidder is positively correlated with the other opposing bids. In this example, if the buyer's value **V** had been independent of the seller's value **A**, then there would have been no winner's curse, because learning that **A** was less than the bid β would convey no new information about the value **V** for the buyer.

Now let us reconsider the Bates case. In such a case of bidding to sell, the winner's curse occurs when the conditional expected cost given that our bid wins would be greater than the expected cost according to our current beliefs. That is, in the notation of Section 7.4, the winner's curse would be the inequality

$$E(\mathbf{C}|\mathbf{A} > \beta) > E(\mathbf{C})$$

In Figure 7.8, the cost random variable in cell A4 depends on its own RAND( ), which is independent of the RAND( )s that drive the opposing bids in A6:A12 of Figure 7.8. So this model implicitly assumes that Gehrig's cost of doing the Bates job would be independent of the opposing bids, which means

Myerson_C07.qxd 09-09-2004 10:35 Page 259

that we have implicitly assumed that there would be no winner's curse in this case. Given the importance of the winner's curse in Akerlof's "lemon" example, this independence assumption should be seriously scrutinized.

Recall that two random variables are said to be independent if learning about one of them would not change our beliefs about the other. Should Gehrig feel this way about his cost of the Bates job and the lowest opposing bid that he has to beat to win the contract? There are very good reasons to think that he might not consider these unknown quantities to be independent. After all, the opposing bids will be generated by other contractors using their own best estimates of the costs of doing the Bates job, and there is some chance that these other contractors may have noticed some difficulties in the job that Gehrig has overlooked. Thus, if all the opposing bids turn out to be surprisingly high, then Gehrig might infer that the other contractors have generated higher cost estimates for the job, which might make him feel that his own cost estimates were probably too low.

To assess a possible winner's curse effect in this case, we could ask the decision maker to think about how he might revise the median value in his cost distribution if he learned that the lowest opposing bid was actually somewhat higher than expected. To be specific, let us apply here the method for subjective assessment of correlations that we developed in Section 5.8. The first step of that method is to ask about how the decision maker's beliefs about cost would change if he learned that the lowest opposing bid was equal to the value at the 0.75 percentile in its probability distribution. We could find any percentile value in the bid-to-beat distribution by sorting the bid-to-beat simulation data in cells C19:C1019 of Figure 7.8, for example, but to keep all the analysis on one page, we have generated bid-to-beat simulation data in the new spreadsheet shown in Figure 7.11. In Figure 7.11, the random number of other bidders in cell A5, the bids from these other bidders in cells A7:A11, and the bid-to-beat (that is, the minimum of the opposing bids, including the buyer's reservation price) in cell B15 are all generated in the same way as in (slightly different cells of) Figure 7.8.

The opposing bids are modeled in cells A7:A11 of Figure 7.11 by random variables that have their own separate RAND( )s, and so we are implicitly assuming here (as in Figure 7.8) that the opposing bids are independent of each other. But we are also going to allow here that these bids may all have a positive correlation with the cost of doing this job, because such positive correlation between opposing bids and the value of winning is the source of the winner's curse. So we might have considered a model where the opposing bid's mutual correlation with the cost makes these bids positively correlated with each other as well. But there is a practical disadvantage to such a model, because positively correlating the opposing bids is effectively similar to reducing the number of opposing bidders. Reducing the number of opposing bidders would be beneficial for our decision maker and could hide the adverse effects of the winner's curse. So our assumption of independent opposing bidders here is a conservative assumption that will help us to see more clearly the effects of the winner's curse.

| | A | B | C | D | E | F | G | H | I | J |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Bid-to-beat distributn | | | Quartiles for Gehrig's Cost of doing job | | | | | | |
| 2 | OtherBids' quartiles | | | 86 | 100 | 120 | | ($1000s) | | |
| 3 | 120 | 140 | 180 | BidToBeat's distribution is generated | | | | | | |
| 4 | #Other bidders | | | from simulation data in B16:B416, | | | | | | |
| 5 | 4 | | | using the PERCENTILE function. | | | | | | |
| 6 | Bids | Bidders | Proby | | | | | | | |
| 7 | 139.814 | 1 | 0.2 | 0.3391 | Correl(BidToBeat,GsCost) | | | | | |
| 8 | 358.246 | 2 | 0.3 | This correlation satisfies: | | | | | | |
| 9 | 166.185 | 3 | 0.3 | P(Cost<106 \| BidToBeat=136) = 1/2 | | | | | | |
| 10 | 134.785 | 4 | 0.1 | GivenValue | | ConditionalMedian | | | | |
| 11 | .. | 5 | 0.1 | Corand | 0.75 | 0.5905 | | Corand | | |
| 12 | 250 | Buyer's reserve | | BidToBeat | 136 | 106 | | Cost | | |
| 13 | | | | | | | | | | |
| 14 | | Bid to beat | | SIMULATION MODEL | | | | G's bid | | |
| 15 | SimTable | 134.785 | | Corands | 0.2135 | 0.1282 | | 150 | | |
| 16 | 0 | 108.784 | | BidToBeat | | G's cost | | | | |
| 17 | 0.0025 | 119.042 | | SimTable | 110.03 | 78.944 | | G wins? | G'sProfit | |
| 18 | 0.005 | 105.262 | | 0 | 109.58 | 87.206 | | 0 | 0 | |
| 19 | 0.0075 | 135.812 | | 0.0005 | 121.91 | 99.873 | | 0 | 0 | |
| 20 | 0.01 | 129.655 | | 0.001 | 115.71 | 87.719 | | 0 | 0 | |
| 21 | 0.0125 | 115.75 | | 0.0015 | 112.18 | 112.16 | | 0 | 0 | |
| 22 | 0.015 | 106.072 | | 0.002 | 136.42 | 106.23 | | 0 | 0 | |
| 23 | 0.0175 | 105.313 | | 0.0025 | 108.15 | 142.92 | | 0 | 0 | |
| 24 | 0.02 | 118.466 | | 0.003 | 136.46 | 104.4 | | 0 | 0 | |
| 25 | 0.0225 | 118.317 | | 0.0035 | 145.12 | 89.386 | | 0 | 0 | |
| 26 | 0.025 | 103.426 | | 0.004 | 224.61 | 86.913 | | 1 | 63.087 | |
| 27 | 0.0275 | 127.864 | | 0.0045 | 111.38 | 71.343 | | 0 | 0 | |
| 28 | 0.03 | 111.019 | | 0.005 | 115.31 | 169.26 | | 0 | 0 | |
| 2016 | [SimTable data to B416] | | | 0.999 | 109.59 | 95.717 | | 0 | 0 | |
| 2017 | | | | 0.9995 | 142.98 | 144.07 | | 0 | 0 | |
| 2018 | | | | 1 | 250 | 101.4 | | 1 | 48.6 | |
| 2019 | | | | | | | | | | |
| 2020 | | | | | | | | Pr(Win) | EProfit | |
| 2021 | | | | | | | | 0.15792 | 4.5233 | |
| 2022 | FORMULAS | | | | | | | RiskTol | CE | |
| 2023 | A5. =DISCRINV(RAND(),B7:B11,C7:C11) | | | | | | | 67.88 | 2.2873 | |
| 2024 | A7. =IF(B7<=$A$5,GENLINV(RAND(),$A$3,$B$3,$C$3),"..") | | | | | | | | | |
| 2025 | A7 copied to A7:A11 | | | | B15. =MIN(A7:A12) | | | | | |
| 2026 | F11. =MIDRAND($D$7,$E$11) | | | | | | | | | |
| 2027 | E12. =PERCENTILE($B$16:$B$416,E11) | | | | | | | | | |
| 2028 | F12. =GENLINV(F11,$D$2,$E$2,$F$2) | | | | | | | | | |
| 2029 | E15:F15. {=CORAND(D7)} | | | | | | | | | |
| 2030 | E17. =PERCENTILE($B$16:$B$416,E15) | | | | | | | | | |
| 2031 | F17. =GENLINV(F15,$D$2,$E$2,$F$2) | | | | | | | | | |
| 2032 | H18. =IF(E18>$H$15,1,0) | | | | I18. =H18*($H$15-F18) | | | | | |
| 2033 | H18:I18 copied down | | | | | | | | | |
| 2034 | H2021. =AVERAGE(H18:H2018) | | | | I2021. =AVERAGE(I18:I2018) | | | | | |
| 2035 | H2023. =RISKTOL(40,-25,0) | | | | I2023. =CE(I18:I2018,H2023) | | | | | |

FIGURE **7.11**
**Bidding model for the Bates case (part B) with winner's curse**

With the simulated bid-to-beat data in cells B16:B416 of Figure 7.11, the bid-to-beat value with cumulative probability 0.75 is computed in cell E12 by the formula

=PERCENTILE(B16:B416,E11)

where cell E11 contains the value 0.75. This 0.75 percentile value is shown to be $136 thousand in cell E12 of Figure 7.11. So to apply the method for subjectively assessing correlations from Chapter 5, we could ask Gehrig the following question:

Myerson_C07.qxd 09-09-2004 10:35 Page 261

Imagine for a moment that the lowest opposing bid turned out to be $136 thousand. Obviously this information would simplify your bidding problem and make it easier for you to win the contract. But think now how this information would affect your beliefs about the cost of fulfilling the contract. After getting this information, would you still feel that your own cost of doing the job was equally likely to be above or below $100 thousand (which was the median in the originally assessed distribution for this cost)? If not, then for what number $X$ would you say that your cost was equally likely to be above or below $\$X$ thousand, given this information that the bid to beat is $136 thousand?

Suppose that in answer to these questions, Gehrig tells us that the conditional median of his cost to do the Bates job would increase to $X$ = $106 thousand, given the information that the lowest opposing bid was $136 thousand. The fact that Gehrig's beliefs about his costs would be affected by information about the lowest opposing bid is an indication that the independence assumption in Figure 7.8 was wrong, and the model needs to be revised to take account of the winner's curse. Such a revised model is developed in cells E17 and F17 of Figure 7.11, which correlated random variables to simulate the lowest opposing bid and Gehrig's cost of doing the job. The correlation in these two cells is achieved by having them depend on a pair of correlated CORAND values in cells E15:F15, which contain the array formula

{=CORAND(D7)}

Let us defer for a moment the question of what correlation is in cell D7, because there is also something new to discuss about the way that the random variables in cells E17:F17 of Figure 7.11 are generated from these CORANDs in E15:F15. The assessed quartile boundary points for Gehrig's cost are in cells D2:F2 here, and so the simulated cost can be generated in cell F17 by the formula

=GENLINV(F15,$D$2,$E$2,$F$2)

We want the other CORAND in cell E15 to generate a simulated bid-to-beat value, but our model for generating simulated bid-to-beat values required six different RAND( )s (in cells A5 and A7:A11). We need another way to generate a simulated bid-to-beat value that can be driven by one CORAND value, through its inverse cumulative function. We can get this inverse cumulative function (or a good estimate of it) by applying the PERCENTILE function to the simulated bid-to-beat data in cells B16:B416 of Figure 7.11 (as also illustrated in Figure 7.12). That is, for any number $p$ between 0 and 1, the formula PERCENTILE(B16:B416, $p$) here would return the bid-to-beat value that has cumulative probability $p$, as estimated from our simulation data in B16:B416 in Figure 7.11. Thus, a simulated bid-to-beat value can be made in cell E17 with the formula

=PERCENTILE($B$16:$B$416,E15)

| | U | V | W | X |
|---|---|---|---|---|
| 1 | %rank | BidToBeat | | |
| 2 | 0 | 101.435 | | |
| 3 | 0.05 | 104.897 | | |
| 4 | 0.1 | 107.166 | | |
| 5 | 0.15 | 108.429 | | |
| 6 | 0.2 | 109.828 | | |
| 7 | 0.25 | 111.373 | | |
| 8 | 0.3 | 112.796 | | |
| 9 | 0.35 | 114.152 | | |
| 10 | 0.4 | 115.688 | | |
| 11 | 0.45 | 117.767 | | |
| 12 | 0.5 | 119.158 | | |
| 13 | 0.55 | 120.882 | | |
| 14 | 0.6 | 123.997 | | |
| 15 | 0.65 | 126.946 | | |
| 16 | 0.7 | 131.066 | | |
| 17 | 0.75 | 135.999 | | |
| 18 | 0.8 | 142.914 | | |
| 19 | 0.85 | 154.087 | | |
| 20 | 0.9 | 171.931 | | |
| 21 | 0.93 | 187.843 | | |
| 22 | 0.95 | 207.598 | | |
| 23 | 0.96 | 228.909 | | |
| 24 | 0.97 | 247.269 | | Chart plots (U2:U27,V2:V27) |
| 25 | 0.98 | 250 | | FORMULAS |
| 26 | 0.99 | 250 | | V2. =PERCENTILE($B$16:$B$416,U2) |
| 27 | 1 | 250 | | V2 copied to V2:V27 |

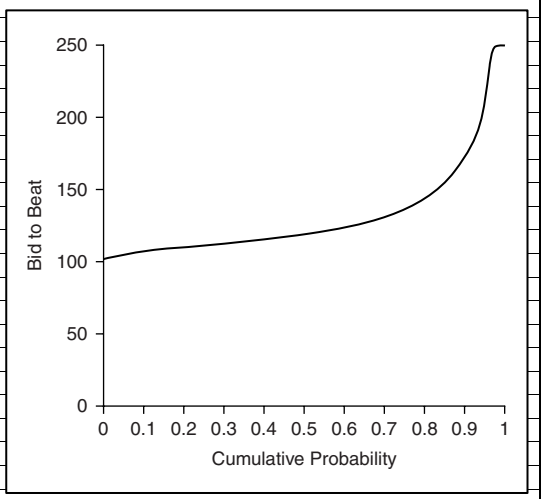


FIGURE **7.12**
**Inverse cumulative distribution for the opposing bid-to-beat in Bates (B)**

This method of using PERCENTILE to generate a random variable from the results of another simulation table may be called *modular simulation*. (See also Section 4.8.7 on nonparametric distributions.)

The next task is to find what correlation between the cost and the bid-to-beat would correspond to the conditional median assessment that Gehrig has given us, and this task is accomplished in the D7:H12 box in Figure 7.11. We can start in this box by guessing any value for the correlation and entering this value into cell D7. With the value 0.75 in cell E11, the formula in cell F11

=MIDRAND($D$7,$E$11)

returns the conditional median value for one of two cells in a CORAND array with correlation D7 when the other cell has the given value 0.75. So the formula in cell F12

=GENLINV(F11,$D$2,$E$2,$F$2)

returns the conditional median of the cost given that the bid-to-beat is at its 0.75 percentile value. According to the decision maker's subjective assessment, this conditional median of the cost should be $x$ = \$106 thousand. So we can use Excel's GoalSeek tool to set the conditional median in cell F12 to the value 106 by changing the correlation in cell D7. The result is that GoalSeek returns the correlation value 0.3391 in cell D7, as shown in Figure 7.11.

Myerson_C07.qxd 09-09-2004 10:35 Page 263

| | K | L | M |
|---|---|---|---|
| 1 | Bid | E($) | CE |
| 2 | (,H15) | 4.52334 | 2.28727 |
| 3 | 100 | -7.1728 | -19.368 |
| 4 | 110 | -0.5706 | -11.481 |
| 5 | 120 | 2.28882 | -5.8921 |
| 6 | 130 | 3.72369 | -2.6993 |
| 7 | 140 | 4.22907 | 1.47469 |
| 8 | 150 | 4.52334 | 2.28727 |
| 9 | 155 | 4.46498 | 2.31678 |
| 10 | 160 | 4.0783 | 2.10404 |
| 11 | 165 | 4.35537 | 2.35407 |
| 12 | 170 | 4.28045 | 2.3328 |
| 13 | 175 | 3.87173 | 2.09502 |
| 14 | 180 | 3.9245 | 2.13533 |
| 15 | 185 | 3.78983 | 2.10687 |
| 16 | 190 | 3.74944 | 2.06668 |
| 17 | 195 | 3.93857 | 2.16082 |
| 18 | 200 | 3.50777 | 1.9459 |
| 19 | 210 | 3.36584 | 1.82446 |
| 20 | 225 | 3.33608 | 1.72372 |
| 21 | 240 | 2.98372 | 1.45282 |
| 22 | 250 | 0 | 0 |
| 23 | | | |
| 24 | max | 4.52334 | 2.35407 |
| 25 | | | |

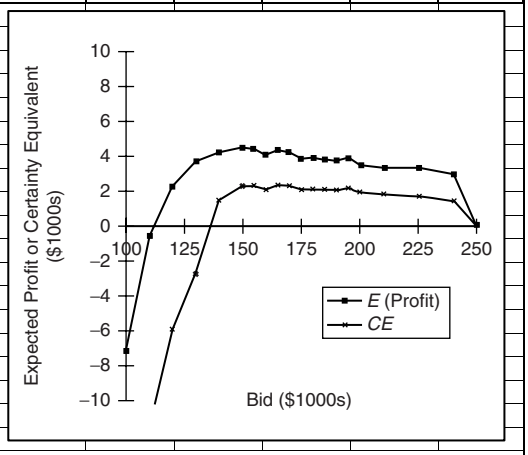


Chart plots (K3:K22,L3:L22),(K3:K22,M3:M22)

FORMULAS

L2. =I2021 M2. =I2023

L3:M22. {=TABLE(,H15)}

L24. =MAX(L3:L22)

M24. =MAX(M3:M22)

FIGURE **7.13**
**Finding optimal bids for the Bates case (part B) with winner's curse**

The rest of the analysis in Figure 7.11 is essentially the same as in Figure 7.8. Figure 7.13 shows the results of applying different bids to 2001 simulations of the bid-to-beat and cost random variables from the model with winner's curse in Figure 7.11. Comparing Figure 7.13 to Figure 7.9, we can see that the winner's curse decreases Gehrig's expected profits. With the winner's curse here, the expected profit is now maximized by a bid of $150 thousand, which yields an expected profit of only $4.52 thousand (compared to a maximal expected profit of $7.72 thousand from a bid of $130 thousand in the model without the winner's curse). When Gehrig's limited risk tolerance is taken into account, the results in Figure 7.13 suggest that Gehrig should can get a maximal *CE* of $2.35 thousand by using an optimal bid of about $165 thousand. So in this case, both the winner's curse and risk aversion have the effects of increasing Gehrig's optimal bid and of reducing his value of participating in this auction.

## 7.6 Analyzing Competitive Behavior

Decision makers often have great uncertainty about what their competitors are likely to do. To model such uncertainty about other people's behavior, we may treat their actions as random variables, and we may try to assess probability distributions for these random variables by the usual methods of subjective assessment. To think carefully about what our competitors are

Myerson_C07.qxd 09-09-2004 10:35 Page 264

## CASE ■ The Valdez Wilderness

The federal government is selling by auction the rights to all unexplored oil reserves in the Valdez Wilderness area. Exploratory drilling by government engineers has shown that oil exists under the Valdez Wilderness, but the amount of oil is still an unknown quantity. The staff geologists at the Alpha Petroleum Company have estimated that the total value of the Valdez oil reserves may be $10 million. But this is only an imperfect estimate. Past experience with such new oil fields has suggested that the ratio of the actual value of the oil to the geologists' estimate of it may be a Lognormal random variable with log-mean 0 and log-standard-deviation 0.35.

In the auction, three other oil companies will be bidding against Alpha for the rights to this oil in the Valdez Wilderness. Each company will submit a sealed bid, and the company that submits the highest bid will pay the amount of its bid and will win the rights to all oil in the Valdez Wilderness. We may assume that each of these companies has its own team of geologists who are independently generating their own similarly imperfect estimates of the value of the Valdez oil.

In Alpha's first meeting to discuss bidding strategies, a senior executive remarked that, in his experience, many people liked to apply the simple rule of bidding 80% of their geologists' estimated value. But one of his colleagues cautioned that although that strategy may be commonly used in auctions for some kinds of older, well-developed oil fields, it might not be so reasonable for new oil fields like Valdez.

likely to do, it is often useful to try to imagine ourselves in their positions, but then a couple of special analytical difficulties arise.

One difficulty is that our competitors may have some information that we do not know. Since this information is unknown to us, it should be modeled (like any other unknown quantity) as a random variable in our decision model. Then we may focus on the question of how each competitor's actions are likely to depend on the random variables that represent this private information. That is, we may ask what strategy each competitor is likely to use for making decisions as a function of personal information, even though we do not have that information ourselves.

When we conjecture that our competitors will use some particular strategy, we should also ask whether they would really want to use such a strategy, given their own preferences and information. That is, unless we think that our competitors are likely to behave foolishly, we should ask whether the strategy that we attribute to them actually would make sense if they did a careful decision analysis of their situation.

But this question may lead to a second basic difficulty. Our competitors' optimal strategies may depend on what we are predicted to do. So when we really imagine ourselves in our competitors' positions, we may realize that they are busy worrying about what we are likely to do, and indeed they may be trying to imagine themselves in our position. So it seems that a vicious cycle may arise. To find our optimal decision, we must analyze our competitors' decision problems, but their decision problems cannot be solved without an analysis of our own decision problem. The methods of game theory may be used to break through this vicious cycle. A game model

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 10 | Alpha's estimated value of oil in Valdez ($millions) | | | | | | | |
| 2 | 10.34575 | Actual value of oil in Valdez | | | | | | | |
| 3 | 13.09807 | Opponent's Estimate (1) | | | | | | | |
| 4 | 6.646099 | Opponent's Estimate (2) | | | | | | | |
| 5 | 9.347459 | Opponent's Estimate (3) | | | | | | | |
| 6 | | | | | | | | | |
| 7 | Fraction of Estimate Bid by Opponents | | | | | A'sBid | | | |
| 8 | | | | 0.8 | | 8 | | | |
| 9 | | Value | Highest | | | | | | |
| 10 | | of Oil | OppEst | | Highest | | | | |
| 11 | SimTable | 10.34575 | 13.09807 | | OppBid | AWins? | AsProfit | V\|AWins | |
| 12 | 0 | 6.754736 | 8.789973 | | 7.031979 | 1 | -1.24526 | 6.754736 | |
| 13 | 0.002 | 12.80586 | 18.48945 | | 14.79156 | 0 | 0 | .. | |
| 14 | 0.004 | 6.914 | 6.27317 | | 5.018536 | 1 | -1.086 | 6.914 | |
| 15 | 0.006 | 5.941167 | 10.23081 | | 8.184651 | 0 | 0 | .. | |
| 16 | 0.008 | 9.534073 | 12.11897 | | 9.695179 | 0 | 0 | .. | |
| 17 | 0.01 | 5.636246 | 13.2329 | | 10.58632 | 0 | 0 | .. | |
| 507 | 0.99 | 4.61466 | 4.776742 | | 3.821393 | 1 | -3.38534 | 4.61466 | |
| 508 | 0.992 | 13.95113 | 17.55869 | | 14.04695 | 0 | 0 | .. | |
| 509 | 0.994 | 7.454461 | 6.874322 | | 5.499458 | 1 | -0.54554 | 7.454461 | |
| 510 | 0.996 | 6.702069 | 11.31236 | | 9.049884 | 0 | 0 | .. | |
| 511 | 0.998 | 10.46207 | 12.68126 | | 10.145 | 0 | 0 | .. | |
| 512 | 1 | 10.92589 | 12.84069 | | 10.27255 | 0 | 0 | .. | |
| 513 | | | | | | | | | |
| 514 | FORMULAS | | | | P(A Wins) | | E(A's $) | E(V\|AWins) | |
| 515 | A2. =A1*EXP(NORMINV(RAND(),0,0.35)) | | | | | 0.2415 | -0.19141 | 7.207488 | |
| 516 | A3. =$A$2/EXP(NORMINV(RAND(),0,0.35)) | | | | | | -0.19141 | | |
| 517 | A3 copied to A3:A5 | | | | | | | | |
| 518 | B11. =A2 | | | | | Bid | E(A'sProfit) | | |
| 519 | C11. =MAX(A3:A5) | | | | | (,F8) | -0.19141 | | |
| 520 | E12. =$D$8*C12 | | | | | 0 | 0 | | |
| 521 | F12. =IF($F$8>E12,1,0) | | | | | 1 | 0 | | |
| 522 | G12. =IF($F$8>E12,B12-$F$8,0) | | | | | 2 | 0 | | |
| 523 | H12. =IF($F$8>E12,B12,"..") | | | | | 3 | 0.001754 | | |
| 524 | E12:H12 copied to E12:H512 | | | | | 4 | 0.009801 | | |
| 525 | F515. =AVERAGE(F12:F512) | | | | | 5 | 0.002494 | | |
| 526 | G515. =AVERAGE(G12:G512) | | | | | 6 | 0.033134 | | |
| 527 | H515. =AVERAGE(H12:H512) | | | | | 7 | 0.002479 | | |
| 528 | G516. =F515*(H515-F8) | | | | | 8 | -0.19141 | | |
| 529 | G519. =G515 | | | | | 9 | -0.47129 | | |
| 530 | G520:G530. {=TABLE(,F8)} | | | | | 10 | -0.86409 | | |

FIGURE **7.14**
**Alpha's bidding for Valdez when others bid 80% of their estimated values**

is an attempt to look at a competitive situation from all sides at once, ours and our competitors'. In such a game model, any quantity that is uncertain to any player must be described by a random variable, so even numbers that we actually know may become random variables in a game model. To illustrate these ideas, let us consider another example of bidding to buy.

Figure 7.14 shows a simulation model of Alpha Petroleum's bidding problem and evaluates the scenario proposed by the senior executive. Alpha's estimated value of the oil ($10 million) is listed in cell A1. Then the unknown actual value of the oil in Valdez is simulated in cell A2 by the formula

=A1*EXP(NORMINV(RAND( ),0,0.35))

This formula makes the ratio of the oil's actual value to Alpha's estimate of it a Lognormal random variable with log-mean 0 and log-standard-deviation 0.35,

Myerson_C07.qxd 09-09-2004 10:35 Page 266

as suggested by past experience with the geologists' estimates of other new oil fields. Any Lognormal random variable with log-mean 0 will be equally likely to be above or below 1, and so the estimate 10 is the median value of the Valdez oil in this model.

The other companies' bids will depend on their own geologists' estimates of the value of the Valdez oil. Of course, each company knows only what its own geologists are saying. The case suggests that each company's estimated value depends on the actual value in the same way, but independently. Thus, the model is constructed so that the ratio of the actual value to each company's estimated value is an independent random variable drawn from the same Lognormal distribution with log-mean 0 and log-standard-deviation 0.35. Given the simulated value of the oil in cell A2, the three other companies' estimates are simulated in cells A3:A5 by the formula

```
=$A$2/EXP(NORMINV(RAND( ),0,0.35))
```

The senior executive in the case suggested that each company may submit a bid equal to 80% of its estimated value of the Valdez oil. In this analysis, we will follow the suggestion of considering strategies equal to some fixed fraction of the estimated value, but we will extend the analysis by considering fixed fractions other than the 80% the senior executive suggested.

**NOTE** (Other strategies in which the bid is some nonlinear function of the estimate could be analyzed by methods that are introduced in Chapter 8, as in Figures 8.3 and 8.7.)

When the other bidders are submitting bids equal to any fixed fraction of their estimated value, then the highest opposing bid will be this fraction of the highest estimate of any other company, which is computed in cell C11 of Figure 7.14 by the formula

```
=MAX(A3:A5)
```

So to compute the profit for Alpha Petroleum in any simulation, we only need to know this highest estimate of any other company (other than Alpha) and the actual value of the Valdez oil. The simulation table here, for 501 independent simulations, shows the value of the Valdez oil (in B12:B512) and the highest estimates of this value by the three opposing bidders (in C12:C512). The fixed fraction (here, 0.8) of their estimates that each company is supposed to bid has been entered into cell D8 of this spreadsheet. So the highest opposing bid for each simulation is computed by entering the formula

```
=$D$8*C12
```

into cell E12 and copying E12 to cells E12:E512. Cell F8 here represents Alpha Petroleum's bid in this auction, and Alpha will win in those simulations where this bid is greater than the highest opposing bid (as indicated by 1s in cells F12:F512). Alpha's profit in the auction will be the actual value of the oil minus Alpha's bid, if Alpha wins the auction, but will be 0

Myerson_C07.qxd 09-09-2004 10:35 Page 267

if Alpha does not win. So Alpha's profit is computed for each simulation here by entering the formula

=IF($F$8>E12,B12–$F$8,0)

into cell G12 and copying G12 to cells G12:G512. The average of these profits, computed in cell G515 by the formula

=AVERAGE(G12:G512)

is our best estimate of Alpha's expected profit from submitting the F8 bid in this auction, when the other three bidders apply the strategy of bidding the G8 fraction of their respective estimates.

This analysis suggests that if the other companies are bidding 80% of their estimated values, then Alpha Petroleum would be very foolish to bid 80% of its own estimate (0.8∗10 = $8 million). Alpha's expected profit from this bid is negative. Of course, if all companies are bidding the same fraction of their estimates, then Alpha wins only if its geologists have generated the highest estimate, which should happen about 1/4 of the time (as indicated in cell F515). But given the additional information that Alpha's $10 million estimate is higher than any of the other three companies' estimates, the conditional expected value of the Valdez oil would be only about $7.2 million (as indicated in cell H515 of Figure 7.14), which is less than Alpha's bid of $8 million here, and so Alpha has an expected net loss in this scenario. The problem is that Alpha tends to win only when its estimate is too high, because the opponents' bids are also highly correlated with the actual value of the Valdez oil, and so the winner's curse reduces the expected value of winning the auction for Alpha.

Cells G520:G530 of Figure 7.14 are generated by a data table, and they show how Alpha's expected profit would change if Alpha used other bids from 0 to $10 million, instead of the $8 million in cell F8. (The assumption that other companies will bid 80% of their estimates is being held fixed here.) According to this table, Alpha could increase its expected profit to a maximum of $0.033 million by bidding $6 million. This lower bid would have a substantially smaller probability of winning the auction, but at least it would be less than the conditional expected value of the Valdez oil given that this bid wins. So this analysis suggests that Alpha should bid $6 million in this auction.

This analysis in Figure 7.14 was based on the assumption that the three other companies would each submit a bid equal to 80% of its estimated value of the Valdez oil. But we have just seen that bidding 80% of its estimate would be a very foolish decision for Alpha. In view of this fact, the decision makers at Alpha Petroleum may want to reexamine the assumption that their competitors would use such a bidding strategy. If the responsible decision maker at Alpha really believes that Alpha's competitors will use such an 80%-of-estimate bidding strategy, even though it may be foolish for them, then Figure 7.14 can be considered a reasonable analysis of Alpha's optimal bidding problem.

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 10 | Alpha's estimated value of oil in Valdez ($millions) | | | | | | | |
| 2 | 10.11281 | Actual value of oil in Valdez | | | | | | | |
| 3 | 6.531616 | Opponent's Estimate (1) | | | | | | | |
| 4 | 10.24085 | Opponent's Estimate (2) | | | | | | | |
| 5 | 12.56646 | Opponent's Estimate (3) | | | | | | | |
| 6 | | | | | | | | | |
| 7 | Fraction of Estimate Bid by Opponents | | | | | A'sBid | | | |
| 8 | | | | 0.6 | | 6 | | | |
| 9 | | Value | Highest | | | | | | |
| 10 | | of Oil | OppEst | | Highest | | | | |
| 11 | SimTable | 10.11281 | 12.56646 | | OppBid | AWins? | AsProfit | V\|AWins | |
| 12 | 0 | 6.754736 | 8.789973 | | 5.273984 | 1 | 0.754736 | 6.754736 | |
| 13 | 0.002 | 12.80586 | 18.48945 | | 11.09367 | 0 | 0 | .. | |
| 14 | 0.004 | 6.914 | 6.27317 | | 3.763902 | 1 | 0.914 | 6.914 | |
| 15 | 0.006 | 5.941167 | 10.23081 | | 6.138488 | 0 | 0 | .. | |
| 16 | 0.008 | 9.534073 | 12.11897 | | 7.271384 | 0 | 0 | .. | |
| 17 | 0.01 | 5.636246 | 13.2329 | | 7.939741 | 0 | 0 | .. | |
| 507 | 0.99 | 4.61466 | 4.776742 | | 2.866045 | 1 | -1.38534 | 4.61466 | |
| 508 | 0.992 | 13.95113 | 17.55869 | | 10.53522 | 0 | 0 | .. | |
| 509 | 0.994 | 7.454461 | 6.874322 | | 4.124593 | 1 | 1.454461 | 7.454461 | |
| 510 | 0.996 | 6.702069 | 11.31236 | | 6.787413 | 0 | 0 | .. | |
| 511 | 0.998 | 10.46207 | 12.68126 | | 7.608754 | 0 | 0 | .. | |
| 512 | 1 | 10.92589 | 12.84069 | | 7.704415 | 0 | 0 | .. | |
| 513 | | | | | | | | | |
| 514 | FORMULAS | | | | P(A Wins) | | E(A's $) | E(V\|AWins) | |
| 515 | A2. =A1*EXP(NORMINV(RAND(),0,0.35)) | | | | | 0.2415 | 0.291629 | 7.207488 | |
| 516 | A3. =$A$2/EXP(NORMINV(RAND(),0,0.35)) | | | | | | 0.291629 | | |
| 517 | A3 copied to A3:A5 | | | | | | | | |
| 518 | B11. =A2 | | | | | Bid | E(A'sProfit) | | |
| 519 | C11. =MAX(A3:A5) | | | | | (,F8) | 0.291629 | | |
| 520 | E12. =$D$8*C12 | | | | | 0 | 0 | | |
| 521 | F12. =IF($F$8>E12,1,0) | | | | | 1 | 0 | | |
| 522 | G12. =IF($F$8>E12,B12-$F$8,0) | | | | | 2 | 0 | | |
| 523 | H12. =IF($F$8>E12,B12,"..") | | | | | 3 | 0.019781 | | |
| 524 | E12:H12 copied to E12:H512 | | | | | 4 | 0.055824 | | |
| 525 | F515. =AVERAGE(F12:F512) | | | | | 5 | 0.227187 | | |
| 526 | G515. =AVERAGE(G12:G512) | | | | | 6 | 0.291629 | | |
| 527 | H515. =AVERAGE(H12:H512) | | | | | 7 | 0.262394 | | |
| 528 | G516. =F515*(H515-F8) | | | | | 8 | 0.104022 | | |
| 529 | G519. =G515 | | | | | 9 | -0.17803 | | |
| 530 | G520:G530. {=TABLE(,F8)} | | | | | 10 | -0.72942 | | |

FIGURE **7.15**
**Alpha's bidding for Valdez when others bid 60% of their estimated values**

But suppose instead that the responsible decision maker at Alpha is not comfortable with an assumption that his opponents will foolishly overbid. Perhaps he might say, "We should not assume that we are smarter than our opponents. Whatever fraction of our estimate is optimal for us to bid in this auction, we should be expecting our opponents to bid the same fraction of their estimates." To satisfy this requirement of consistent rationality, we need to find a bidding fraction in cell D8 for the other bidders such that using this fraction of 10 (Alpha's estimate) as Alpha's bid in cell F8 would give Alpha a higher expected profit (in cell G515) than any other possible bid in cell F8.

If we assume (for simplicity) that Alpha only considers bidding whole multiples of 0.1 times its estimate, then bidding 60% of estimates yields a consistent rational scenario, as shown in Figure 7.15. Figure 7.15 differs

Myerson_C07.qxd 09-09-2004 10:35 Page 269

from the previous figure in that cell D8 has been changed to 0.6, which expresses an assumption that each of Alpha's opponents will submit a bid equal to 60% of the estimated value. Then the data table in cells G520:G530 in Figure 7.15 shows that the highest expected profit for Alpha ($0.29 million) is achieved by having Alpha bid $6 million, which is also 60% of Alpha's estimate. This optimal bid of $6 million has also been entered into cell G8. So this analysis yields the same recommended bid for Alpha as before ($6 million), but the rational opponents' lower bidding here gives Alpha a higher expected profit ($0.29 million).

In game theory, a *Nash equilibrium* is a scenario in which each player uses a strategy that maximizes the expected payoff, given the strategies that all other players are using in the scenario. In this sense, the scenario in Figure 7.15 might be called a Nash equilibrium of the Valdez bidding game. But this terminology would be technically incorrect, because the model in Figure 7.15 does not quite properly apply the required structure of a game model as it is defined in game theory. The problem is that any proper game model must describe the uncertainty of all players. But the other bidders in this auction do not know Alpha's estimate, and yet this estimate in cell A2 is treated as a known number. In our simulation models, we express uncertainty about a quantity by making it a random variable. So if we want to make a proper game model of the Valdez auction, then we need to make Alpha's estimate into a random variable, just like all the others' estimates. The only information that we can put into our model is that which may be common knowledge among all four bidders.

Such a game model is shown in Figure 7.16. In cell A1 here, the unknown value of oil at Valdez is modeled as a random variable with a probability distribution that represents what all bidders might have believed about the amount of oil based only on the government's public announcement that some oil had been found at this site. To express uncertainty over a wide range of positive values, an Exponential distribution with mean 20 has been assumed here, which is simulated by the formula =EXPOINV(RAND(),20) in cell A2. Then, as in our previous model, each company's geologists are assumed to generate an estimate of this value that is equal to the value divided by an independent Lognormal random variable with log-mean 0 and log-standard-deviation 0.35. (So the ratio of the true value divided by each company's estimate is an independent random variable with this Lognormal distribution.)

So the four bidders are treated symmetrically in this game model. But the following analysis of the model focuses on the decision problem from Alpha's perspective. The simulation table in Figure 7.15 lists data from 501 simulations about the actual value of the oil (in B12:B512), the highest estimate of any company other than Alpha (in C12:C512), and Alpha's estimate (in D12:D512). The fraction of their estimates that Alpha's opponents are assumed to bid is entered into cell F8, and the fraction of its own estimate that Alpha will bid is entered into cell G8. So for the first

| | A | B | C | D | E | F | G | H | I | J |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 11.3597 | Actual value of oil in Valdez | | | | | | | | |
| 2 | 10.0869 | Alpha's Estimate of value of oil in Valdez ($millions) | | | | | | | | |
| 3 | 14.1927 | Opponent's Estimate (1) | | | | | | | | |
| 4 | 17.4051 | Opponent's Estimate (2) | | | | | | P(A wins) | E(A'sProfit) | |
| 5 | 5.50987 | Opponent's Estimate (3) | | | | | | 0.23952 | 0.63266 | |
| 6 | | | | | | | | | | |
| 7 | | Fraction of Estimate Bid by Opponents | | | | | Fractn of Est Bid by A | | | |
| 8 | | | | | | 0.6 | 0.6 | | | |
| 9 | | Value | Highest | | | | | | | |
| 10 | | of Oil | OppEst | A'sEst | | Highest | | | | |
| 11 | SimTabl | 11.3597 | 17.4051 | 10.0869 | | OppBid | A's bid | A wins? | A's Profit | |
| 12 | 0 | 12.6665 | 26.6843 | 12.1377 | | 16.0106 | 7.28263 | 0 | 0 | |
| 13 | 0.002 | 4.37504 | 4.94433 | 2.97999 | | 2.9666 | 1.788 | 0 | 0 | |
| 14 | 0.004 | 11.2059 | 24.5507 | 12.3381 | | 14.7304 | 7.40287 | 0 | 0 | |
| 15 | 0.006 | 12.0287 | 20.227 | 12.1933 | | 12.1362 | 7.31599 | 0 | 0 | |
| 16 | 0.008 | 8.23894 | 14.4699 | 6.89961 | | 8.68193 | 4.13977 | 0 | 0 | |
| 17 | 0.01 | 1.90694 | 2.1656 | 1.64608 | | 1.29936 | 0.98765 | 0 | 0 | |
| 507 | 0.99 | 64.842 | 80.4346 | 61.8151 | | 48.2608 | 37.089 | 0 | 0 | |
| 508 | 0.992 | 22.5812 | 24.74 | 35.1287 | | 14.844 | 21.0772 | 1 | 1.50403 | |
| 509 | 0.994 | 22.8165 | 26.0281 | 13.9273 | | 15.6169 | 8.35638 | 0 | 0 | |
| 510 | 0.996 | 9.35569 | 12.2408 | 13.9542 | | 7.34448 | 8.37252 | 1 | 0.98317 | |
| 511 | 0.998 | 21.491 | 21.0438 | 12.9286 | | 12.6263 | 7.75715 | 0 | 0 | |
| 512 | 1 | 19.9347 | 22.4771 | 22.1839 | | 13.4863 | 13.3103 | 0 | 0 | |
| 513 | | | | | | | | | | |
| 514 | FORMULAS | | | | | | | Bid/Est | E(A'sProfit) | |
| 515 | A1. =EXPOINV(RAND(),20) | | | | | | | (,G8) | 0.63266 | |
| 516 | A2. =$A$1/EXP(NORMINV(RAND(),0,0.35)) | | | | | | | 0 | 0 | |
| 517 | A2 copied to A2:A5 | | | | | | | 0.1 | 0 | |
| 518 | B11. =A1 | | | | | | | 0.2 | 0 | |
| 519 | C11. =MAX(A3:A5) | | | | | | | 0.3 | 0.01679 | |
| 520 | D11. =A2 | | | | | | | 0.4 | 0.27308 | |
| 521 | F12. =C12*$F$8 | | | | | | | 0.5 | 0.52387 | |
| 522 | G12. =D12*$G$8 | | | | | | | 0.6 | 0.63266 | |
| 523 | H12. =IF(G12>F12,1,0) | | | | | | | 0.7 | 0.27441 | |
| 524 | I12. =IF(G12>F12,B12-G12,0) | | | | | | | 0.8 | -0.5741 | |
| 525 | F12:I12 copied to F12:I512 | | | | | | | 0.9 | -1.6127 | |
| 526 | H5. =AVERAGE(H12:H512) | | | | | | | 1 | -2.9499 | |
| 527 | I5. =AVERAGE(I12:I512) | | | | | | | | | |
| 528 | I515. =I5 | | | | | | | | | |
| 529 | I516:I526. {=TABLE(,G8)} | | | | | | | | | |

FIGURE **7.16**
**Game model of bidding for Valdez**

data row (row 12) in the simulation table, Alpha's bid can be computed in cell G12 by the formula

=D12*$G$8

the highest opposing bid can be computed in cell F12 by the formula

=C12*$F$8

and Alpha's profit can be computed in cell I12 by the formula

=IF(G12>F12,B12–G12,0)

Then copying these formulas down 501 rows gives us Alpha's profits from the auction in all stored simulations. The expected profit is estimated in cell I5 by the formula

=AVERAGE(I12:I512)

Myerson_C07.qxd 09-09-2004 10:35 Page 271

A data table shows how this expected profit would change if the fraction that Alpha bids (in cell G8) were changed to bid any multiple of 0.1 between 0 and 1 when the other bidders continue to bid the fraction that is shown in cell F8. In this table we can see that when 0.6 is the fraction that others are bidding (as shown in cell F8), Alpha's highest expected profit ($0.63 million) is achieved when Alpha bids 60% of its estimate. Of course, because the four companies have symmetric positions in this auction, if Alpha wants to bid 60% of its estimate in this scenario, then each of the other three companies should want to bid 60% of their estimates also, and so this scenario is a Nash equilibrium.

This analysis has assumed that each company must choose a strategy of bidding some constant fraction of its geologists' estimate, and we limited consideration to fractions that are multiples of 0.1. A more sophisticated analysis of other nonlinear bidding strategies (where each company's bid is some nonlinear function of its geologists estimate) would be more difficult to do (see Figures 8.3 and 8.7). But even with the restrictive linear assumption of this game model, our analysis here can give realistic insights into the degree to which bidders should discount their bids below their estimates of the common value that the winner will take from the auction.

## 7.7 Summary

This chapter considered problems where the decision maker has to choose a decision variable that can take any positive numerical value. We considered problems where the decision variable is a quantity to produce and decision problems where the decision variable is a price to bid.

We discussed three general methods of using simulations to compare alternative decisions or strategies:

1. making a separate simulation table of profits for each alternative
2. making a simulation table of profits from several alternatives at once
3. making a simulation table of all payoff-relevant random variables that can then be used to compute payoffs for any alternative strategy

The third method is generally best because it allows us to compare all possible alternatives, knowing that the differences we find are due only to different decisions and not to different outcomes of random variables. One of the other two methods may be necessary, however, when the number of payoff-relevant random variables is large.

A strategy is a plan for making decisions now and in the future, where future decisions may depend on the information that will be available to the decision maker. Getting information about a random variable would enable the decision maker to use a strategy that selects different decisions

Myerson_C07.qxd 09-09-2004 10:35 Page 272

with different values of the random variable. For a random variable that has finitely many possible values, we can describe a strategy by a table that lists the random variable's possible values in the left column and the strategy's corresponding decisions in the right column. With such a strategy table in a spreadsheet, the strategy can then be implemented in a simulation analysis by a VLOOKUP formula.

The expected value of information about a random variable is the maximum amount that the decision maker should be willing to pay now (before learning the value of the unknown quantity) to learn the value of this random variable before making decisions. To compute the expected value of information, we must compare the decision maker's expected payoffs from the optimal strategy that uses information about the random variable and from the optimal decision without this information. The optimal strategy with information would select, for each possible value of the random variable, the decision that maximizes the decision maker's conditional expected payoff given this value of the random variable.

Decision-tree diagrams are a useful way of illustrating the sequence of observations and decisions in a decision problem. But decision trees have the disadvantage of forcing us to consider only a small number of alternatives if we want to see the whole tree.

We considered bidding problems where the bidder must choose a bid $\beta$ that will be the price at which the bidder can buy or sell if this bid $\beta$ wins. In the case of bidding to sell, the bid wins if it is less than the lowest opposing bid. We let the random variable $\mathbf{A}$ denote the lowest opposing bid and let $\mathbf{C}$ denote the bidder's actual cost of fulfilling the contracted sale. Then the bidder's profit from the bid $\beta$ is

$$\text{Profit} = \beta - \mathbf{C} \quad \text{if } \beta < \mathbf{A}$$
$$\text{Profit} = 0 \quad \text{if } \beta > \mathbf{A}$$

and so the bidder's expected profit from submitting the bid $\beta$ is

$$E(\text{Profit}) = P(\mathbf{A} > \beta) * (\beta - E(\mathbf{C} | \mathbf{A} > \beta))$$

In the case of bidding to buy, the bid wins if it is greater than the highest opposing bid. We let the random variable $\mathbf{V}$ denote the actual value to the bidder of the object that the bid would buy and let $\mathbf{A}$ denote the highest opposing bid. Then the bidder's profit from the bid $\beta$ is

$$\text{Profit} = \mathbf{V} - \beta \quad \text{if } \beta > \mathbf{A}$$
$$\text{Profit} = 0 \quad \text{if } \beta < \mathbf{A}$$

and so the bidder's expected profit from submitting the bid $\beta$ is

$$E(\text{Profit}) = P(\mathbf{A} < \beta) * (E(\mathbf{V} | \mathbf{A} < \beta) - \beta)$$

Correlation between the opposing bids and the bidder's actual value of buying or cost of selling may cause an effect called the winner's curse. In the case of bidding to buy, there is a winner's curse when $E(\mathbf{V} | \mathbf{A} < \beta)$ is less

Myerson_C07.qxd 09-09-2004 10:35 Page 273

than $E(\mathbf{V})$. In the case of bidding to sell, there is a winner's curse when $E(\mathbf{C}|\mathbf{A} > \beta)$ is more than $E(\mathbf{C})$.

We developed a model to find optimal bids in situations where a winner's curse may exist. In the construction of this model for a case of bidding to sell, the correlation between the bid-to-beat **A** and the decision maker's cost **C** was subjectively assessed by asking the decision maker about how the conditional median for **C** would be affected by information about **A**. A method of modular simulation was used to simulate the bid-to-beat **A** by a random variable that depends on a single CORAND value so that we could easily achieve the assessed correlation between **C** and **A**. A case of bidding to buy could be handled similarly (with the value **V** in place of the cost **C** here).

To model uncertainty about other people's information and behavior, we may treat their observations and decisions as random variables, and we may try to assess probability distributions for these random variables by the usual methods of subjective assessment. But when we hypothesize that a competitor's decisions will depend on observations according to some particular strategy, we should check to see whether the competitor would really want to use such a strategy. So we may try to imagine ourselves in our competitors' positions, but then we may realize that they are busy worrying about what we are likely to do. A game model is an attempt to look at a competitive situation from all sides at once, ours and our competitors'. In such a game model, any quantity that is uncertain to any player must be described by a random variable. So that even quantities that we actually know may become random variables in a game model if our competitors do not know them.

## Exercises

**1.** Our local heroes are going to the Rose Bowl next month, and we have a licence to sell official Rose Bowl sweatshirts. The sweatshirts cost $10 each to produce, and we will sell them before the Rose Bowl at a price of $25. At this price, the anticipated demand for our sweatshirts before the Rose Bowl should be a Normal random variable with mean 9000 and standard deviation 2000. After the Rose Bowl, things depend on whether our local heroes win or not. If they win, then we will continue to sell the sweatshirts for $25, and demand in the month after the game will be a Normal random variable with mean 6000 and standard deviation 2000. If they lose, then we will cut the price to $12.50, and demand at that price should be a Lognormal random variable with mean 2000 and standard deviation 1000. The probability of our local heroes winning the Rose Bowl game is 0.4. Undemanded sweatshirts will be discarded.

**(a)** What production quantity would maximize our expected profit?

**(b)** How would our optimal quantity change if we want to maximize our certainty equivalent when we evaluate risky profits with a risk tolerance of $50,000?

**(c)** Assuming that we want to maximize our expected profit (as in part (a)), what would be the expected value of perfect information about whether our local heroes will win the Rose Bowl?

**2.** For the Scotia problem from Section 7.1, do a sensitivity analysis on the remainder value that was

assessed to be \$8 in the original case. Show how the maximal expected profit and the optimal production decision would change when this remainder value is reassessed to different values ranging from \$0 to \$15.

**3.** A movie studio is about to release a new film. This studio has a regular policy of spending \$5 million on advertising in local media to promote any new film that it releases. But for some films that are considered to have strong potential, the studio may also invest in a special marketing campaign on national television to increase interest in the film. The special marketing campaign on national television would cost an extra \$10 million (thus increasing the cost of promoting the film to a total of \$15 million).

The marketing staff at the studio has tried to forecast this new film's potential revenue. The experts at the studio believe that if the studio follows its regular marketing policy (without any special national television campaign), the studio's revenue from this film has probability 1/2 of being greater than \$18 million, probability 1/4 of being greater than \$30 million, and probability 1/4 of being less than \$12 million. Let us assume that this revenue is drawn from a Generalized Lognormal distribution. The marketing experts also believe that, whatever the revenue would be under the regular marketing policy, having a special marketing campaign on national television would increase revenue by 50%.

A more complicated strategy has been suggested for this film: to first release it only in the New York metropolitan area (with only the regular local advertising), and then make a decision about national marketing based on the results of the first week in New York. In previous films that were distributed with the regular policy, the ratio of the first week's revenue in New York divided by the total national revenue has had mean 0.02 and standard deviation 0.01. So let us assume that the first week's revenue in New York is a random fraction, drawn from a Lognormal distribution with mean 0.02 and standard deviation 0.01, multiplied by the total national revenue that the film would earn with the regular marketing policy.

**(a)** Make a simulation model to analyze the studio's possible profit under these alternative strategies. Then answer parts (b) and (c) based on a table of output from at least 400 simulations of your model.

**(b)** Consider first the two simple alternatives (regardless of the New York results) of (1) distributing the film with regular marketing and (2) distributing the film with a special marketing campaign on national television. For each of these two alternatives, estimate the *mean* and *standard deviation* of the studio's profit from distributing the film, and show the *cumulative risk profile* of the studio's profit from distributing the film. (Define profit as revenue minus marketing costs.)

**(c)** Now consider more complicated strategies of the following form: "first release the film only in the New York metropolitan area, then spend an extra \$10 million for a special marketing campaign on national television if the first week's revenue in New York is greater than $X$, but otherwise (if less than $X$) continue with regular marketing only."

Consider different values of $X$ from \$0 up to \$1 million at intervals of \$0.1 million (that is, consider $X$ = \$0, $X$ = \$0.1 million, $X$ = \$0.2 million, $X$ = \$0.3 million, . . . , $X$ = \$1.0 million). Make a table and a chart showing, for each of these strategies, the *expected* profit and the profit levels that have 10% and 90% cumulative probabilities for the studio from this film. If the studio wants to maximize its expected profit, which of these strategies should it use?

**4.** Suppose that we are negotiating to buy the Lepton Company's rights to a process for making a chemical that has great medical potential which Lepton has recently developed and patented. If the chemical is approved by the FDA as a drug, then its value to us will be a Normal random variable with mean \$120 million and standard deviation \$40 million. If the chemical is not approved by the FDA, then its value to us will be \$0. Our best guess now is that the probability of FDA approval is 2/3.

But experts at Lepton would have much more accurate information about this probability. Let **Q** denote the probability of FDA approval as it is assessed by the experts at Lepton. For us, **Q** is an unknown quantity. Suppose that we have assessed a Beta distribution with mean 2/3 and standard deviation 0.2 for this unknown **Q**.

Lepton's willingness to sell us the patent will depend on its assessment of **Q**. If it thinks that **Q** is 1/2, then we figure that it would be willing to sell for \$40 million; but if it thinks that **Q** is 3/4, it would not sell for less than \$60 million. More generally, we figure that the risk-averse owners of Lepton would be willing to sell us the patent if our bid for the patent is more than **Q**∗\$80 million, but otherwise they would refuse to sell.

Based on at least 2000 simulations of a model of this situation, estimate our expected profit from submitting a bid of $\beta$ for Lepton's patent, for each

number β that is a multiple of $5 million from $5 million to $80 million. Also, for each β, show the probability of β winning and our conditional expected value of the drug given that the bid β would win. To illustrate your results, make a chart that shows how our expected profit depends on this bid β, and make another chart that shows the cumulative risk profile of our net profit from submitting the bid that gives us the highest expected profit. (*Hint:* Use BETINV to simulate a Beta random variable with a given mean and standard deviation.)

**5.** The legal staff for Acme Inc. is deciding how much money Acme should offer now to settle a lawsuit for breach of contract so as to avoid the costs and risks of going to trial. Given their current information, Acme's staff would say now that the probability of the plaintiff winning this suit against Acme is about 0.5. But Acme's lawyers recognize that the plaintiff can assess his probability of winning this lawsuit much more accurately, because the outcome of this trial will depend mostly on the plaintiff's quality as a witness, which he and his lawyer are much better able to predict. From Acme's perspective, the plaintiff's assessed probability of winning is an unknown quantity, and Acme's current beliefs about this unknown probability can be described by a Beta distribution with mean 0.5 and standard deviation 0.18.

If they go to trial and the plaintiff wins the lawsuit, then Acme will pay a $500,000 damage award to the plaintiff. But if they go to trial and the plaintiff loses, then he gets $0 from Acme. In a trial, regardless of who wins, each side will have to pay its own litigation cost. The staff at Acme expect that their own litigation cost in this trial would be $80,000. The staff at Acme do not know what the plaintiff thinks his expected litigation cost would be in a trial, but Acme's beliefs about the plaintiff's expected cost can be described by a Uniform distribution on the interval from $75,000 to $100,000 (independent of his probability of winning). So the plaintiff's net payoff in a trial would be his damage award ($500,000 if he wins, $0 if he loses) minus the plaintiff's litigation cost. But Acme's total cost in a trial would be this same damage award ($500,000 if the plaintiff wins, $0 if he loses) *plus* Acme's $80,000 litigation cost.

The litigation costs and risks of a trial can be avoided if the parties agree to settle out of court now. Both parties are risk-neutral. So the plaintiff will agree to settle out of court only if Acme offers to pay him now an amount that is greater than his expected net payoff in a trial. If Acme's offer is accepted and they settle out of court, then Acme's total cost would just be the amount that it has offered to pay in this settlement.

Suppose that Acme can make just one offer to settle with the plaintiff and it will go to trial if he rejects it. Also, to keep things simple, suppose also that Acme is only considering possible offers that are whole multiples of $25,000 between $0 and $400,000. Make a model of this situation to evaluate different offers that Acme might make, and answer the following questions based on at least 2000 simulations of your model.

**(a)** Make a table that lists, for each amount $*X* that is a multiple of $25,000 between $0 and $400,000, Acme's expected total cost if Acme offers to pay $*X* for settling the lawsuit now.

**(b)** From the table, find the offer amount that is optimal for Acme, in the sense of minimizing Acme's expected total cost. What is the probability of the plaintiff accepting this optimal offer?

**(c)** What would be the conditional probability that the plaintiff would have won the lawsuit if we were given the information that the plaintiff is willing to accept Acme's optimal offer now?

**(d)** What would be the conditional expected value of Acme's total cost in a trial if we were given the information that the plaintiff is *not* willing to accept Acme's optimal offer now?

**(e)** Imagine now that a spy in the plaintiff's law office could tell Acme whether the plaintiff's probability of winning is greater than 1/2 or less than 1/2 (but this spy knows nothing else). Assume (as above) that Acme will only consider making offers that are multiples of $25,000. To minimize Acme's expected cost, how much would Acme offer the plaintiff to settle the lawsuit if the spy reported that the plaintiff's probability of winning is more than 1/2? How much would Acme offer if the spy reported that the plaintiff's probability of winning is less than 1/2? What is the most that Acme should be willing to pay this spy now to get the spy's report?

**6.** Suppose that we are helping Crafty Foods to prepare a bid to buy the Gold Star Cola brand from Aquacom, a diversified conglomerate that has been selling off some of its unrelated divisions. Let **V** denote the net present value of profits that Crafty Foods could get from owning Gold Star Cola. Based on summary sales data supplied by Aquacom and our own expertise in other beverage products, we have assessed a subjective probability distribution for this value **V** that is Normal with mean $40 million and standard deviation of $10 million.

Myerson_C07.qxd 09-09-2004 10:35 Page 276

Let **A** denote the lowest price at which Aquacom would be willing to sell Gold Star Cola. Of course, Aquacom's negotiators have not revealed much information about this reservation price. So we think of Aquacom's reservation price **A** as a Normal random variable, with mean \$30 million and standard deviation \$5 million.

Aquacom's negotiators have suggested that they may be getting another bid for Gold Star Cola. Let **B** denote the unknown value of this other bid. We think that it is possible that Aquacom may be bluffing, without any prospect of another serious prospect. Specifically, we estimate that there is a probability 0.4 that Aquacom will not get any other serious bid, in which case **B** will just be 0. In the event that they actually get another bid, which has probability 0.6, we figure this other bid **B** will be drawn from a Normal distribution with mean \$32 million and standard deviation \$7 million.

Our problem is to recommend how much Crafty Foods should bid for Gold Star Cola. Let *C* denote our bid. If our bid *C* is greater than Aquacom's reservation price **A** and the other bid **B**, then Crafty Foods will pay *C* to buy Gold Star Cola, and the ultimate profit for Crafty Foods from the deal will be the value **V** minus our bid *C*. On the other hand, if the reservation price **A** or the other bid **B** is greater than our bid *C*, then Crafty Foods will not buy Gold Star Cola, and Crafty Foods will get zero profit from the deal.

**(a)** Make a model of this situation, assuming that **V**, **A**, and **B** are independent random variables, and generate data from at least 1000 simulations of this model. Make a table and a chart to show Crafty Foods' *expected profit* from bidding any number *C* between \$20 million and \$40 million that is a whole multiple of \$1 million. For the optimal bid in this range that maximizes expected profit for Crafty Foods, show the *cumulative risk profile* of profit for Crafty Foods, and estimate the *probability* that this optimal bid will win the Gold Star deal for Crafty Foods.

**(b)** Aquacom's reservation price **A** might be based on other information that Aquacom has not shared with us, which might be relevant to the potential value **V** of Gold Star Cola for Crafty Foods. If so, then it may be wrong to assume that **A** and **V** are independent. To take account of this possibility, make a revised model of this situation in which **V** and **A** have correlation 0.5. You should continue to assume that the other bid **B** is independent of **A** and **V**, however. Using data from at least 1000 simulations of your revised model, make a table and a chart to show Crafty Foods' *expected profit* from bidding any number *C* between \$20 million and \$40 million that is a multiple of \$1 million. For the optimal bid in this range that maximizes expected profit for Crafty Foods, show the *cumulative risk profile* of profit for Crafty Foods. What would be the *conditionally expected value* of **V** given the event that this optimal bid wins Gold Star Cola for Crafty Foods?

**7.** Suppose that we are bidding to do a construction job, which is being allocated by a sealed-bid auction. We will win the auction if our bid is lower than any other bidder, in which case our profit will be the amount of our bid minus our cost of doing the job. (If we do not win, then we get zero profit from the auction, of course.)

We have substantial uncertainty about our cost of doing the job. We figure that if the cost is equally likely to be greater or less than \$4.0 million, it has probability 1/4 of being less than \$3.5 million and has probability 3/4 of being less than \$5.0 million. We also have uncertainty about the number of bidders. We feel sure that there will be at least one other bidder for the job, and we figure that the expected number of other bidders is three. So let us suppose that the number of other bidders will be 1 plus a Poisson random variable with mean 2. (*Note:* 1 + POISINV(0.999, 2) = 9, and so the probability of ten or more opposing bidders under this assumption is negligible, being less than 0.001.)

When we think about the bids that our competitors might submit, we figure that any single opposing bid is equally likely to be above or below \$7 million, with a probability 1/4 of being below \$5.5 million and a probability 3/4 of being below \$9 million.

**(a)** Assume first that the cost and all the opposing bids are independent random variables. Based on a simulation analysis, estimate what *optimal bid* would maximize our expected profit from participating in this auction. For this bid, what would be our *expected profit*, and what would be our *probability of winning*? Make a chart showing how our expected profit would vary with our bid for bids from \$4 million to \$18 million.

**(b)** To check for the winner's curse effect, we should think about a possible relationship between the opposing bids and the actual cost of doing the job. Suppose that if we learned that the lowest opposing bid was \$7.5 million, then we would revise our beliefs about the unknown cost of the job and would say that the cost was equally likely to be above or below \$4.75 million (instead

of $4.0 million). Make a new model in which the probability distribution of the lowest opposing bid is the same as in your model for part (a) but the cost of doing the job is now positively correlated with the lowest opposing bid. Based on a large simulation analysis of this revised model, estimate what *optimal bid* would maximize our expected profit from participating in this auction. For this bid, what would be our *expected profit*, and what would be our *probability of winning*? Make a chart showing how our expected profit would vary with our bid for bids from $4 million to $18 million.

**8.** NASA has announced the technical specifications for a new solar explorer satellite. The contract to produce the solar explorer satellite for NASA will be awarded to the lowest bidder in a sealed-bid auction next month. The winner of the auction will then have a contract to produce a solar explorer satellite, with NASA paying the amount of the winning bid on delivery of the completed satellite. Furthermore, after this first solar explorer satellite is delivered and launched into orbit, NASA may request additional solar explorer satellites, and the winner of this auction will have the option to sell these additional satellites to NASA for the same price.

Lockman Aerospace Company has been invited to bid in the auction for the solar explorer satellite, and managers at Lockman have assessed probability distributions for the important unknown quantities that they face in this auction. Managers at Lockman estimate that among the bids submitted by their competitors, the lowest opposing bid is equally likely to be above or below $36 million, has probability 1/4 of being below $30 million, and has probability 1/4 of being above $42 million. Lockman's cost to produce the first solar explorer probe for NASA has expected value $32 million and standard deviation $7 million. To take account of possible winner's curse effects, let us assume that the lowest opposing bid and Lockman's cost to produce the satellite have correlation 0.4.

The number of additional satellites that NASA may request after the first satellite is launched has the following probability distribution (independently of the bids and the cost of production):

| *Number of Additional Satellites* | *Probability* |
|---|---|
| 0 | 0.4 |
| 1 | 0.3 |
| 2 | 0.2 |
| 3 | 0.1 |

From past experience, Lockman managers know that the cost of producing each additional satellite would be 10% less than the cost of producing the first satellite. So in the event that Lockman wins the auction, Lockman would agree to supply additional satellites only if 0.9 times its cost to produce the first satellite is less than its bid.

**(a)** Make a model for analyzing Lockman's bidding problem as a decision under uncertainty, and generate a table of simulation data from at least 2000 simulations of your model to study the probability distribution of the profit that Lockman could get from submitting various bids.

**(b)** Considering bid amounts from $30 million to $50 million in incremental steps of $1 million show for each of these 21 bids that Lockman could submit in the auction

**(1)** Lockman's expected profit from submitting the bid

**(2)** Lockman's certainty equivalent if it has a risk tolerance constant of $15 million

**(3)** the profit level having 10% cumulative probability for Lockman with this bid

**(4)** the profit level having 90% cumulative probability

**(5)** the probability that this bid will win the auction

**(c)** Make a chart of the inverse cumulative distribution of Lockman's profit from submitting the bid that maximizes Lockman's expected profit from part (b).

**9.** Ajax Sprocket Company is building a new factory to go into the widget business. Ajax will make two kinds of widgets: regular and super. The variable cost of labor and materials is $250 for each regular widget, and $350 for each super widget. Up to the levels of its capacity and demand, Ajax will sell regular widgets for $650 each and super widgets for $850 each.

The factory can have three kinds of assembly lines: lines that are designed to make only regular widgets, lines that are designed to make only super widgets, and flexible lines that can produce either kind of widget. The cost of an assembly line depends on capacity, which is measured in the number of widgets that it can produce. An assembly line that makes only regular widgets costs $100 per unit of capacity, a line that makes only super widgets costs $120 per unit of capacity, and a flexible line that can make either kind of widget costs $140 per unit of capacity.

The three types of capacity in the factory must be determined now, before we learn the amounts

Myerson_C07.qxd 09-09-2004 10:35 Page 278

demanded. Demand for each type of Ajax's widgets next summer is an unknown quantity, but the demand for regular widgets has expected value 20,000 and standard deviation 5000, while the demand for super widgets has expected value 30,000 and standard deviation 7000. The number of units that Ajax actually sells will be determined by the minimum of demand and available capacity. Because the profit margin is higher on super widgets, any flexible capacity should first be applied to super demand and then, if there is excess capacity for supers, to regular demand. So total capacity for super widgets is the sum of capacities in the dedicated super assembly lines and the flexible assembly lines. Total capacity for regular widgets will be the capacity in dedicated regular assembly lines plus any excess capacity in flexible assembly lines after meeting demand for super widgets.

**(a)** Suppose that the demands for each kind of widget are independent Normal random variables. To maximize Ajax's expected profit, how much capacity should Ajax build in each of the three kinds of assembly line? What is the probability of having excess capacity to make more super widgets than is demanded? What is the probability of having excess capacity to make more regular widgets than is demanded? What is the probability of having excess capacity to make more of at least one kind of widget?

**(b)** How would your answers to part (a) change if the demands for super and regular widgets are not independent, but instead are Multivariate Normal random variables with correlation –0.7? (Thanks to Jan Van Mieghem for suggesting this problem.)

**10.** *Case:* The Riverfront Development.

Louise Capet had sold insurance and real estate, first in Springfield and then in Shelbyville, for a decade when she bought an extensive undeveloped riverfront property south of Shelbyville. A year later, she got approval from government regulators to develop over 80% of the property, setting aside only a small portion for flood plain and wetlands conservation. At this point, Capet went to see Chuck Martel, a partner in a big Springfield real estate firm. "The approval is good news, but now the project has gotten too big for me," Capet explained. Martel had worked with Capet on some small projects in the past, and now they were discussing how to share the risks of this development.

Capet and Martel agreed that building access roads and other basic infrastructure on the property would cost about $200,000. The revenue that would be earned from selling the subdivided property was harder to estimate. Capet suggested that the revenue from selling the lots should be considered to have expected value $500,000 and standard deviation $100,000.

Martel thought that these revenue estimates would be accurate if this property were near Springfield, but he suspected that Shelbyville values might be somewhat lower. Martel knew that Capet had the best expertise for making this assessment, but Capet also had an incentive to pretend to be more optimistic than she really was. Thus Martel felt that he should not necessarily believe what Capet said about the expected value. Instead, Martel believed that Capet's true expected value of the project's revenue was somewhere between $300,000 and $500,000. To be precise, let us suppose that his beliefs about her true expected value can be described by a Uniform distribution over this interval.

On the other hand, Martel felt that he could trust the $100,000 standard deviation as an accurate assessment of Capet's uncertainty about the project's revenue. With such uncertainty, Capet's risk aversion would give her a good reason to sell some or all of the riverfront development project to Martel's firm. Martel estimated that Capet's risk-tolerance index was somewhere between $40,000 and $60,000. Let us suppose that Martel's beliefs about Capet's risk tolerance can be described by a Uniform random variable on this interval (independently of her beliefs about the expected value of revenue).

If Martel could learn Capet's true assessment of the expected value of the project's revenue, then he would agree that the revenue should be considered a Normal random variable with this mean and standard deviation $100,000. But for risks involving a standard deviation of $100,000, Martel's firm would be essentially risk-neutral and would care only about getting the highest possible expected monetary value. Thus, Martel was optimistic that some mutually beneficial deal could be found in which his firm would buy some or all of the development project from Capet. But he wondered what share they should offer to buy, and at what price. Whatever offer he might make, Capet would accept it only if it increased her certainty equivalent, as she would compute it with her true expected value of the revenue and her true risk tolerance (applying the formula for certainty equivalents of Normal lotteries from Section 4.7).

**(a)** Build a simulation model of this situation from Martel's perspective to analyze his decision problem.

Myerson_C07.qxd 09-09-2004 10:35 Page 279

**(b)** Suppose that Martel made a fixed take-it-or-leave-it offer to buy the entire riverfront development project from Capet now for some fixed price. Considering prices between \$0 and \$250,000 at intervals of \$10,000, estimate the price that would maximize Martel's expected profit from making this offer. What is the probability that Capet will accept this offer?

**(c)** Suppose instead that Martel offered to pay the \$200,000 costs of building roads and other infrastructure in exchange for a 2/3 share of the project's revenues, keeping Capet as a partner with a 1/3 share of the revenues. What is Martel's expected profit from making this offer? What is the probability that Capet will accept it?

Myerson_C07.qxd 09-09-2004 10:35 Page 280

Myerson_C08.qxd 07-09-2004 09:45 Page 281

# 8

# Risk Sharing and Finance

In our study of decision analysis, we initially assumed (starting in Chapter 2) that the basic criterion for defining optimal decisions under uncertainty is maximization of the decision maker's expected payoff. This assumption seemed questionable, however, when we observed that individual decision makers are often risk-averse. Then utility theory (in Chapter 3) taught us to use expected utility values, rather than expected monetary values, as our general criterion for optimal decision making. We have focused on the theory of utility functions with constant risk tolerance as a practical framework for analyzing the effect of risk aversion on people's decisions. But utility theory was still about decision making by individuals. In this chapter we make the transition from individual decision making to decision making in partnerships and corporations.

A large business enterprise generally has a chief executive officer, but it typically has many owners (partners or stockholders), each of whom may have a different risk tolerance. How should such partnerships or corporations make decisions under uncertainty? Should they use the utility function of the chief executive officer, or of the owners? If the owners, how do we resolve their differences when they have different utility functions?

We begin this chapter with a general analysis of optimal risk sharing among individuals who have constant risk tolerance. We find that in an optimal allocation of risks, more risk-tolerant people should hold more risks, in proportion to their risk tolerance. We show that a partnership with such optimal risk sharing should evaluate investment opportunities by applying a risk tolerance that is the sum of the partners' individual risk tolerances.

We then consider the effect of incentive constraints that may prevent managers from achieving such optimal risk sharing with outside investors. The trade-off between risk sharing and incentives is analyzed in simple principal-agent problems. This analysis teaches us that senior managers of big businesses should be expected to bear a personally significant share of the corporate risks. Such managers may then want corporate decision

Myerson_C08.qxd 07-09-2004 09:45 Page 282

making to be guided by their own personal utility functions applied to their personal shares of the corporate risks. But a very different approach is needed if we want to analyze corporate decision making on behalf of the stockholders.

When a publicly held corporation makes decisions on behalf of its stockholders, it should generally assume that its stock is held by investors as a part of a well-diversified portfolio and that these stockholders want the corporation to maximize the value of its shares in the stock market. So to understand optimal corporate decision making for the stockholders, we need a theory of how the prices of financial assets are determined in the stock market. We develop here a model of financial asset pricing, using the assumption that the stock market includes many investors who have constant risk tolerance. This model is somewhat different from the well-known capital asset pricing model (CAPM), but it yields similar and closely related results. Both these asset pricing models teach us that the magnitude of a corporation's risks alone may be less important than the relationship between these risks and the greater aggregate risks of the whole stock market.

Any system of asset pricing that does not create arbitrage opportunities must be consistent with a generalized expected value criterion that applies some modified probabilities which may be determined in the stock market. These general results of arbitrage pricing theory are introduced at the end of this chapter and are shown to include our asset pricing model as a special case.

## 8.1 Optimal Risk Sharing in a Partnership of Individuals with Constant Risk Tolerance

To introduce the basic ideas of optimal risk sharing, let us begin with an example of two individuals (numbered 1 and 2) who are considering a real-estate development project. Suppose that they have an option to buy a tract of land for $125,000, after which they would then need to spend an additional $40,000 on improvements (including an allocation for the cost of their own time in supervising the project) before they could sell the land in subdivided lots. The total revenue that they could then earn from selling these lots would be uncertain but has an expected value of $200,000 and a standard deviation of $25,000. For simplicity, let us assume here that the time to complete this real-estate project is small enough that we can ignore the interest costs of borrowing money to cover the expenses before the revenues come in. So the net returns from this real-estate project next year will have expected value

$$\mu = 200{,}000 - (125{,}000 + 40{,}000) = \$35{,}000$$

Myerson_C08.qxd 07-09-2004 09:45 Page 283

and standard deviation

$$\sigma = \$25{,}000$$

Suppose that each of these two individuals evaluates risky incomes using a utility function with constant risk tolerance, where individual 1 has risk tolerance

$$r_1 = \$20{,}000$$

and individual 2 has risk tolerance

$$r_2 = \$30{,}000$$

They must decide whether to undertake this real-estate project and if so, how to divide the returns between them. Let us assume that the uncertainty about profits from this project can be described by a Normal distribution.

In Section 4.7 we saw that when an individual with constant risk tolerance $r$ has a gamble that will pay a random amount of money drawn from a Normal probability distribution with mean $\mu$ and standard deviation $\sigma$, the certainty equivalent for the gamble is

$$CE = \mu - (0.5/r)*\sigma\text{^}2$$

So if individual 1 were to undertake this project alone, the certainty equivalent would be

$$\begin{aligned}\mu - (0.5/r_1)*(\sigma\text{^}2) &= 35{,}000 - (0.5/20{,}000)*(25{,}000\text{^}2)\\ &= 35{,}000 - 15{,}625 = \$19{,}375\end{aligned}$$

That is, the option to buy this land and undertake this project would be worth \$19,375 to individual 1, if undertaking all the risks of the project alone.

If individual 2 were to undertake this project alone, then its value would be

$$\begin{aligned}\mu - (0.5/r_2)*(\sigma\text{^}2) &= 35000 - (0.5/30{,}000)*(25{,}000\text{^}2)\\ &= 35{,}000 - 10{,}417 = \$24{,}583\end{aligned}$$

So if individual 1 had the option to buy this land, then individual 2 would be willing to pay up to \$24,583 to buy the option, and individual 1 would be glad to sell the option for any price above \$19,375. Of course, it is not surprising that this risky project should be more valuable to the individual who has greater risk tolerance.

But even though individual 2 is strictly more risk-tolerant than individual 1, the project could be even more valuable to these individuals if they undertake the project as partners, with individual 1 taking a positive share of the project's risks. For example, if they each took 50% of the net profits from the project, then each individual would anticipate a payment drawn

Myerson_C08.qxd 07-09-2004 09:45 Page 284

from a Normal distribution with mean 0.50∗35,000 = $17,500 and standard deviation 0.50∗25,000 = $12,500. For the 50% share, individual 1 would have certainty equivalent

$$CE(1) = 17{,}500 - (0.5/20{,}000)*(12{,}500\text{^}2) = 17{,}500 - 3906 = \$13{,}594$$

For a 50% share, individual 2 would have certainty equivalent

$$CE(2) = 17{,}500 - (0.5/30{,}000)*(12{,}500\text{^}2) = 17{,}500 - 2604 = \$14{,}896$$

So the total certainty-equivalent value of the project to the two individuals when they share it equally is

$$CE(1) + CE(2) = 13{,}594 + 14{,}896 = \$28{,}490$$

Thus, the project is worth more to them when it is shared equally than when the more risk-tolerant individual 2 owns it completely.

Such risk sharing is beneficial because each individual *j*'s risk premium $(0.5/r_j)*\sigma\text{^}2$ is proportional to the square of the standard deviation (the variance) of individual income. So halving individual 2's share from 100% to 50% would halve the standard deviation of the monetary returns from $25,000 to $12,500, which in turn would reduce the risk premium to a quarter of its former value, from $10,417 to $2604. This decrease in individual 2's risk premium from giving up 50% of the project (10,417 – 2604 = 7813) is much greater than the increase in individual 1's risk premium when taking on 50% of the project (3906 – 0).

The spreadsheet in Figure 8.1 is set up to analyze the effect on the individuals' certainty of other ways of sharing the risks of this project. When we enter individual 1's share of the risks into cell C7, then the expected value and standard deviation of 1's income are calculated in cells D7 and E7 by the formulas =C7∗$B$2 and =C7∗$B$3, where cells B2 and B3 contain the mean 35,000 and standard deviation 25,000 of the project's total profits. Then individual 1's certainty equivalent is calculated in cell F7 by the formula = D7 – (0.5/B7)∗E7^2, where B7 contains individual 1's risk tolerance 20,000. Individual 2's share is calculated by =1–C7 in cell C8, and copying D7:F7 to D8:F8 yields the certainty equivalent for individual 2's share in cell F8. Cell F11 computes the sum of the individuals' certainty equivalents by the formula =SUM(F7:F8).

Now we can use Solver in this spreadsheet to maximize the sum of the computed certainty equivalents in cell F11 by changing individual 1's percentage share of the project in cell C7. The result is that Solver returns the value 0.4 in cell C7, as shown in Figure 8.1. When individual 1 takes a 40% share, the expected monetary value is 0.40∗35,000 = $14,000 and the standard deviation is 0.40∗25,000 = $10,000, and so the certainty equivalent is

$$CE(1) = 14{,}000 - (0.5/20{,}000)*(10{,}000\text{^}2) = 14{,}000 - 2500 = \$11{,}500$$

Myerson_C08.qxd 07-09-2004 09:45 Page 285

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | Suppose profits will be drawn from a Normal distribution | | | | | | | |
| 2 | Mean | 35000 | | | | | | |
| 3 | Stdev | 25000 | | | | | | |
| 4 | | | | | | | | |
| 5 | Profits | can be shared by individuals 1 and 2 | | | | | | |
| 6 | Individ | RiskTol | %Share | Mean | Stdev | CE | RiskPremium | |
| 7 | 1 | 20000 | 0.4 | 14000 | 10000 | 11500 | 2500 | |
| 8 | 2 | 30000 | 0.6 | 21000 | 15000 | 17250 | 3750 | |
| 9 | | | | | | | | |
| 10 | | Sum(RTs) | CE(total,sumRTs) | | | Sum(CEs) | Sum(RPs) | |
| 11 | | 50000 | 28750 | | | 28750 | 6250 | |
| 12 | SOLVER(1): Maximize F11 | | | by changing C7. | | | | |
| 13 | FORMULAS | | | | | | | |
| 14 | C8. =1-C7 | | | | B33. =NORMINV(RAND(),B2,B3) | | | |
| 15 | D7. =C7*$B$2 | | | | D27. =1-C27 | | D28. =-C28 | |
| 16 | E7. =C7*$B$3 | | | | C34. =C$28+C$27*$B34 | | | |
| 17 | F7. =D7-(0.5/B7)*E7^2 | | | | D34. =D$28+D$27*$B34 | | | |
| 18 | G7. =D7-F7 | | | | C34:D34 copied to C34:D534 | | | |
| 19 | D7:G7 copied to D8:G8 | | | | C31. =CE(C34:C534,C30) | | | |
| 20 | F11. =SUM(F7:F8) | | | | D31. =CE(D34:D534,D30) | | | |
| 21 | G11. =SUM(G7:G8) | | | | F31. =SUM(C31:D31) | | | |
| 22 | B11. =SUM(B7:B8) | | | | F30. =SUM(C30:D30) | | | |
| 23 | C11. =B2-(0.5/B11)*B3^2 | | | | F33. =CE(B34:B534,F30) | | | |
| 24 | | | | | F28. =C28+D31 | | | |
| 25 | SOLVER(2): Maximize F31 | | | by changing C27:C28. | | | | |
| 26 | Individual | | 1 | 2 | | | | |
| 27 | Sharing rate | | 0.4 | 0.6 | | C28 value for CE2=0 | | |
| 28 | Fixed payment | | 17000 | -17000 | | 17475 | | |
| 29 | | | | | | | | |
| 30 | RiskTols | | 20000 | 30000 | | 50000 | Sum(RiskTols) | |
| 31 | CE | | 28650 | 475 | | 29125 | Sum(CEs) | |
| 32 | (sim'd) Total$ | | Net incomes | | | | | |
| 33 | SimTable | 72184.7 | Pay 1 | Pay 2 | | 29125 | CE(total$,sumRTs) | |
| 34 | 0 | 59273 | 40709 | 18564 | | | | |
| 35 | 0.002 | 396 | 17158 | -16763 | | | | |
| 36 | 0.004 | 34138 | 30655 | 3483 | | | | |
| 37 | 0.006 | 48928 | 36571 | 12357 | | | | |
| 38 | 0.008 | 5586 | 19235 | -13648 | | | | |

FIGURE **8.1**
**Sharing a Normal gamble**

When individual 2 takes a 60% share, the expected monetary value is 0.60∗35,000 = $21,000 and the standard deviation is 0.60∗25,000 = $15,000, and so the certainty equivalent is

$$CE(2) = 21{,}000 - (0.5/30{,}000)*(15{,}000\text{^}2) = 21{,}000 - 3750 = \$17{,}250$$

When they plan to share the risks in this way, their total certainty equivalent of the project is

$$CE(1) + CE(2) = 11{,}500 + 17{,}250 = \$28{,}750$$

This total of $28,750 is the maximal sum of certainty equivalents that the partners can achieve by sharing the profits of this project.

In this optimal sharing rule, the ratio of 2's share to 1's share is 0.6/0.4 = 1.5. Notice that the ratio of 2's risk tolerance to 1's risk tolerance is exactly the same 30,000/20,000 = 1.5. This result is not a coincidence, as the following general fact asserts.

Myerson_C08.qxd 07-09-2004 09:45 Page 286

**Fact 1** Suppose that a group of individuals have formed a partnership to share the risky profits from some joint venture or gamble, and each individual $j$ in this group has a constant risk tolerance that we may denote by $r_j$. Let $R$ denote the sum of all the partners' risk tolerances ($R = \sum_j r_j$). Then these individuals can maximize the sum of their certainty equivalents by sharing the risky profits among themselves in proportion to their risk tolerances, with each individual $j$ taking the fractional share $r_j/R$ of the risky profits.

For this example, Fact 1 yields the same optimal shares that Solver returned in Figure 8.1. The sum of the partners' risk tolerances here is

$$R = r_1 + r_2 = 20{,}000 + 30{,}000 = \$50{,}000$$

So the optimal share for individual 1 is 20,000/50,000 = 0.4, the same share that Solver generated in cell C7.

For any such partnership, we may define the *total risk tolerance* of a partnership to be the sum of the risk tolerances of the individual partners. For this example, we have seen that the partnership's total risk tolerance is $R$ = \$50,000. Now, if we considered the partnership as a corporate person with constant risk tolerance equal to this total $R$, then a Normal lottery with mean \$35,000 and standard deviation \$25,000 would have certainty equivalent

$$\mu - (0.5/R)*\sigma\text{^}2 = 35{,}000 - (0.5/50{,}000)*(25{,}000\text{^}2) = \$28{,}750$$

for this partnership, as is calculated in cell C11 of Figure 8.1. Notice that this corporate certainty equivalent is exactly the same as the maximized sum of the partners' individual certainty equivalents in cell F11 under the optimal sharing rule. The following general fact asserts that this result is also not a coincidence.

**Fact 2** Consider a group of individuals who have formed a partnership to share the risky profits from some joint venture or gamble, where each individual has constant risk tolerance, as assumed in Fact 1. Let $R$ denote the sum of all the partners' individual risk tolerances ($R = \sum_j r_j$). Then the maximal sum of the partners' certainty equivalents that can be achieved by optimal risk sharing (as described in Fact 1) is equal to the certainty equivalent of the whole gamble to an individual who has a constant risk tolerance equal to the sum of these partners' risk tolerances. Thus, to maximize the sum of their certainty equivalents, the partnership should evaluate gambles according to its total risk tolerance whenever the partners have a choice about which gambles to undertake.

Facts 1 and 2 here do not require the gamble to be Normal. In illustrating these two facts, we have used the special formula for certainty equivalents of Normal gambles, but the same results can also be obtained with simulation analysis, as shown in the lower portion of Figure 8.1 (row 25 and below). A simulation table here holds a sample of 501 independent simulations of the Normally distributed profits for this project. We consider sharing rules where each partner gets a fractional share of the profits as

Myerson_C08.qxd 07-09-2004 09:45 Page 287

listed in cell C27 or D27 plus a fixed payment listed in cell C28 or D28 (for 1 or 2, respectively). A negative payment in D27 represents a payment from individual 2 to individual 1, as when individual 2 must buy into a project that was initially owned by individual 1. The fractional shares in C27:D27 must sum to 1 because the partners must share 100% of the profits, and the fixed payments in C28:D28 must sum to 0 because any fixed payment to one partner must come from the other. These constraints are represented in this spreadsheet by the formulas =1–C27 in cell D27 and =–C28 in cell D28. Under the sharing rule in C27:D28, the net incomes from the project's simulated profits for individuals 1 and 2 are listed below in cells C34:C534 and D34:D534, and the corresponding certainty equivalents are computed in cells C31 and D31 with the formulas

=CE(C34:C534,C30) and =CE(D34:D534,D30)

(where C30 and D30 contain the individuals' risk tolerances). The sum of the individuals' certainty equivalents is computed in cell F31.

If we ask Solver to maximize the sum of the individuals' certainty equivalents in cell F31 of Figure 8.1 by changing the sharing-rule parameters in cells C27:C28, then Solver will report that individual 1 should keep a 40% share of the profits (as shown in cell C27) and individual 2 should take the remaining 60% (D27), as Fact 1 predicts. Solver will leave 1's fixed payment in C28 at any arbitrary value because changing it would not affect the sum of the certainty equivalents in G31. (In making Figure 8.1, 17,000 was arbitrarily entered into C28 before running Solver, and Solver left it unchanged.)

Cell F33 in Figure 8.1 calculates the certainty equivalent of the total profits from this project, based on the simulation data in B34:B534, by the formula

=CE(B34:B534,F30)

where F30 contains the sum of the partners' risk tolerances $R$ = \$50,000. The value in cell F33 (\$29,125) is exactly the same as the maximized sum of the partners' individual certainty equivalents in cell F31, as Fact 2 predicts. Cells F33 and F31, being estimates from simulation data that only approximates the given Normal distribution, are slightly different from the values in cells F11 and C11, which use the exact formula for certainty equivalents of Normal gambles. But these simulation estimates also confirm Facts 1 and 2, because these facts do not depend on Normality.

Fact 2 can give us some sense of why businesses are typically more risk-tolerant than individuals, because the risks of a business may be shared among many investors. When shares of a company are owned by 50 people whose average risk tolerance is \$20,000, then Fact 2 asserts that the company itself should evaluate risks with a risk tolerance of \$1,000,000. Fact 1 tells us that among these 50 people, the ones with greater risk tolerance should have a greater share of the company.

This discussion assumes that partners should want to maximize the sum of their certainty equivalents. This is a good assumption, but it needs some

Myerson_C08.qxd 07-09-2004 09:45 Page 288

defense. After all, any single partner may care only about his own certainty equivalent of what he gets from the partnership. Why should anyone care about maximizing this sum of all certainty equivalents? The answer is given by the following fact.

**Fact 3** Consider a risk-sharing partnership where all partners have constant risk tolerance. If the partners were planning to share risks according to a sharing rule that does not maximize the sum of the partners' certainty equivalents, then any partner $j$ could propose another sharing rule that would increase $j$'s own certainty equivalent and would not decrease the certainty equivalents of any other partners.

To understand Fact 3, notice first that adding any fixed payment from one partner to another partner would not change the sum of the partners' certainty equivalents. A net payment of $x$ dollars from partner 2 to partner 1 (when there is no uncertainty about this amount $x$) would decrease 2's certainty equivalent by $x$ and would increase 1's certainty equivalent by $x$, because each partner is assumed to have constant risk tolerance. Thus the net payment of $x$ dollars would leave the sum of their certainty equivalents unchanged.

Now, suppose the partners were originally planning to use some sharing rule that does not maximize the sum of the partners' certainty equivalents. Then consider any other sharing rule that is optimal, in the sense of maximizing the sum of the partners' certainty equivalents. Changing to this "optimal" sharing rule would increase some partners' certainty equivalents, but it might also decrease other partners' certainty equivalents. But let us now modify this optimal rule by adding some net payments that will cancel out these changes for all partners except one, say partner $j$. Any partner whose certainty equivalent would decrease should receive an additional payment equal to the amount of the decrease, to be paid by this partner $j$. Any other partner whose certainty equivalent would increase should make an additional payment equal to the amount of the increase, paying it to partner $j$. So when these payments have been added into the optimal sharing rule, everybody other than partner $j$ is getting exactly the same overall certainty equivalent as under the original plan. But adding these fixed payments does not change the sum of the partners' certainty equivalents. So our modified optimal plan (with the additional payments) still maximizes the sum of the partners' certainty equivalents, and so it must generate a strictly greater sum of certainty equivalents than the original plan. Thus, with everybody else's certainty equivalent unchanged, partner $j$ must be enjoying a strictly greater certainty equivalent under this new plan. This proves Fact 3.

Fact 3 tells us that it is always optimal for partners to maximize the sum of their certainty equivalents. To apply Fact 3, consider our sharing example from the perspective of individual 1 in a situation where the option to buy and develop the land could originally be done alone, and so individual 1

Myerson_C08.qxd 07-09-2004 09:45 Page 289

has the option to undertake the project without any participation from individual 2. Individual 2, of course, has the alternative of not participating in the project, in which case the return would be \$0. Any sharing rule that gives 2 a certainty equivalent more than \$0 would be better than nonparticipation, and so could be accepted. The best possible sharing rule for individual 1 would be one that maximizes 1's certainty equivalent subject to the constraint that 2's certainty equivalent should not be less than \$0. Fact 3 tells us that this can be achieved by sharing in the optimal proportions, to maximize the sum of the individuals' certainty equivalents, with an additional payment from individual 2 to individual 1 that reduces 2's certainty equivalent to \$0 (or to some value slightly greater than \$0). By Fact 1, the optimal share for individual 2 is 60% of this project, because 30,000/(20,000 + 30,000) = 0.6, and we have seen that a 60% share with no additional payment would have a certainty equivalent \$17,250 to individual 2 (see cell F8 in Figure 8.1). So the best possible sharing rule for individual 1 would be to sell individual 2 a 60% share of this project for an initial payment of \$17,250 (or slightly less than this), which just exhausts 2's perceived gains from participating in the partnership. After selling 60% of the investment to individual 2 for this maximal price, individual 1 would have \$17,250 in cash plus a risky investment that is worth \$11,500 (the certainty equivalent for a 40% share). Thus, selling 60% to individual 2 for \$17,250 would make individual 1's overall certainty equivalent from the project 17,250 + 11,500 = \$28,750. This is the most 1 could possibly hope for in any sharing rule because it allocates to 1 all the maximal sum of certainty equivalents that the two partners can get from this project.

Of course, individual 2 would prefer to pay less than \$17,250 for a 60% share and might try to negotiate for a lower price in this situation. Recall that \$19,375 was 1's certainty equivalent for undertaking the project alone, and so 1 would not accept any certainty equivalent less than \$19,375 when the alternative is owning 100% of the project. Because the certainty equivalent for 40% of the project is \$11,500, 1 needs an additional payment of 19,375 – 11,500 = \$7875 to raise the certainty equivalent to this level. So the best possible sharing rule for individual 2 here would be to buy 60% of the project (2's optimal share) for just a bit above \$7875, which is the lowest price that individual 1 would be willing to accept.

But regardless of who initially owns the project, the partners can agree that they should maximize the sum of their certainty equivalents by sharing the risky returns in proportion to their risk tolerances. How this maximal value is divided between them is a bargaining problem. If one of them initially owns more than his or her optimal share of the project, there will exist a range of transfer prices at which the individuals could both gain by changing to their optimal shares. In this situation, the price that individual 2 may actually pay to buy 60% of the project must be a question of bargaining between the two individuals, and without a theory of bargaining

Myerson_C08.qxd 07-09-2004 09:45 Page 290

| | A | B | C | D | E | F | G | H | I | J |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Sharing profits drawn from a Gen-Lognormal distribution | | | | | | | | | |
| 2 | with quartiles: | | | 50 | 80 | 120 | ($1000s) | | | |
| 3 | Total $profits (sim'd) | | | | | | | | | |
| 4 | | 53 | | | | | | | | |
| 5 | | | Partners | | | | | | | |
| 6 | | | 1 | 2 | 3 | | Sums | | | |
| 7 | Sharing rates | | 0.3 | 0.3 | 0.4 | | 1 | Sum(Rates) | | |
| 8 | Fixed payment | | 0 | 0 | 0 | | 0 | Sum(Payments) | | |
| 9 | | | | | | | | | | |
| 10 | RiskTols | | 30 | 30 | 40 | | 100 | Sum(RiskTols) | | |
| 11 | | CE | 24.959 | 24.959 | 33.279 | | 83.198 | Sum(CEs) | | |
| 12 | | Total$ | | | | | | | | |
| 13 | SimTabl | 53.341 | Partners' incomes | | | | 83.198 | CE(total$s,sumRTs) | | |
| 14 | 0 | 66.709 | 20.013 | 20.013 | 26.683 | | | | | |
| 15 | 0.002 | 70.729 | 21.219 | 21.219 | 28.291 | | | | | |
| 16 | 0.004 | 19.302 | 5.7906 | 5.7906 | 7.7208 | FORMULAS | | | | |
| 17 | 0.006 | 97.585 | 29.276 | 29.276 | 39.034 | B4. =GENLINV(RAND(),D2,E2,F2) | | | | |
| 18 | 0.008 | 68.36 | 20.508 | 20.508 | 27.344 | B13. =B4 | | | | |
| 19 | 0.01 | 141.21 | 42.363 | 42.363 | 56.484 | C7. =1-SUM(D7:E7) | | | | |
| 20 | 0.012 | 86.192 | 25.858 | 25.858 | 34.477 | C8. =-SUM(D8:E8) | | | | |
| 21 | 0.014 | 84.812 | 25.444 | 25.444 | 33.925 | C14. =C$8+C$7*$B14 | | | | |
| 22 | 0.016 | 105.66 | 31.697 | 31.697 | 42.262 | C14 copied to C14:E514. | | | | |
| 23 | 0.018 | 137.95 | 41.385 | 41.385 | 55.18 | C11. =CE(C14:C514,C10) | | | | |
| 24 | 0.02 | 136.17 | 40.851 | 40.851 | 54.468 | C11 copied to C11:E11. | | | | |
| 25 | 0.022 | 87.173 | 26.152 | 26.152 | 34.869 | G7. =SUM(C7:E7) | | | | |
| 26 | 0.024 | 52.619 | 15.786 | 15.786 | 21.048 | G7 copied to G8,G10:G11. | | | | |
| 27 | 0.026 | 155.76 | 46.729 | 46.729 | 62.305 | G13. =CE(B14:B514,G10) | | | | |
| 28 | 0.028 | 55.537 | 16.661 | 16.661 | 22.215 | SOLVER: maximize G11 | | | | |
| 29 | 0.03 | 193.94 | 58.183 | 58.183 | 77.577 | by changing D7:E8 | | | | |

FIGURE **8.2**
**Optimal risk sharing among three partners with constant risk tolerance**

we can only say here that it should be somewhere between $7875 and $17,250.

Facts 1, 2, and 3 here require the assumption that all partners have constant risk tolerance, but they do not require any assumption about the probability distribution from which the partnership's profits will be drawn. Normality here was only used to compute exact certainty equivalents in the top 11 rows of Figure 8.1.

Figure 8.2 shows an example of optimal linear sharing among three partners where the partnership's profits are generated by a Generalized Lognormal distribution that is not Normal. A linear sharing rule for partners 2 and 3 is parameterized in cells D7:E8 here, with 1's share being determined in cells C7 and C8 so as to keep the sum of shares equal to 100% and the sum of the fixed payments equal to $0. When Solver is asked to adjust these sharing rules so as to maximize the sum of the partners' certainty equivalents in cell G11, then Solvers' optimal solution should give the partners fractional shares in C7:E7 that are proportional to their respective risk tolerances in cells C10:E10, as predicted by Fact 1. Then as predicted by Fact 2, the maximized sum of the three partners' certainty equivalents in cell G11 is equal to the partnerships' certainty equivalent for the total risky profit, which is computed in cell G13 using a risk tolerance for the partnership that is the sum of the partners' individual risk tolerances.

Myerson_C08.qxd 07-09-2004 09:45 Page 291

## 8.2 Optimality of Linear Rules in the Larger Class of Nonlinear Sharing Rules

In Figures 8.1 and 8.2, when we asked Solver to find an optimal sharing rule, we implicitly assumed that the two partners would share the profits linearly. Here a partner's share is *linear* if each extra dollar of profit would increase the partner's income by the same amount. But this linearity assumption is not necessary. Even when we allow that a partner's income may be a nonlinear function of the total profit earned, the linear sharing rule that we described in Fact 1 is still optimal for maximizing the sum of the certainty equivalents among partners who all have constant risk tolerance. (For an illustration of a nonlinear sharing rule, see the inset graph in Figure 8.7.)

To show that nonlinear sharing rules cannot do better for the partners in this example is more complicated. Even with profits coming from a Normal random variable, the individuals' incomes will not be Normal with nonlinear sharing rules, and so we cannot use the simple quadratic formula to compute exact certainty equivalents. So we must use a simulation model. Furthermore, searching among nonlinear sharing rules is much harder, because there are so many nonlinear rules. But Figure 8.3 shows a spreadsheet in which we can evaluate a large set of nonlinear sharing rules and show the optimality of the (40%, 60%) linear sharing rule in this set.

Data from 501 simulations of the partnership's total profit are contained in cells B22:B522 of Figure 8.3. (Rows 25 to 519 have been hidden in Figure 8.3, using the menu command Data:Group.) The smallest simulated profit (–40,750) is shown in cell B5, and a value slightly above the largest simulated profit (104,222) is shown in cell B15. Cells B6:B14 have been filled with an increasing sequence of values that were chosen (somewhat arbitrarily) between these smallest and largest profits (0, 10,000, . . . , 80,000).

We will consider continuous sharing rules such that 1's income depends linearly on profit inside the interval between each pair of adjacent values in cells B5:B15 of Figure 8.3, but a different linear function may be used in each of these intervals. That is, 1's income could be specified by one linear formula for profits in the interval from –40,750 to 0, by another linear formula for profits in the interval from 0 to 10,000, and so on. For continuity, we require that the linear formulas on the intervals below and above 0 must specify the same income when profit is $0, with a similar requirement at each of the other interval separators in B6:B14. Such rules are called *piecewise linear*. By using more small intervals, we could closely approximate any nonlinear sharing rule by such piecewise-linear rules. (For a picture of a piecewise linear function for a similar example, see the inset graph in Figure 8.7.)

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | Sharing profits drawn from a Normal distribution | | | | | | | | |
| 2 | with Mean | | 35000 | and Stdev | | 25000 | | | |
| 3 | | | | | | | | | |
| 4 | | Total | PayTo1 | Slope1 | Intercept1 | | | | |
| 5 | | -40750 | 700 | 0.4 | 17000 | | | | |
| 6 | | 0 | 17000 | 0.4 | 17000 | | | | |
| 7 | | 10000 | 21000 | 0.4 | 17000 | | | | |
| 8 | | 20000 | 25000 | 0.4 | 17000 | | | | |
| 9 | | 30000 | 29000 | 0.4 | 17000 | | | | |
| 10 | | 40000 | 33000 | 0.4 | 17000 | | | | |
| 11 | | 50000 | 37000 | 0.4 | 17000 | | | | |
| 12 | | 60000 | 41000 | 0.4 | 17000 | | | | |
| 13 | | 70000 | 45000 | 0.4 | 17000 | | | | |
| 14 | | 80000 | 49000 | 0.4 | 17000 | | | | |
| 15 | | 104222 | 58689 | | | | | | |
| 16 | | | | | SOLVER: Max F19 by changing C5:C15 | | | | |
| 17 | | | Partner1 | Partner2 | | | | | |
| 18 | RiskTols | | 20000 | 30000 | | 50000 | Sum(RiskTols) | | |
| 19 | | CE | 28746 | 618.46 | | 29364 | Sum(CEs) | | |
| 20 | (sim'd) | Total$ | | | | | | | |
| 21 | SimTabl | 24223 | PayTo1 | PayTo2 | | CE(total$s,sumRTs) | | | |
| 22 | 0 | 68354 | 44342 | 24012 | | 29364 | | | |
| 23 | 0.002 | 18508 | 24403 | -5895 | | | | | |
| 24 | 0.004 | 11732 | 21693 | -9961 | | | | | |
| 520 | 0.996 | 34555 | 30822 | 3733 | | | | | |
| 521 | 0.998 | 24301 | 26720 | -2420 | | | | | |
| 522 | 1 | 46912 | 35765 | 11147 | | | | | |
| 523 | | | | | | | | | |
| 524 | FORMULAS | | | | | | | | |
| 525 | B21. =NORMINV(RAND(),C2,F2) | | | | | | | | |
| 526 | B5. =MIN(B22:B522) | | | | | | | | |
| 527 | B15. =MAX(B22:B522)+1 | | | | | | | | |
| 528 | D5. =(C6-C5)/(B6-B5) | | | | D5 copied to D5:D14 | | | | |
| 529 | E5. =C5-D5*B5 | | | | E5 copied to E5:E14 | | | | |
| 530 | C22. =VLOOKUP(B22,$B$5:$E$14,4)+B22*VLOOKUP(B22,$B$5:$E$14,3) | | | | | | | | |
| 531 | D22. =B22-C22 | | | | | | | | |
| 532 | C22:D22 copied to C22:C522 | | | | | | | | |
| 533 | C19. =CE(C22:C522,C18) | | | | D19. =CE(D22:D522,D18) | | | | |
| 534 | F18. =SUM(C18:D18) | | | | F19. =SUM(C19:D19) | | | | |
| 535 | F22. =CE(B22:B522,F18) | | | | | | | | |

FIGURE **8.3**
**A spreadsheet to evaluate nonlinear sharing rules**

Such a piecewise-linear sharing rule can be specified in Figure 8.3 by listing the income that individual 1 would get for each of the profit values listed in cells B5:B15. In this spreadsheet, these incomes for 1 are listed in cells C5:C15. That is, each cell in C5:C15 specifies the income that individual 1 would get, under this profit-sharing rule, if the profit were equal to the corresponding value in B5:B15. For profits in the interval between any adjacent pair of values in B5:B15, we will determine 1's income by linear interpolation, that is, by applying the linear function that matches the specified income for 1 at each of the two endpoints of the interval. So for any profit $x$ in the interval between B5 and B6, 1's income is supposed to be a linear function of $x$ that has the form $A*x + B$, where the slope $A$ is computed in cell D5 by the formula

=(C6–C5)/(B6–B5)

Myerson_C08.qxd 07-09-2004 09:45 Page 293

and the intercept B is computed in cell E5 by the formula

=C5–D5*B5

Cells D5:E5 have been copied down the range D5:E14 to show the corresponding slope and intercepts for the linear sharing rule that is applied in the interval between each value in B5:B14 and the next value below it.

Now to apply this piecewise-linear sharing rule to the simulated profits in cell B22 of Figure 8.3, cell C22 contains the formula

=VLOOKUP(B22,$B$5:$E$14,4) +B22*VLOOKUP(B22,$B$5:$E$14,3)

The first VLOOKUP in this formula finds the lowest row in the range B5:E15, where the B cell's value is not greater than B22, and then returns the intercept listed in the E column (the fourth column of B5:E14) in that row. The second VLOOKUP in this formula returns the corresponding slope in the third column of the B5:E14 table, which is then multiplied by B22 and added to the intercept. Cell D22 computes the corresponding income for individual 2 from the B22 profit value, by the formula

=B22–C22

Then copying C22:D22 down to C22:D522, we get the incomes for the two partners when the piecewise-linear sharing rule from B5:E14 is applied to the simulated profits in B22:B522. The corresponding certainty equivalents for individuals 1 and 2 are computed in cells C19 and D19, using the CE function, and the sum of these certainty equivalents is computed in cell F19.

Now we can ask Solver to maximize cell F19 by changing cells C5:C15 in Figure 8.3, which specify 1's income values in this piecewise-linear function. Having so many cells to adjust makes this a hard problem for Solver, and it may need to work through 40 trial solutions, which could take an hour of computing time on an older computer from the 1990s but which can be done in less than a minute on newer machines in 2004. The result that Solver returns, as shown in Figure 8.3, has the same linear sharing rule applied in all intervals, and the slope of this rule in cells D5:D14 always gives individual 1 the optimal share of risky profits as specified by Fact 1. (If you happen to specify values in C5:C15 that depend on the B5:B15 profits according to a linear formula that has the optimal slope 0.4, then Solver will quickly report that these initial values constitute an optimal solution and will leave them unchanged. The intercept that Solver returns in cells E5:E15 may be any number and will depend on the initial values that you specified in cells C5:C15 before applying Solver.)

Cell F22 in Figure 8.3 applies the CE function to estimate the value of the total profits of the project (as sampled in B22:B522) to the partnership, when the partnership is treated as a corporate person with a constant risk tolerance equal to the sum of the partners' individual risk tolerances ($50,000, computed in cell F18). Fact 2 tells us that the optimal sum of certainty equivalents in cell F19, after it has been maximized by Solver, must be equal to this value in cell F22, and this equality can be seen in Figure 8.3.

Myerson_C08.qxd 07-09-2004 09:45 Page 294

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | | | Partner 1 | Partner 2 | | | |
| 2 | Risk Tolerance | | 20000 | 30000 | | | |
| 3 | | | | | | | |
| 4 | POSSIBLE OUTCOMES | | | | | | |
| 5 | Total $ | Proby | PayTo1 | PayTo2 | Rate1 | | |
| 6 | 0 | 0.2 | 17857 | -17857 | 0.4 | | |
| 7 | 25000 | 0.3 | 27858 | -2858 | 0.4 | | |
| 8 | 50000 | 0.4 | 37857 | 12143 | 0.4 | | |
| 9 | 75000 | 0.1 | 47859 | 27141 | | | |
| 10 | | | | | | | |
| 11 | Sum(RTs) | | CE(1) | CE(2) | | | |
| 12 | 50000 | | 29763 | 0 | | | |
| 13 | CE(total$,sumRTs) | | | Sum of CEs | | | |
| 14 | 29763 | | | 29763 | | | |
| 15 | | | | | | | |
| 16 | SOLVER 1 (no moral hazard): | | | | | | |
| 17 | Max C12 by changing C6:C9 subject to D12>=0. | | | | | | |
| 18 | | | | | | | |
| 19 | | | | | | | |
| 20 | FORMULAS | | | | | | |
| 21 | D6. =A6-C6 | | | D6 copied to D6:D9 | | | |
| 22 | E6. =(C7-C6)/(A7-A6) | | | E6 copied to E6:E8 | | | |
| 23 | C12. =CEPR(C6:C9,$B$6:$B$9,C2) | | | | | | |
| 24 | D12. =CEPR(D6:D9,$B$6:$B$9,D2) | | | | | | |
| 25 | D14. =SUM(C12:D12) | | | | | | |
| 26 | A12. =SUM(C2:D2) | | | A14. =CEPR(A6:A9,B6:B9,A12) | | | |

FIGURE **8.4**
**Optimal risk sharing in a discrete example**

Finally, let us consider in Figure 8.4 a discrete example where two partners with constant risk tolerances of $20,000 and $30,000 have to share a gamble that will pay an amount of money drawn from the following discrete distribution:

| *Partnership's Total Profit* | *Probability* |
|---|---|
| $0 | 0.2 |
| $25,000 | 0.3 |
| $50,000 | 0.4 |
| $75,000 | 0.1 |

These profit values and probabilities are shown in cells A6:A9 and B6:B9 of Figure 8.4. Cells C6:C9 are used to specify how much income individual 1 should get from the partnership for each possible amount of profit. The corresponding net incomes for individual 2 are computed in cells D6:D9 by entering the formula =A6 – C6 in cell D6 and then copying D6 to D6:D9. The resulting certainty equivalent for individual 1 is computed in cell C12 by the formula

=CEPR(C6:C9,$B$6:$B$9,C2)

where cell C2 contains 1's risk tolerance, $20,000. (Recall from Section 3.1 that CEPR(*values*, *probabilities*, *risktolerance*) returns the certainty equivalent of a discrete gamble, where the given values have the given probabilities, for an individual with the given constant risk tolerance.) Copying C12 to

Myerson_C08.qxd 07-09-2004 09:45 Page 295

D12 yields individual 2's certainty equivalent of income from this discrete gamble. The CEPR function is also applied in cell A14 to compute the certainty-equivalent value of the whole gamble to an individual whose risk tolerance is the sum of these partners' risk tolerances.

In Figure 8.4, Solver has been asked to maximize 1's certainty equivalent in cell C12 by changing the sharing-rule parameters in cells C6:C9, subject to the constraint that 2's certainty equivalent in cell D12 must satisfy D12 >= 0. So cells C6:D9 here show the best possible sharing rule for individual 1, subject to the constraint that individual 2 should be willing to stay in the partnership, when 2's alternative is to leave the partnership and get nothing ($0) from this gamble. In this optimal sharing rule, individual 1 gets a fixed payment of $17,857 from individual 2 (1's income in C6 when the gamble pays $0), and then individual 1 gets $0.40 of each dollar that is earned from the gamble (computed in cells E6:E8). Individual 2 gets the remaining $0.60 of each dollar earned from the gamble, and this 60% share is just worth the fixed payment of $17,857 to 2. Thus, the partners' optimal shares are linear in profits and are proportional to their risk tolerances, as predicted by Fact 1. Also, as predicted by Fact 2, the optimal sum of the partners' certainty equivalents (in cell D14) is equal to the certainty equivalent of the whole gamble to an individual with the total risk tolerance of the partners (in cell A14).

## 8.3 Risk Sharing Subject to Moral-Hazard Incentive Constraints

In real life, people do not always share every risk in proportion to their individual risk tolerances. One basic reason is that people who are well insured against risks sometimes do not work hard enough to avoid them. This problem is called *moral hazard* in the insurance industry. To avoid such moral-hazard problems, workers and managers in an enterprise are often forced to bear more of the enterprise's risks than they would bear under an ideal risk-sharing system.

To introduce the ideas of moral hazard, let us reconsider in Figure 8.5 an extension of the discrete example from Figure 8.4 in the previous section. In this example, individuals 1 and 2, who have constant risk tolerances of $20,000 and $30,000, respectively, are sharing a gamble that will pay a total dollar amount to be drawn from the probability distribution shown in cells A5:B9. Cells C5:D9 in Figure 8.4 show the sharing rule that maximizes 1's certainty equivalent, subject to the constraint that 2's certainty equivalent should be at least $0. This sharing rule is equivalent to 1 selling a 60% share of the gamble to 2 for $17,857.

But suppose now that the distribution of returns listed in cells A5:B9 can be achieved only if individual 1 attends to some managerial duties that

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | | | Partner 1 | Partner | 2 | | MORAL-HAZARD INCENTIVES | | |
| 2 | Risk Tolerance | | 20000 | 30000 | | | 1's private$ if negligent | | |
| 3 | | | | | | | 6000 | | |
| 4 | POSSIBLE OUTCOMES | | | | | | Probys if 1 is negligent | | |
| 5 | Total $ | Proby | PayTo1 | PayTo2 | Rate1 | | Total $ | Proby | |
| 6 | 0 | 0.2 | 6823 | -6823 | 0.95 | | 0 | 0.33333 | |
| 7 | 25000 | 0.3 | 30449 | -5449 | 0.54 | | 25000 | 0.33333 | |
| 8 | 50000 | 0.4 | 43956 | 6044 | 0.67 | | 50000 | 0.33333 | |
| 9 | 75000 | 0.1 | 60739 | 14261 | | | 75000 | 0 | |
| 10 | | | | | | | | | |
| 11 | Sum(RTs) | | CE(1) | CE(2) | | | If 1 neglects his duties | | |
| 12 | 50000 | | 27184 | 0 | | | his share has CE | | |
| 13 | CE(total$,sumRTs) | | | Sum of CEs | | | 21184 | | |
| 14 | 29763 | | | 27184 | | | CE(1) including private $ | | |
| 15 | | | | | | | 27184 | | |
| 16 | SOLVER 1 (no moral hazard): | | | | | | | | |
| 17 | Max C12 by changing C6:C9 subject to D12>=0. | | | | | | | | |
| 18 | SOLVER 2 (moral hazard): | | | | | | | | |
| 19 | Max C12 by changing C6:C9 subject to D12>=0, C12>=G15. | | | | | | | | |
| 20 | FORMULAS | | | | | | | | |
| 21 | D6. =A6-C6 | | | D6 copied to D6:D9 | | | | | |
| 22 | E6. =(C7-C6)/(A7-A6) | | | E6 copied to E6:E8 | | | | | |
| 23 | C12. =CEPR(C6:C9,$B$6:$B$9,C2) | | | | | | | | |
| 24 | D12. =CEPR(D6:D9,$B$6:$B$9,D2) | | | | | | | | |
| 25 | D14. =SUM(C12:D12) | | | | | | | | |
| 26 | A12. =SUM(C2:D2) | | | A14. =CEPR(A6:A9,B6:B9,A12) | | | | | |
| 27 | G13. =CEPR(C6:C9,H6:H9,C2) | | | | | | | | |
| 28 | G15. =G13+G3 | | | | | | | | |

FIGURE **8.5**
**Optimal risk sharing with moral hazard, in a discrete example**

individual 2 cannot directly observe. Suppose that if individual 1 neglected these duties, then there would be no chance of the partnership earning $75,000 and the profit would instead be either $0 or $25,000 or $50,000, each with probability 1/3, as shown in cells G5:H9 of Figure 8.5. So 1's neglect of duties would reduce the partnership's expected profit from $35,000 to $25,000 (computable here by SUMPRODUCT(A6:A9,B6:B9) and SUMPRODUCT(G6:G9,H6:H9), respectively). But suppose that this neglect of duties would enable individual 1 to take up another private project that would be worth $6000.

Under the sharing rule that was shown in Figure 8.4, if individual 1 neglected these duties, then 1's income would be either $17,857 or $27,857 or $37,857, each with probability 1/3, and this gamble would have a certainty equivalent of $26,224 to 1 (given the constant risk tolerance of $20,000). So when the additional $6000 that could be earned privately is taken into account, neglecting the duties to the partnership would enable individual 1 to get an overall certainty-equivalent value of $32,224, which is better than the certainty-equivalent value of $29,763 that 1 would get by properly fulfilling duties to the partnership (shown in cell C12 of Figure 8.4). Thus, under the sharing rule that is shown in Figure 8.4, individual 1 would prefer to neglect these duties. But if 1 neglects them, then individual 2 should not be willing to pay $17,857 for a 60% share of the profits!

Myerson_C08.qxd 07-09-2004 09:45 Page 297

To find a sharing rule that avoids this difficulty, we must add a constraint that individual 1 should not prefer to neglect the duties. Such a constraint may be called a *moral-hazard incentive constraint*, because it says that the risk sharing should not insure so well that 1 does not want to exert appropriate efforts to avoid bad outcomes. This moral-hazard incentive constraint can be expressed in Figure 8.5 by the inequality C12 >= G15, where C12 is the certainty equivalent for 1's share of the partnership when the duties are fulfilled (applying the probabilities in B6:B9) and G15 is 1's certainty equivalent for private income (in G3) plus the share of the partnership when the duties are neglected (applying the probabilities in H6:H9).

Figure 8.5 shows the results when Solver is asked to maximize 1's certainty equivalent in cell C12 by changing the sharing-rule parameters in cells C6:C9, subject to the constraints D12 >= 0 and C12 >= G15. The first constraint here says that individual 2 should not prefer to quit the partnership, and the second constraint says that individual 1 should not prefer to neglect the duties. Under the optimal sharing rule in cells C6:D9 of Figure 8.5, if profit is $0, then individual 1 gets a payment of $6823 from individual 2, as shown in cell C6 of Figure 8.5. This payment is smaller than the corresponding payment in cell C6 of Figure 8.4, but now individual 1 keeps more than half of the partnership's risky profits. As shown in cells E6:E8 of Figure 8.5, individual 1 gets $0.95 from each dollar of the partnership's profit between $0 and $25,000, $0.54 from each dollar of profit between $25,000 and $50,000, and $0.67 from each dollar of profit between $50,000 and $75,000. So individual 1 here holds a larger share of the risks than the ideal share (40%) that Fact 1 would predict, because the Facts in the preceding section assumed that there were no moral-hazard incentive constraints.

Now let us consider a moral-hazard incentive problem in a more realistic example where the profit that an investment may earn is a continuous random variable that has infinitely many possible values. Suppose that a large group of investors is hiring an agent to manage some investment for them. The investors will not be able to monitor the manager directly to see whether he is working or shirking, but they will observe the profit that the investment earns under his management. If the manager works diligently, then this profit will be a Generalized Lognormal random variable with quartile points $230,000, $280,000, and $340,000. But if the manager shirks his responsibilities, then the profit will instead be a Generalized Lognormal random variable with quartile points $190,000, $230,000, and $280,000. Shirking his responsibilities would allow the manager to attend to some personal affairs that would generate private rewards worth $10,000 to him. The manager has constant risk tolerance $20,000. The investors who are hiring this manager have total risk tolerance $480,000. (You may think of these investors as a group of 24 partners, each of whom also has constant risk tolerance $20,000.) The manager's alternative employment opportunities would pay him $50,000 during the period when he is being

Myerson_C08.qxd 07-09-2004 09:45 Page 298

asked to manage this investment, so his certainty equivalent when he agrees to manage this investment cannot be lower than $50,000. Furthermore, no matter how badly the investment turns out, the investors cannot subject the manager to a penalty worse than a $0 wage.

**NOTE** (This last minimum-wage condition may deserve some explanation. It says that the investors cannot ask their agent to put money in an escrow account that he would forfeit in the event of a particularly bad investment performance. Such punitive conditions may be unacceptable because the agent might simply not have the money to put in such an account. Or they may be unacceptable because if the investors requested such punitive conditions in bad events, then the agent would become suspicious that the investors might actually know something bad about this investment project which would increase the probability of the events where they would profit at his expense.)

Among the feasible compensation plans that would give the manager an incentive to work diligently on managing this investment, let us try to find the plan that would yield the highest sum of certainty equivalents for these investors. This is a difficult optimization problem. So we begin by considering a simpler class of compensation plans in Figure 8.6, and then we can go on to consider a more general class of compensation plans later in Figure 8.7.

The given parameters of the optimal compensation problem are listed in the range A1:C15 of Figure 8.6. The profit quartiles if the manager works diligently are in B3:B5, the profit quartiles if the manager shirks are in C3:C5, the risk tolerances of the manager and the investors are in B9 and C9, the manager's certainty equivalent under his best alternative employment option is in B11, the manager's private gain from shirking is in B13, and the required minimum wage is in B15. (All monetary values in Figure 8.6 are in $1000s.)

Cells B20:B520 contain 501 simulated values of the profit returned by this investment when the manager works diligently, and cells C20:C520 contain the corresponding simulated profit values when the manager shirks. This simulation data was generated by recalculations of the random variables in cells B19 and C19, where we have assumed that the working and shirking profits are maximally correlated by having the same RAND( ) (in cell A18) drive both of these random variables.

**NOTE** (Because working and shirking are alternatives that cannot both happen at once, it would not have been wrong to simulate the working and shirking profits by independent random variables. But using correlated random variables here improves the expected accuracy of our estimates from any limited number of simulations, because it reduces the probability that the no-shirking constraint may be distorted by a false contrast between unusually high simulated profits from one alternative and unusually low simulated profits from the other alternative.)

| | A | B | C | D | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | | Profit quartiles | | | | | | | | | |
| 2 | | Work | Shirk | | | | | | | | |
| 3 | Q1 | 230 | 190 | | | | | | | | |
| 4 | Q2 | 280 | 230 | | | | | | | | |
| 5 | Q3 | 340 | 280 | | | | | | | | |
| 6 | | (in $1000s) | | | | | | | | | |
| 7 | Risk tolerances | | | | | | | | | | |
| 8 | | Manager | Investors | | SOLVER: Max G17 by changing E12:E13 | | | | | | |
| 9 | | 20 | 480 | | subject to E17>=B11, E17>=F17. | | | | | | |
| 10 | Mgr's alternative CE | | | | | | | | | | |
| 11 | | 50 | | | Mgr's compensation plan | | | | | | |
| 12 | Mgr's extra $ if shirks | | | | 0.2658 | Slope | | | | | |
| 13 | | 10 | | | -16.71 | Intercept | | | | | |
| 14 | Mgr's required minimum | | | | (Minimum is paid when profit < | | | | | 62.843 | ) |
| 15 | | 0 | | | Mgr's CE | | Investors' CE | | | | |
| 16 | | | | | Work | Shirk | with work | | | | |
| 17 | (rand) | Sim'd profit | | | 50 | 50 | 226.05 | | | | |
| 18 | 0.5636 | Work | Shirk | | Mgr's income | | Investors' income | | | | |
| 19 | SimTabl | 293.28 | 240.89 | | Work | Shirk | with work | | | | |
| 20 | 0 | 423.17 | 352.41 | | 95.79 | 76.981 | 327.38 | | | | |
| 21 | 0.002 | 324.7 | 267.06 | | 69.615 | 54.291 | 255.09 | | | | |
| 22 | 0.004 | 312.85 | 257.13 | | 66.464 | 51.649 | 246.39 | | | | |
| 518 | 0.996 | 282.43 | 231.99 | | 58.377 | 44.966 | 224.06 | | | | |
| 519 | 0.998 | 227.04 | 187.68 | | 43.65 | 33.188 | 183.39 | | | | |
| 520 | 1 | 363.59 | 300.19 | | 79.951 | 63.098 | 283.64 | | | | |
| 521 | | | | | | | | | | | |
| 522 | FORMULAS | | | | | | | | | | |
| 523 | J14. | =(B15-E13)/E12 | | | | | | | | | |
| 524 | A18. | =RAND() | | | | | | | | | |
| 525 | B19. | =GENLINV($A$18,B3,B4,B5) | | | | | | | | | |
| 526 | C19. | =GENLINV($A$18,C3,C4,C5) | | | | | | | | | |
| 527 | E20. | =MAX($E$13+$E$12*B20,$B$15) | | | | | | | | | |
| 528 | F20. | =MAX($E$13+$E$12*C20,$B$15) | | | | | | | | | |
| 529 | G20. | =B20-E20 | | | | | | | | | |
| 530 | E20:G20 copied to E20:G520 | | | | | | | | | | |
| 531 | E17. | =CE(E20:E520,$B$9) | | | | | | | | | |
| 532 | F17. | =CE(F20:F520,$B$9)+B13 | | | | | | | | | |
| 533 | G17. | =CE(G20:G520,C9) | | | | | | | | | |
| 534 | SOLVER: Max G17 by changing E12:E13 | | | | | | | | | | |
| 535 | subject to E17>=B11, E17>=F17. | | | | | | | | | | |

FIGURE **8.6**
**Optimal linear incentive plan for an agent with moral hazard**

In Figure 8.6, we consider compensation plans where the manager is paid a linear function of the total profit that is returned by the investment, except that the manager can never be paid less than the required minimum base. The slope and intercept of this linear function are entered into cells E12 and E13, and the required minimum base ($0) has been specified in cell B15. (We will ask Solver to choose E12 and E13 optimally, so we may start by putting any arbitrary values in E12 and E13.) Then for our simulated profit data, the manager's corresponding wages when he works diligently can be computed in cells E20:E520 by entering the formula

=MAX($E$13+$E$12∗B20, $B$15)

into cell E20 and copying E20 to E20:E520. The manager's simulated wages from shirking are similarly calculated in by entering the formula

=MAX($E$13+$E$12∗C20, $B$15)

Myerson_C08.qxd 07-09-2004 09:45 Page 300

into cell F20 and copying F20 to F20:F520. The remaining profits that will be paid to the investors, in the case where the manager works diligently, are computed by entering the formula

=B20–E20

into cell G20 and copying G20 to G20:G520. Then the manager's certainty equivalent from working can be estimated in cell E17 by the formula

=CE(E20:E520,$B$15)

Then the manager's certainty equivalent from shirking (including the private rewards listed in cell B13) can be estimated in cell F17 by the formula

=CE(F20:F520,$B$15)+B13

By Fact 2 (which can be applied to the investors because they have no moral-hazard incentive constraints among themselves), the sum of the investors' certainty equivalents can be computed in cell G17 by the formula

=CE(G20:G520,C9)

In Figure 8.6, the investors' optimal linear compensation plan has been found by asking Solver to maximize the investors' total certainty equivalent in cell G17 by changing the compensation parameters in cells E12:E13, subject to the constraints E17 >= B11 and E17 >= F17. The first constraint here says that the manager should not prefer his alternative employment opportunities to working for these investors. The second constraint says that the manager should not prefer shirking to working. The results in E12:E13 tell us that the manager should be paid 26.58% of the profits, minus a constant deduction of $16,710, except that he would get the $0 minimum wage if this linear formula yielded an amount less than $0. Cell J14 computes that this minimum wage would apply if profit were less than $62,843, and so in effect the manager is getting a 26.58% share of all profits above $62,843. When the manager works under this optimal linear compensation plan, the investors' total certainty equivalent is $226,050.

Fact 1 does not apply here because of the moral-hazard problem. To see why, notice first that Fact 1 would suggest that the manager's share of profits be 20,000/(20,000 + 480,000) = 0.04, because the manager's risk tolerance is $20,000 and the investors' total risk tolerance is $480,000. With this simulation data, the manager's certainty equivalent from working here could be made equal to the competitive $50,000 certainty equivalent by offering the manager 4% of the profits plus a fixed payment of $38,662, and the resulting total certainty equivalent for the investors would then be increased to $233,450. That is, if we entered 0.04 and 38.662 into cells E12 and E13 of Figure 8.6, then we would get the values 50 in cell E17 and 233.45 in cell G17. But we would also get the value 58.101 in cell F17, which would mean that the manager could increase his own certainty equivalent from $50,000 to $58,101 by shirking. If the manager shirked,

Myerson_C08.qxd 07-09-2004 09:45 Page 301

then the investors would get only the shirking profits (simulated in C20:C520) minus the corresponding payments to the manager (in F20:F520), which would actually yield a certainty equivalent of less than $190,000 for the investors. Thus, because of the moral-hazard problem, Solver has recommended instead that the manager should be paid 26.58% of all profits above $62,843 to make sure that he has an incentive to work diligently.

This optimal linear incentive plan is relatively insensitive to different assumptions about the investors' total risk tolerance. If the investors were risk-neutral, then the investors' total certainty-equivalent formula in cell G17 in Figure 8.6 would be changed to

```
=AVERAGE(G20:G520)
```

If Solver were asked to change E12:E13 so as to maximize this new formula in cell G17, with the same constraints as before, the optimal linear incentive plan would actually remain the same.

**NOTE ON SOLVER** Numerical difficulties may cause Solver to stall at a false solution when different changing cells take values that have very different magnitudes. Such difficulties could occur in an example like Figure 8.6 when the optimal intercept in cell E13 is many times larger than the slope in cell E12. If you find that Solver seems unable to change such cells in a sensible way, then you should try Solver's automatic-scaling option, which is accessed by clicking the Options button in the Solver Parameters dialog box, and then selecting the "Use Automatic Scaling" option.

## 8.4 Piecewise-Linear Sharing Rules with Moral Hazard

Figure 8.7 reconsiders the same problem as Figure 8.6, but it allows nonlinear compensation plans that are piecewise linear on the intervals between adjacent profits in cells E3:E13. The basic parameters of the problem, as listed in the range A1:C15, are the same in Figure 8.7 as in Figure 8.6. But in the range E1:H13 in Figure 8.7, we can construct a nonlinear sharing rule similarly to Figure 8.3.

Cell E3 in Figure 8.7 contains the lowest profit level that is observed under working or shirking in our simulation data in B20:C520, and cell E13 contains a value above the highest profit level in this simulation data. Cells E4:E12 contain a sequence of selected profit values that break the interval between E3 and E13 into 10 similar-sized subintervals. A compensation plan can be specified in cells F3:F13 by listing here the manager's wage for each of the corresponding profit values in E3:E13. (We will ask Solver to choose F3:F13 optimally, so we could start by entering any arbitrary initial values in F3:F13.) For profits between any adjacent pair in E3:E13, the

Myerson_C08.qxd 07-09-2004 09:45 Page 302

| | A | B | C | D | E | F | G | H | I | J |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Profit quartiles | | | | Mgr's compensation plan | | | | | |
| 2 | | Work | Shirk | | Profit | PayMgr | Slope | Intercept | | |
| 3 | Q1 | 230 | 190 | | 114.58 | 0 | 0.5675 | -65.02 | | |
| 4 | Q2 | 280 | 230 | | 150 | 20.102 | 0.5384 | -60.66 | | |
| 5 | Q3 | 340 | 280 | | 200 | 47.021 | 0.1138 | 24.253 | | |
| 6 | | (in $1000s) | | | 250 | 52.714 | 0.1763 | 8.6494 | | |
| 7 | Risk tolerances | | | | 300 | 61.526 | 0.0153 | 56.931 | | |
| 8 | | Manager | Investors | | 350 | 62.292 | 0.0437 | 47 | | |
| 9 | | 20 | 480 | | 400 | 64.477 | 0.1045 | 22.662 | | |
| 10 | Mgr's alternative CE | | | | 450 | 69.704 | 0 | 69.704 | | |
| 11 | | 50 | | | 500 | 69.704 | 0.0678 | 35.787 | | |
| 12 | Mgr's extra $ if shirks | | | | 550 | 73.095 | 0.0345 | 54.131 | | |
| 13 | | 10 | | | 685.35 | 77.762 | | | | |
| 14 | Mgr's required minimum | | | | | | | | | |
| 15 | | 0 | | | Mgr's CE | | Investors' CE | | | |
| 16 | | | | | Work | Shirk | with work | | | |
| 17 | (rand) | Sim'd profit | | | 50 | 50 | 229.99 | | | |
| 18 | 0.7439 | Work | Shirk | | Mgr's income | | Investors' income | | | |
| 19 | SimTabl | 338.15 | 278.43 | | Work | Shirk | with work | | | |
| 20 | 0 | 423.17 | 352.41 | | 66.898 | 62.398 | 356.27 | | | |
| 21 | 0.002 | 324.7 | 267.06 | | 61.905 | 55.721 | 262.8 | | | |
| 22 | 0.004 | 312.85 | 257.13 | | 61.723 | 53.969 | 251.13 | | | |
| 518 | 0.996 | 282.43 | 231.99 | | 58.43 | 50.663 | 224 | | | |
| 519 | 0.998 | 227.04 | 187.68 | | 50.099 | 40.39 | 176.94 | | | |
| 520 | 1 | 363.59 | 300.19 | | 62.886 | 61.529 | 300.7 | | | |
| 521 | | | | | | | | | | |
| 522 | | | | | | | | | | |
| 523 | | | | | | | | | | |
| 524 | | | | | | | | | | |
| 525 | FORMULAS | | | | | | | | | |
| 526 | A18. =RAND() | | | | | | | | | |
| 527 | B19. =GENLINV($A$18,B3,B4,B5) | | | | | | | | | |
| 528 | C19. =GENLINV($A$18,C3,C4,C5) | | | | | | | | | |
| 529 | E3. =MIN(B20:C520) | | | | | | | | | |
| 530 | E13. =MAX(B20:C520)+1 | | | | | | | | | |
| 531 | G3. =(F4-F3)/(E4-E3) | | | | | | | | | |
| 532 | H3. =F3-G3*E3 | | | | | | | | | |
| 533 | G3:H3 copied to G3:H12 | | | | | | | | | |
| 534 | E20. =VLOOKUP(B20,$E$3:$H$12,4)+B20*VLOOKUP(B20,$E$3:$H$12,3) | | | | | | | | | |
| 535 | F20. =VLOOKUP(C20,$E$3:$H$12,4)+C20*VLOOKUP(C20,$E$3:$H$12,3) | | | | | | | | | |
| 536 | G20. =B20-E20 | | | | E20:G20 copied to E20:G520 | | | | | |
| 537 | E17. =CE(E20:E520,$B$9) | | | | F17. =CE(F20:F520,$B$9)+B13 | | | | | |
| 538 | G17. =CE(G20:G520,C9) | | | | | | | | | |
| 539 | SOLVER: Max G17 by changing F3:F13 subject to | | | | | | | | | |
| 540 | F3>=B15, G3:G12>=0, G3:G12<=1, E17>=F17, E17>=B11. | | | | | | | | | |

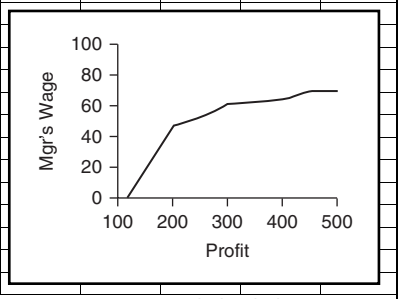


FIGURE **8.7**
**Optimal piecewise-linear incentive plan for an agent with moral hazard**

manager's wage is computed by a linear interpolation between the corresponding pair of wages in F3:F13. So for profits between E3 and E4, the manager's wage is to be determined by a linear function with the slope and intercept computed in cells G3 and H3, respectively, by the formulas

=(F4–F3)/(E4–E3) and =F3–G3∗E3

Then copying G3 and H3 to G3:G12 and H3:H12 yields the slopes and intercepts that are applied in each of the 10 subintervals. Thus, cells E3:H12 form a table that lists the low ends of our 10 subintervals in the first column (E) and that lists the slope and intercepts of the piecewise-linear wage function for each subinterval in the third and fourth columns (G and H).

Myerson_C08.qxd 07-09-2004 09:45 Page 303

So the manager's wage for the simulated profit in cell B20 can be computed in cell E20 of Figure 8.7 by the formula

=VLOOKUP(B20,$E$3:$H$12,4)+B20*VLOOKUP(B20,$E$3:$H$12,3)

Copying cell E20 to E20:F520, we get the manager's simulated wages from working in cells E20:E250, and we get his simulated wages from shirking in cells F20:F520. Entering the formula

=B20–E20

in cell G20, and copying G20 to G20:G520, we get in cells G20:G520 the profits that are paid to the investors when the manager works diligently. Then the certainty equivalents that the manager would get from working and from shirking are computed in cells E17 and F17, respectively, by the formulas

=CE(E20:E520,$B$9) and =CE(F20:F520,$B$9)+$B$13

The investors' total certainty equivalent is computed in cell G17 by the formula

=CE(G20:G520,C9)

In Figure 8.7, the investors' optimal piecewise-linear compensation plan has been found by asking Solver to maximize the investors' total certainty equivalent in cell G17 by changing the compensation parameters in cells F3:F13, subject to the constraints F3 >= B15, G3:G12 >= 0, G3:G12 <= 1, E17 >= F17, and E17 >= B11. The constraint G3:G12 >= 0 asserts that in each subinterval, the slope of the wage function must not be negative. That is, the manager's wage must never be decreased by an increase in the investment's total profits, because otherwise the manager could have an incentive to incur unnecessary costs simply to reduce profits where this reduction would increase his wage. The constraint G3:G12 <= 1 says that the manager must never gain more than $1 for an additional $1 of profit, because otherwise the manager could have an incentive to artificially inflate profits by covertly adding some of his own money. The constraint F3 >= B15 guarantees that the manager's wage is never below the required minimum wage of $0 in cell B15, because F3 is the manager's wage at the lowest profit and the nonnegative slopes in G3:G12 imply that the wages cannot be lower at higher profits. The constraint E17 >= B11 says that the manager must not prefer his alternative employment opportunities (which pay the amount $50,000 shown in B11) over working to manage this investment. The constraint E17 >= F17 says that the manager must not prefer shirking to working here.

With so many variables to adjust and so many constraints to satisfy, this Solver problem can take a long time on older computers, but newer computers can do all this in a few minutes. The optimal plan, as shown in Figure 8.7, pays the minimum wage at the lowest possible profit level in our simulation data, but then increases the manager's wage with a share of profits that

Myerson_C08.qxd 07-09-2004 09:45 Page 304

starts above 50% and declines gradually as profit increases. For profits above $300,000, this Solver solution gives the manager a share of incremental profits that varies somewhat erratically but that averages around 4% over the interval from $300,000 to $685,000. (The wrinkles in the compensation curve for profits above $300,000 could be caused by small random clusters in our simulated profit data, which Solver may exploit by wrinkling the compensation curve to pay less at the shirking-profit clusters and more at the working-profit clusters. Such effects would tend to disappear if we used a larger simulation sample or wider subintervals.) The investors' total certainty equivalent with this optimal piecewise-linear incentive plan is $229,990, shown in cell G17, which is about $4000 more than the certainty equivalent that they got from their best linear incentive plan in Figure 8.6. (*Note:* Figure 8.7 actually uses the same simulation data in B20:C520 as Figure 8.6, so this increase is really due to the change of allowing some nonlinearity in the manager's compensation plan.)

If the investors were risk-neutral, or if there were so many investors that the sum of their risk tolerances was too large to measure, then the investors' total certainty equivalent would be their expected share of profits, which could be estimated in cell G17 by the formula

```
=AVERAGE(G20:G520)
```

If we ask Solver to maximize this measure of the investors' expected profit by changing the same variables with the same constraints as we used in Figure 8.7 (using the same simulation data as used in Figures 8.6 and 8.7), then the optimal compensation plan (in $1000s) would become:

| *Profit* | *Pay Manager* | *Slope* | *Intercept* |
|---|---|---|---|
| 114.58 | 0 | 0.474 | –54.323 |
| 150 | 16.795 | 0.711 | –89.821 |
| 200 | 52.334 | 0.04 | 44.396 |
| 250 | 54.319 | 0.133 | 20.963 |
| 300 | 60.990 | 0 | 60.990 |
| 350 | 60.990 | 0 | 60.990 |
| 400 | 60.990 | 0.046 | 42.446 |
| 450 | 63.308 | 0 | 63.308 |
| 500 | 63.308 | 0 | 63.308 |
| 550 | 63.308 | 0.003 | 61.931 |
| 685.35 | 63.646 | | |

The main difference between this compensation plan and the one in Figure 8.7 is that the manager here gets almost no wage increases for profit increases above $300,000. So the positive wage increases for profit increases above $300,000 that we see in cells G7:G12 of Figure 8.7 are due mainly to the advantages of adding the manager as another partner in risk sharing. Most of the required incentives for working instead of shirking are generated in Figure 8.7 by the big pay increases that the manager can get from increasing profits up to $300,000.

Myerson_C08.qxd 07-09-2004 09:45 Page 305

The optimality of the nonlinear incentive plan in Figure 8.7 may depend, however, on our implicit assumption that the manager must decide once whether to work or shirk, and only thereafter can he learn how much profit will be earned. But in a real world that is more complex than our simple model, the profits from such an investment project might actually be earned over a period of weeks or months, and the manager might be able to decide about working or shirking every day, knowing how much profit has already been earned so far. If our optimal incentive plan from Figure 8.7 were applied in such a world, then the manager might work diligently only until the accumulated profits reach $300,000, and thereafter he might start shirking, which would significantly reduce the investors' chances of earning any higher profits above $300,000. To avoid this dynamic moral-hazard problem, we may recommend that the investors should restrict their attention to linear compensation plans, as assumed in Figure 8.6, because the incentives to continue working under linear compensation plans would be less affected by good news in the early weeks of the project.

## 8.5 Corporate Decision Making and Asset Pricing in the Stock Market

The preceding sections can provide the basis for understanding optimal decision making under uncertainty in big publicly held corporations. Actually, we should have two theories of decision analysis in corporations, depending on whether the decision analysis is supposed to be for the benefit of senior managers who control the corporation or for the benefit of the stockholders who legally own the corporation. The latter assumption is the main focus of this section, but let us first devote one paragraph to the former assumption.

Under a well-designed compensation plan, as we saw in Sections 8.3 and 8.4, senior managers may anticipate that their personal rewards will depend substantially on the profits that they generate for the corporation, because they should have strong incentives to increase these profits. But then corporate risks will generate substantial personal risks for senior managers. So a senior manager may want to evaluate corporate risks using a risk-averse corporate utility function similar to those we have studied for individual decision making. If $F(\mathbf{X})$ denotes the senior manager's income when $\mathbf{X}$ is the corporation's total profit (under the manager's compensation plan) and if $U$ denotes the manager's personal utility function for income, then the manager should personally prefer that corporate decisions under uncertainty should maximize $E(U(F(\mathbf{X})))$. In effect, the manager would want the corporation to evaluate risky profits by the expected value of a utility function $V$ such that $V(x) = U(F(x))$ for any amount of corporate profits $x$. But as

Myerson_C08.qxd 07-09-2004 09:45 Page 306

agents of the corporation, senior managers are not supposed to define the corporation's interests so blatantly in terms of their own, and so decision analysts rarely use such an approach to assessing a corporate utility function. Nevertheless, the sensitivity of a corporation to risks cannot be well understood unless we recognize that corporate risks typically generate substantial personal risks for the risk-averse senior managers who are responsible for making corporate decisions. In this context, decision analysts have found that senior managers may be comfortable with an expected-utility analysis of corporate risks, using risk tolerances for corporate profits that are sometimes about 1/6 of the total value of the corporation.

The perspective of the stockholders on corporation decision making is often very different. Individual stockholders may be very risk-averse, but a single corporation's risks may have only a very small impact on the overall net worth of a typical well-diversified individual stockholder. From the perspective of such well-diversified stockholders, the goal of corporate decision making should be to adopt strategies that, when generally understood by investors, will increase the market value of the corporations's stock. To apply this criterion, we need some theory to predict how prices in the stock market are determined. So we will discuss here the elements of such a theory, which is based on a simplified, idealized version of the stock market but which can offer real practical insights into the pricing of capital assets.

The most important fact about the stock market is that the future prices of stocks and other financial assets are unknown quantities. So any individuals' beliefs about a collection of asset prices at some future date should be described by a joint probability distribution. In our financial theory, we assume that all investors' beliefs about future prices at some future date (say, a year from now) can be described by some joint probability distribution. That is, we assume that everybody shares the same beliefs, which can be measured by a joint probability distribution for these unknown quantities. The end result of our asset pricing theory will be to show how to compute what an asset should be worth now in the stock market given such a joint probability distribution of future asset prices.

Let us assume here that investors in the stock market have constant risk tolerance, so that we can apply what we have already learned about optimal risk sharing among such individuals. The *market portfolio* is the collection of all the risky assets that are available to be bought and sold in the stock market (or in other financial markets). All the risks in this market portfolio must be shared in some way among the investors, and we may think of the entire stock market as a risk-sharing partnership that includes all investors. So Fact 1 tells us that when the stock market achieves an optimal allocation of risks, investors should have a share in the whole market portfolio of risky assets that is proportional to their risk tolerance. If any investor owned more or less of some asset than that investor's proportional share, then, by Fact 3, investors could make mutually beneficial trades of money and risky assets until an optimal allocation of risks is achieved. So in a market equilibrium, the prices of all financial assets should be such that

| | A | B | C | D | E | F | G | H | I | J |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | M = $Returns next year to $1 now in the Market Portfolio | | | | | | | | | |
| 2 | X,Y = $Values (per share) of Assets 1 and 2 next year | | | | | | | | | |
| 3 | | M | X | Y | | SOLVER: max F13 by changing F11. | | | | |
| 4 | LogMean | 0.09 | 20 | 20 | Q1 | | | | | |
| 5 | LogStdev | 0.25 | 25 | 25 | Q2 | | | | | |
| 6 | | | 32 | 32 | Q3 | | | What if prices now are: | | |
| 7 | | Correl(M,X) | | | | | | for X | | for Y |
| 8 | | 0.8 | | | | | | 25 | | 25 |
| 9 | 1+interest | | RiskTol | | | Invest now | | ...invest more? | | |
| 10 | 1.06 | | 10000 | | | $ in M | | $ in X | or | $ in Y |
| 11 | Y is independent of M and X | | | | | 9703 | | 1 | | 1 |
| 12 | corands | 0.9703 | 0.9772 | | | CE | | CE change | | |
| 13 | | M | X | Y | | 332.713 | | -0.0377 | | 0.04125 |
| 14 | SimTable | 1.753 | 54.96 | 23.66 | | Income | | | | |
| 15 | 0 | 1.024 | 17.74 | 23.53 | | -348.09 | | -348.45 | | -348.21 |
| 16 | 0.002 | 1.091 | 30.98 | 41.16 | | 299.817 | | 299.996 | | 300.403 |
| 17 | 0.004 | 1.097 | 20.80 | 48.35 | | 362.924 | | 362.696 | | 363.798 |
| 18 | 0.006 | 1.189 | 23.84 | 18.74 | | 1252 | | 1251.9 | | 1251.69 |
| 19 | 0.008 | 1.419 | 32.84 | 24.72 | | 3483.57 | | 3483.82 | | 3483.5 |
| 511 | 0.992 | 1.508 | 49.79 | 19.99 | | 4348.74 | | 4349.67 | | 4348.48 |
| 512 | 0.994 | 0.689 | 19.60 | 36.04 | | -3597.1 | | -3597.4 | | -3596.7 |
| 513 | 0.996 | 1.005 | 29.54 | 25.08 | | -537.17 | | -537.05 | | -537.22 |
| 514 | 0.998 | 0.912 | 18.92 | 19.57 | | -1434.4 | | -1434.7 | | -1434.7 |
| 515 | 1 | 1.110 | 32.08 | 22.38 | | 486.686 | | 486.909 | | 486.521 |
| 516 | | | | | | | | | | |
| 517 | FORMULAS | | | | | | | | | |
| 518 | B12:C12. {=CORAND(B8)} | | | | | | | | | |
| 519 | B14. =EXP(NORMINV(B12,B4,B5)) | | | | | | | | | |
| 520 | C14. =GENLINV(C12,C4,C5,C6) | | | | | | | | | |
| 521 | D14. =GENLINV(RAND(),D4,D5,D6) | | | | | | | | | |
| 522 | F15. =$F$11*(B15-$A$10) | | | | | | | | | |
| 523 | F15 copied to F15:F515 | | | | | | | | | |
| 524 | F13. =CE(F15:F515,C10) | | | | | | | | | |
| 525 | H15. =F15+$H$11*(C15/$H$8-$A$10) | | | | | | | | | |
| 526 | J15. =F15+$J$11*(D15/$J$8-$A$10) | | | | | | | | | |
| 527 | H15:J15 copied to H15:J515 | | | | | | | | | |
| 528 | H13. =CE(H15:H515,$C$10)-$F$13 | | | | | | | | | |
| 529 | J13. =CE(J15:J515,$C$10)-$F$13 | | | | | | | | | |

FIGURE **8.8** **Comparing investments that differ in correlation with the market portfolio**

an investor with constant risk tolerance is willing to hold the same share of every risky asset in the market portfolio, and the only difference among investors should be that their shares differ in proportion to their risk tolerances. Thus, all investors could simply buy shares in well-diversified mutual funds that hold all the stocks in the market portfolio. But every corporation's stock is a part of this overall market portfolio. So the price of a stock today must be such that when investors have bought their optimal share of the market portfolio, then they do not want to buy or sell any additional amounts of this stock now.

To understand all these complex ideas of asset pricing, let us apply them to a specific example, as illustrated in Figures 8.8 and 8.9.

In this example, we consider probability information about values next year of the market portfolio and a couple of selected assets in it. We assume that a dollar invested today in the market portfolio (that is, in a well-diversified mutual fund that holds this market portfolio on behalf of its investors)

Myerson_C08.qxd 07-09-2004 09:45 Page 308

will be worth an amount next year that has a Lognormal probability distribution with log-mean 0.09 and log-standard-deviation 0.25. That is, the logarithmic growth rate of an investment in the market portfolio has expected value 0.09 and standard deviation 0.25. We also assume that investors can borrow or lend at a risk-free interest rate such that $1 today returns $1.06 next year, as indicated in cell A10 of Figure 8.8.

Cells B15:B515 of Figure 8.8 contain 501 simulated values of the returns next year per dollar now in the market portfolio, generated by a random variable in cell B14 that has this Lognormal distribution. (Notice that rows 20 to 510 are hidden in Figure 8.8.) Now consider an individual who has a constant risk tolerance of $10,000, as indicated in cell C10. Suppose that cell F11 denotes the amount of money that this individual chooses to invest now in the market portfolio. Let us account for the results of this individual's investments in the stock market in terms of gains or losses next year relative to the wealth that would be produced next year if all the wealth now were deposited into a risk-free bank account that pays $1.06 next year for each $1 invested now. Relative to this safe strategy, each dollar invested in the market portfolio costs $1.06 next year (regardless of whether the dollar is taken out of the bank account or is borrowed given our simplifying assumption that the investor can both borrow and lend at the same risk-free rate). So if B15 were the market portfolio's growth ratio over the next year, then the net gain next year from investing F11 dollars now in the market portfolio can be computed in cell F15 by the formula

```
=$F$11*(B15–$A$10)
```

Copying cell F15 to F15:F515 yields the investor's net gains next year from investing F11 dollars now in the market portfolio for all of our simulated market outcomes. The certainty equivalent of these gains is computed in cell F13 by the formula

```
=CE(F15:F515,C10)
```

Now we can ask Solver to maximize the investor's certainty equivalent in cell F13 by changing the investment in cell F11, and the resulting optimal solution is as shown in Figure 8.8. That is, the optimal strategy (with this simulation data) is to invest $9703 in the market portfolio now, which will yield a net gamble that is as good as the certainty equivalent gain of $332.71 (relative to what the investor could get next year by putting all the wealth now in the bank). If we changed the risk tolerance and ran Solver again, then the optimal investment in cell F11 would change in proportion to the risk tolerance. For example, if we doubled the risk tolerance to $20,000 into cell C10 and ran Solver again, then Solver would maximize the certainty equivalent F13 by doubling the investment amount to $19,406 in cell F11.

The optimal investment quantity for each investor in these calculations depends on the risk-free interest rate that banks are assumed to pay (in cell

Myerson_C08.qxd 07-09-2004 09:45 Page 309

A10). If this interest rate were lower, then the net gains from putting money in the stock market would be greater, and the optimal investment in the risky market portfolio would increase for any individual investor. In a market equilibrium, the risk-free interest rate has been determined by the condition that the sum of all individuals' investments in shares of the market portfolio must equal the whole portfolio of stocks and other financial assets available in the market.

So in a market equilibrium, all financial assets must be priced now in such a way that when individual investors solve the optimal investment problem that we have just described, they buy different shares of the general market portfolio, and no one wants to buy or sell any further shares of any financial asset. This no-further-trade condition can be used to characterize the prices that must be paid for financial assets in a market equilibrium.

Two other financial assets are considered in Figure 8.8. You should think of these as being just two among the very large number of stocks in the stock market that make up the market portfolio. Let **X** denote the price per share of the first stock next year, and let **Y** denote the price per share of the second stock next year. To be specific, we suppose here that these two stock prices next year have the same marginal probability distribution, which is a Generalized Lognormal distribution with quartile points \$20, \$25, and \$32. But suppose that the first stock's future price **X** has a correlation of 0.8 with the growth ratio of the market portfolio whereas the second stock's price **Y** is independent of the market portfolio. These two random variables are simulated in cells C14 and D14 of Figure 8.8, where the array formula {=CORAND(0.8)} in cells B12:C12 provides the specified correlation between the simulated **X** in cell C14 and the simulated growth of the market portfolio (**M**) in cell B14. The simulated **Y** in cell D14 here depends independently on its own RAND( ).

Column H analyzes the effects of making another small investment now in the first stock, in addition to the individual's F11 optimal investment in the market portfolio. The effect of such an additional investment would depend, of course, on the price that the investor has to pay now for a share of this stock. Our goal in this exercise is to compute what this price should be (based on the probability information about **X**), but let us begin here by entering in cell H8 some arbitrary guess about what this current stock price might be. Next, let us specify in cell H11 any small amount of money that the investor might consider adding now to the investment in this stock. Figure 8.8 considers the possibility of investing an additional \$1 (H11) to this stock when its price per share now is \$25 (H8). For the first outcome of our simulation model, stored in row 15 of the simulation table, this additional investment would change the investor's net gains next year from the amount in cell F15 to

```
=F15+$H$11*(C15/$H$8–$A$10)
```

which has been entered into cell H15. To understand this formula, notice that each of the H11 (1) additional dollars invested now buys 1/H8 (1/25)

Myerson_C08.qxd 07-09-2004 09:45 Page 310

shares of this stock at the current price, which will yield next year C15/H8 dollars in this simulated market outcome, but investing the dollar now also reduces the individual's bank account by A10 (1.06) dollars next year. Then copying the formula from H15 to cells H15:H515, we get the net gains next year with this additional investment for all the outcomes in our simulation table. So the change in the investor's certainty equivalent next year that would be generated by this additional investment in the first stock is computed in cell H15 by the formula

=CE(H15:H515,$C$10)–$F$13

A similar calculation is done for the second stock in column J of Figure 8.8. Here cell J8 contains our guess about the current price per share of this stock, and cell J11 contains an amount of money that the investor is thinking of investing in this stock now, as an addition to the optimal investment in the market portfolio. Then, by copying cells H13:H515 to J13:J515, we similarly compute in cell J13 the change in the investor's certainty equivalent next year that would be generated by this additional investment in the second stock. (These additional investments are being considered here as alternative possibilities, and so the simulated gains in J15:J515 include the J11 additional investment in the second stock but do not include the H11 additional investment in the first stock.)

In Figure 8.8, the computed changes of certainty equivalent in cells H13 and J13 are –0.0377 and 0.04125, respectively. This means that if the prices of these stocks were both $25 now (as assumed in cells H8 and J8), then investing an additional dollar in the first stock would decrease the investor's certainty equivalent next year by almost 4 cents, but investing an additional dollar in the second stock would increase the investor's certainty equivalent next year by a bit more than 4 cents. But remember, in a market equilibrium, investors should hold an optimal share in the whole market portfolio that is proportional to their risk tolerance and should not want to add any other disproportionate investment in particular stocks. So if the positive value of cell J13 indicates that this investor could do strictly better by adding another $1 (J11) in the second stock here when its current price is as specified in cell J8 ($25) , then the assumed price in cell J8 must be too low. So in an equilibrium of the stock market now, the price per share of the second stock must be strictly more than $25, because investors like this one would want to buy more of this stock for $25 now.

The negative value of cell H13 in Figure 8.8 indicates that this investor could not improve the portfolio by adding another $1 (H11) in the first stock. But what about selling some of the first stock? Remember, every stock is a small part of the market portfolio, and so the optimal investment in the market portfolio must implicitly include some small amount of money in this stock, which (in principle) could be sold separately now. Selling the first stock can be represented in this spreadsheet by entering a negative number in cell H11. For example, if we entered the value –1 into cell H11, to represent the investor selling $1's worth of the first stock now, then the

Myerson_C08.qxd 07-09-2004 09:45 Page 311

calculations in cell H13 would show the certainty equivalent increasing next year by \$0.0377. (The increase of *CE* from selling \$1's worth would be almost the same size as the decrease of *CE* from buying \$1's worth.) This result indicates that in an equilibrium of the stock market now, the price per share of the first stock must be strictly less than \$25, because investors like this one would want to sell their shares of this stock for \$25 now.

So the current price per share of the first stock should be less than \$25 and the current price per share of the second stock should be more than \$25, even though investors anticipate that the prices per share of these stocks next year have the same marginal probability distribution. The only difference between these stocks is that the first (**X**) has a positive correlation with the market portfolio, whereas the second (**Y**) is independent and so has zero correlation with the market portfolio. This observation illustrates an important general fact: A higher correlation with the market portfolio tends to decrease the current value of the asset (when the marginal probability distribution of its value at some future date is held fixed).

Intuitively, the second stock is worth more than the first stock now because the second stock can offer better insurance against the risks in the market portfolio that the investors are holding. The investors, being risk-averse, would get more extra utility from an extra dollar of income next year in an event where their overall portfolio has done badly than in an event when their overall portfolio has done well. But the first stock, being highly correlated with the market portfolio, tends to pay its highest returns in the same events where the market portfolio does well. Both of these stocks have the same probability of selling for more than \$32 next year (0.25), but there is a much smaller probability of the first stock being so profitable when the overall market portfolio does badly, because of its positive correlation with the market portfolio.

Now let us see how you could use the spreadsheet in Figure 8.8 to compute the current price of the first stock. Remember, cell F11 contains the optimal investment in the market portfolio that Solver generated to maximize the certainty equivalent in cell F13. Keeping this value in cell F11, you could now ask Solver to change the additional investment in cell H11 to maximize the additional certainty equivalent that it generates in cell H13. If Solver returns a positive investment in cell H11, then the price for the first stock that you assumed in cell H8 is too low. If Solver returns a negative investment in cell H11 (indicating that investors would want to sell this stock at this price), then the price for the first stock that you assumed in cell H8 is too high. When the right current price for the first stock is entered into cell H8, Solver should report that the additional certainty equivalent in H13 is maximized by setting the additional investment in H11 equal to \$0. For the simulation data in this spreadsheet, the correct current price of the first stock turns out to be \$24.11. (Results may differ slightly with other simulation data.)

Similarly, you can find the current equilibrium price for the second stock in Figure 8.8 by adjusting the price in cell J8 so that when Solver is asked to

Myerson_C08.qxd 07-09-2004 09:45 Page 312

| | A | B | C | D | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | M = $Returns next year to $1 now in the Market Portfolio | | | | | | | | | | |
| 2 | X = $Value (per share) of Asset next year | | | | | | | | | | |
| 3 | | M | X | | | SOLVER(1): max E13 by changing E11. | | | | | |
| 4 | LogMean | 0.09 | 20 | Q1 | | SOLVER(2): max I13 by changing I11. | | | | | |
| 5 | LogStdev | 0.25 | 25 | Q2 | | | (Should find I11=0 optimal!) | | | | |
| 6 | | | 32 | Q3 | | | | | | | |
| 7 | | Correl(M,X) | | | | Value of Asset now = E'(X)/(1+i) | | | | | |
| 8 | | 0.8 | | | | | | | 24.11 | | |
| 9 | 1+interest | | RiskTol | | Invest now | | | | Invest more? | | |
| 10 | 1.06 | | 10000 | | $ in M | | | | $ in X | | |
| 11 | | | | | 9703 | | | | 0 | | |
| 12 | corands | 0.5563 | 0.4176 | | CE | UINV(EU) | | | CE change | | |
| 13 | | M | X | | 332.71 | 332.71 | | | 0 | | |
| 14 | SimTable | 1.134 | 23.27 | | Income | Util | Weights | | Income | | |
| 15 | 0 | 1.024 | 17.74 | | -348.1 | -1.035 | 0.0021 | | -348.1 | | |
| 16 | 0.002 | 1.091 | 30.98 | | 299.82 | -0.97 | 0.0020 | | 299.82 | | |
| 17 | 0.004 | 1.097 | 20.80 | | 362.92 | -0.964 | 0.0020 | | 362.92 | | |
| 18 | 0.006 | 1.189 | 23.84 | | 1252 | -0.882 | 0.0018 | | 1252 | | |
| 19 | 0.008 | 1.419 | 32.84 | | 3483.6 | -0.706 | 0.0015 | | 3483.6 | | |
| 511 | 0.992 | 1.508 | 49.79 | | 4348.7 | -0.647 | 0.0013 | | 4348.7 | | |
| 512 | 0.994 | 0.689 | 19.60 | | -3597 | -1.433 | 0.0030 | | -3597 | | |
| 513 | 0.996 | 1.005 | 29.54 | | -537.2 | -1.055 | 0.0022 | | -537.2 | | |
| 514 | 0.998 | 0.912 | 18.92 | | -1434 | -1.154 | 0.0024 | | -1434 | | |
| 515 | 1 | 1.110 | 32.08 | | 486.69 | -0.952 | 0.0020 | | 486.69 | | |
| 516 | | | | | | | | | | | |
| 517 | FORMULAS | | | | | | | | | | |
| 518 | B12:C12. {=CORAND(B8)} | | | | | | | | E'(X) | | |
| 519 | B14. =EXP(NORMINV(B12,B4,B5)) | | | | | | | | 25.56 | | |
| 520 | C14. =GENLINV(C12,C4,C5,C6) | | | | | | | | | | |
| 521 | E15. =$E$11*(B15-$A$10) | | | | | | | | E(X) | | |
| 522 | E15 copied down | | | | | | | | 27.66 | | |
| 523 | E13. =CE(E15:E515,C10) | | | | | | | | | | |
| 524 | F15. =UTIL(E15,$C$10) | | | | | | | | | | |
| 525 | F15 copied down | | | | | | | | | | |
| 526 | F13. =UINV(AVERAGE(F15:F515),C10) | | | | | | | | | | |
| 527 | G15. =F15/SUM($F$15:$F$515) | | | | | | | | | | |
| 528 | G15 copied down | | | | | | | | | | |
| 529 | I8. =SUMPRODUCT(C15:C515,G15:G515)/A10 | | | | | | | | | | |
| 530 | I15. =E15+$I$11*(C15/$I$8-$A$10) | | | | | | | | | | |
| 531 | I15 copied down | | | | | | | | | | |
| 532 | I13. =CE(I15:I515,C10)-E13 | | | | | | | | | | |
| 533 | I519. =SUMPRODUCT(C15:C515,G15:G515) | | | | | | | | | | |
| 534 | I522. =AVERAGE(C15:C515) | | | | | | | | | | |

FIGURE **8.9**
**Computing an asset's value in a market with constant-risk-tolerant investors**

maximize J13 by changing J11, Solver will report an optimal additional investment of $0 in cell J11. For the simulation data in this spreadsheet, the correct current price of the second stock turns out to be $25.97. The difference between this $25.97 for the second stock and the price $24.11 for the first stock is caused by the difference between these stocks' correlations with the market portfolio.

Running Solver so many times may be awkward, so Figure 8.9 shows a way to directly calculate the current price of a stock from its distribution of future returns. To simplify the spreadsheet, we only consider the first stock (the one that will be worth **X** next year), but the simulation model and data in columns B and C of Figure 8.9 are the same as in Figure 8.8. The interest

Myerson_C08.qxd 07-09-2004 09:45 Page 313

rate in A10 and the investor's risk tolerance in C10 are also the same in both figures. Column E of Figure 8.9 repeats the analysis of the individual's optimal investment in the market portfolio that we already saw in column F of Figure 8.8. That is, cell E11 contains the individual's monetary investment now in the market portfolio, cells E15:E515 contain the resulting monetary gains next year from this investment (compared to the alternative of putting all the wealth in bank accounts) with the simulated market outcomes, and cell E13 computes the certainty equivalent for this gamble with the formula

```
=CE(E15:E515,C10)
```

The investment amount shown in E11 ($9703) is the result of asking Solver to maximize E13 by changing E11.

In column F in Figure 8.9, you can find a review of the calculations by which the CE function computes the certainty equivalent in cell E13. First, the utility values of the monetary incomes in cells E15:E515 are computed, by entering the formula

```
=UTIL(E15,$C$10)
```

into cell F15, and then copying F15 to F15:F515. Then these utility values are averaged, and the average utility is converted back to the equivalent monetary value by UINV with the formula

```
=UINV(AVERAGE(F15:F515),C10)
```

in cell F13. The result in cell F13 is the same as the certainty equivalent value that is computed by the CE function in cell E13.

But now let us do something else with these utility values in F15:F515. Recall from Chapter 3 that the constant-risk-tolerance utility functions are defined by the following mathematical formula:

$$\text{UTIL}(x, r) = -\text{EXP}(-x/r)$$

for any monetary amount $x$ and any positive risk tolerance $r$. This exponential formula has the special property that as the amount of money $x$ is varied with a fixed risk tolerance, the slope of the utility function is proportional to the absolute value of the utility function itself. If you look back at the graph in Figure 3.1, you can see that the slope of a constant-risk-tolerance utility curve is proportionally steeper in regions where the utility value is farther below zero. Thus, with these special exponential utility functions, the investor's utility benefit from a small increase in income should be greater in the bad outcomes where the utility is farther below zero, in direct proportion to the distance of the utility below zero. So the investor's relative sensitivity to income changes in the various simulated outcomes here may be measured by the weights shown in G15:G515 of Figure 8.9, which are computed by entering the formula

```
=F15/SUM($F$15:$F$515)
```

Myerson_C08.qxd 07-09-2004 09:45 Page 314

into cell G15 and then copying G15 to G15:G515. The utility values in F15:F515 are all negative numbers, but then their sum is also negative, and so these weights in G15:G515 are all positive numbers. (A negative number divided by a negative number yields a positive number.) Furthermore, the sum of these weights in G15:G515 is equal to 1. In effect, these weights look like a probability distribution over the various outcomes in our simulation table, where bad outcomes are given more probability than the good outcomes, in proportion to the slopes of the investor's utility curve at these different outcomes. For example, row 511 in Figure 8.9 gets a small weight, 0.0013, in cell G11 because this row represents a good outcome where the market portfolio goes up and the investor gains $4349 (E11). But row 512 gets a relatively large weight, 0.0030, in cell G12 because this row represents a bad outcome where the market portfolio goes down and the investor loses $3597, which would make the investor much more sensitive to additional income than would a good outcome.

Until now, we have always analyzed simulation data as if every outcome in every row of our simulation table represented an equally likely outcome from our probability model. But now let us pretend that the outcomes in the various rows have instead the probabilities shown in G15:G515, which we may call the *market-adjusted probabilities*. Let us use $E'$ to denote the expected value of a random variable when we apply these mysterious market-adjusted probabilities. With these market-adjusted probabilities, the expected price-per-share of our selected stock next year would be computed by the formula

$$E'(\mathbf{X}) = \text{SUMPRODUCT(C15:C515, G15:G515)}$$

because cells C15:C515 contain the simulation data for this stock's future value **X**. With the simulation data, this market-adjusted expected value is $E'(\mathbf{X}) = \$25.56$. For comparison, if we treated all rows as equally likely and conventionally estimated the expected value $E(\mathbf{X})$ by AVERAGE(C15:C515), then we would get $E(\mathbf{X}) = \$27.66$ with this simulation data. The market-adjusted expected value $E'(\mathbf{X})$ is smaller than $E(\mathbf{X})$ here because **X** is positively correlated with the market portfolio **M**, and so the high values of **X** tend to come in rows where the market portfolio does well and thus where the market-adjusted probability is relatively small.

The price **X** is a future value of this stock, and so it is denominated in future dollars, a year from today. But such future dollars can be exchanged for dollars today by borrowing or lending at a bank, at the ratio defined by the bank's interest rate, here $1 + i = 1.06$, as shown in cell A10 in Figure 8.9. So to convert future monetary values into current monetary values here, we should divide by A10. In particular, the current monetary equivalent of the market-adjusted expected value of the price per share of this stock next year is

$$E'(\mathbf{X})/(1 + i) = \text{SUMPRODUCT(C15:C515, G15:G515)/A10}$$

Myerson_C08.qxd 07-09-2004 09:45 Page 315

This important formula,

=SUMPRODUCT(C15:C515,G15:G515)/A10

has been entered into cell I8, and it is the answer to our question of what should be the equilibrium price of this stock today.

When today's price per share of this stock is as computed in cell I9 of Figure 8.9, the investor will not want to buy or sell any additional amounts of this stock now, once the optimal (E11) investment has been made in the market portfolio. To demonstrate this result, cell I13 in Figure 8.9 computes the net change in the certainty equivalent next year if the investor were to invest in this stock today the additional amount of money that is entered in cell I11, given the current price per share in cell I9. If we asked Solver to maximize this certainty equivalent in I13 by changing the investment amount in I11 (holding fixed the previously optimized market-portfolio investment in E11), then Solver cannot do better than by letting the investment amount in I11 be $0. So the formula in cell I9 gives us the correct equilibrium price today for this stock that will have the price **X** next year, because individual investors will not want to change their optimal investment in the market portfolio by buying or selling this stock at this price today.

The price that we have computed in cell I9 of Figure 8.9 does not depend on the particular risk tolerance of $10,000 that we assumed in cell C10. For example, suppose we doubled the risk tolerance by entering $20,000 into cell C10, and we then asked Solver again to maximize the new certainty equivalent in cell E13 by changing cell E11. Then the new optimal market-portfolio investment in cell E11 would double, and all the net incomes in cells E15:E515 would double, relative to the amounts shown in Figure 8.9. But then the utility values in cells F15:F515 would remain the same as shown in Figure 8.9, because UTIL($x$, $t$) = –EXP(–$x$/$r$) depends only on the ratio of the income $x$ to the risk tolerance $r$, and so doubling them both would leave the utility unchanged. So the market-adjusted probabilities in cells G15:G515 (which depend only on F15:F515) and the computed stock price in cell I9 (which depends only on G15:G515 and the simulation data and the interest rate) would also be unchanged. Thus, the market-adjusted probabilities and the computed stock price in Figure 8.9 would not be affected by a change of the risk tolerance in cell C10, once Solver has changed the optimal investment in E11 to maximize the new certainty equivalent in E13. Our results here depend only on the assumption that investors are risk-averse and have some constant risk tolerance.

This financial asset pricing model has important applications to corporate decision making. To develop these applications, we just need to extend our interpretation of the random variable **X** in cell C14.

Suppose that a big publicly held corporation is considering a major new investment project that would cost $25 million now. For simplicity, suppose that all the monetary returns from this project would be realized one

Myerson_C08.qxd 07-09-2004 09:45 Page 316

year from now, and that these returns next year would have a Generalized Lognormal probability distribution with quartile points \$20, \$25, and \$32 million. Suppose also that these returns would have a correlation 0.8 with the growth ratio of the market portfolio next year. Thus, the returns from this project next year would have the same marginal probability distribution and the same relationship with the market portfolio as the random variable **X** that we considered in Figures 8.8 and 8.9. So let us now reinterpret **X** as this unknown quantity that the project will earn next year.

Should this project be recommended now as being to the benefit of the corporation's stockholders? This is a decision under uncertainty, because the payoff **X** next year is an unknown quantity that could be much greater or much less than the cost of the project. But current stockholders in this corporation could sell shares tomorrow, and so all stockholders can benefit from decisions that will increase the current market value of the corporations's stock. So we should ask whether a decision to undertake this project would tend to increase or decrease the value of this corporation in the stock market now (once the decision becomes widely known to investors).

According to our analysis in Figure 8.9, a stock coupon (or a lottery ticket) that will be worth **X** dollars next year should have a value of \$24.11 in the stock market today. When this corporation invests in this project that will pay **X** million dollars next year, it is essentially buying a million such coupons. So the corporation's investment in this project should be worth \$24.11 million in the stock market today. That is, the prospect of earning **X** million dollars next year from this project should add \$24.11 million to the total value of the corporation's stock. But undertaking this project requires an expense of \$25 million today. This expense would make the corporation increase its debt (or decrease its liquid assets) by \$25 million, which should have the effect of subtracting \$25 million now from the total market value of the corporation. So according to our analysis, a decision to undertake this project should change the current total value of this corporation by 24.11 – 25 = –\$0.89 million as soon as the decision becomes understood by investors in the stock market. Thus, the project should not be recommended on behalf of the stockholders because it would tend now to decrease the value of their shares in the corporation.

This conclusion depends critically on our assumption that the project's future earnings have correlation 0.8 with the future performance of the market portfolio. If the project's future earnings were considered instead to be independent of the future performance of the market portfolio, then the project's future earnings would have a current value that is greater than \$25 million, because it would be like the gamble that pays **Y** in Figure 8.8. In general, the value of a project or risky investment to investors in the stock market cannot be determined without analyzing how the project's risks are related to the broader macroeconomic risks that investors bear in the market portfolio. To do such analysis, we must use a model of the joint probability distribution of the investment's returns **X** and the market portfolio **M**, as in cells B14:C14 of Figure 8.9.

Myerson_C08.qxd 07-09-2004 09:45 Page 317

(Financial analysts often estimate asset values using the *capital asset pricing model*, or *CAPM*, which is similar to the constant-risk-tolerance model that we have developed here. CAPM is based on an assumption that investors care only about the mean and standard deviation of their portfolio's value, and so it yields an asset pricing formula that depends only on the asset's expected value and covariance of returns with the market portfolio. When future asset returns are drawn from a Multivariate Normal distribution, the CAPM model is equivalent to the model that we have developed here. Even for the non-Normal data in Figure 8.9, a CAPM analysis would yield a stock price differing by only \$0.05 from the value that we computed in cell I9. But CAPM can yield nonsensical results for some extremely skewed non-Normal distributions that are avoided by the method that we have developed here.)

## 8.6 Fundamental Ideas of Arbitrage Pricing Theory

In arbitrage pricing theory, we assume that there is some list of possible *states* of the world such that exactly one of these states will occur and the future values of financial assets will be determined by the state that occurs. We would say that an *arbitrage opportunity* existed if there were some portfolio of loans and assets that an individual could acquire with zero initial personal investment but that would guarantee a positive future return in all states. That is, an arbitrage opportunity would be an opportunity to get something for nothing without any risks. Arbitrage pricing theory characterizes financial asset prices when such arbitrage opportunities do not exist.

Let us begin by considering a simple two-state example in which State 1 is a state of the world with high oil prices and State 2 is a state of the world with low oil prices. Suppose that the value of some automotive company's stock will be \$80 per share next year if State 1 occurs, but it will be \$140 if State 2 occurs. Similarly, suppose that the future value of some oil company's stock will be \$140 per share next year if State 1 occurs, but it will be \$80 if State 2 occurs. Suppose also that investors can borrow and lend at a 10% annual interest rate, so that 1.1 is the annual return ratio for risk-free bonds. If these two stocks are currently selling for \$90 per share, then an arbitrage opportunity exists. For each \$180 that we borrow, we could buy one share of each of these two stocks now, and then next year (after selling the shares and repaying the debt) we could take a sure profit of (80 + 140) – 180∗1.10 = \$22 in all possible states. By borrowing more money, we could make as much money as we like, with no risk to ourselves.

If the current prices for these two stocks are each \$100 per share, however, then such arbitrage opportunities would not exist. At this price, if we assessed a probability 0.5 for each of the two states, then the expected returns next

Myerson_C08.qxd 07-09-2004 09:45 Page 318

year per dollar invested now would be (0.5∗80 + 0.5∗140)/100 = $1.10, which exactly equals the cost next year of borrowing a dollar now. If every financial asset that we can buy or sell has an expected return of 10%, then, no matter how we mix investments and debts, the expected value of our portfolio next year will be 10% more than its value this year. In particular, if we start with no initial investment of our own funds, then the expected total value of our portfolio must be $0 (= 0∗1.10) next year, and so there cannot exist any arbitrage strategy that offers positive returns in all states with zero net initial investment.

In general, if there exists some way of assigning probabilities to the various possible states such that every financial asset offers the same expected return ratio as risk-free bonds, then arbitrage opportunities cannot exist. With each investment offering the same expected return ratio as the risk-free bonds, no portfolio can offer a higher (or lower) expected return ratio than this risk-free return ratio on the net initial investment. But an arbitrage strategy that guarantees a positive return with $0 initial investment would be offering an infinite expected return ratio. So there cannot exist any arbitrage strategy for generating positive returns in all states with zero net initial investment.

Conversely, if arbitrage opportunities do not exist in a financial market, then there must exist some way of assigning probabilities to the possible states such that when we compute expected returns using this probability distribution, the expected return ratio of every financial asset is equal to the risk-free return ratio. This important fact is the main result of arbitrage pricing theory.

To illustrate this result, consider the example shown in Figure 8.10, which describes a simple imaginary financial market in which three stocks are traded, and the return ratios for these stocks over the next year will depend on which of four possible states of the world occurs. The table in cells B5:E7 lists the return ratio for each stock in each possible state of the world. The return ratio for risk-free bonds is 1.10, which is listed in cell A1.

Given this financial data, Figure 8.10 shows how to use Solver to find an arbitrage opportunity if one exists. Cells G5:G7 represent the money to be invested in each of the three stocks now in our investment strategy. If any of these cells becomes negative, it can be interpreted as the amount of money to be raised now by selling the corresponding stock short. The net investment in these stocks is assumed to come out of bonds that pay the risk-free interest rate of 10% per year, and so the formula

=SUMPRODUCT(B5:B7,$G$5:$G$7) –SUM($G$5:$G$7)∗$A$1

in cell B9 represents the net returns from our investments if State 1 occurs. Copying cell B9 to B9:E9 gives us cells representing the net returns from our investments in each of the four possible states. In cell G10, we enter a "goal" of returns that our investment strategy will try to achieve in all

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | 1.1 | Risk-free return ratio ($ next year per $1 invested now) | | | | | | |
| 2 | | | | | | | | |
| 3 | $Returns next year per $1 invested now | | | | | | | |
| 4 | | State 1 | State 2 | State 3 | State 4 | | Invest now | |
| 5 | Stock A | 0.95 | 0.90 | 1.30 | 1.20 | | 0 | |
| 6 | Stock B | 0.95 | 1.35 | 0.85 | 1.00 | | 0 | |
| 7 | Stock C | 1.30 | 1.10 | 1.15 | 1.00 | | 0 | |
| 8 | | | | | | | | |
| 9 | Net $return | 0 | 0 | 0 | 0 | | Goal | |
| 10 | | | | | | | 0 | |
| 11 | Shortfall | 0 | 0 | 0 | 0 | (+ is bad!) | | |
| 12 | | | | | | | | |
| 13 | SOLVER (with Options:AssumeLinearModel): | | | | | | | |
| 14 | Maximize G10 by changing G5:G7,G10 subject to B11:E11<=0. | | | | | | | |
| 15 | Select SensitivityReport when Solver finishes. | | | | | | | |
| 16 | | | | | | | | |
| 17 | | State 1 | State 2 | State 3 | State 4 | | | |
| 18 | ShadowProby | 0.01504 | 0.43609 | 0.34586 | 0.20301 | | | |
| 19 | | ShadowE($Returns) | | | | | | |
| 20 | Stock A | 1.1 | | | | | | |
| 21 | Stock B | 1.1 | | | | | | |
| 22 | Stock C | 1.1 | | | | | | |
| 23 | | | | | | | | |
| 24 | FORMULAS FROM RANGE A1:G22 | | | | | | | |
| 25 | B9. =SUMPRODUCT(B5:B7,$G$5:$G$7)-SUM($G$5:$G$7)*$A$1 | | | | | | | |
| 26 | B9 copied to B9:E9 | | | | | | | |
| 27 | B11. =$G$10-B9 | | | | | | | |
| 28 | B11 copied to B11:E11 | | | | | | | |
| 29 | B20. =SUMPRODUCT($B$18:$E$18,B5:E5) | | | | | | | |
| 30 | B20 copied to B20:B22 | | | | | | | |
| 31 | Shadow probabilities (or Lagrange Multipliers) in B18:E18 are copied | | | | | | | |
| 32 | (with paste-special,transpose) from Solver's SensitivityReport. | | | | | | | |
| 33 | | | | | | | | |
| 34 | Fact: If arbitrage opportunities do not exist, then Solver must | | | | | | | |
| 35 | terminate with 0 optimal value in cell G10, and the Shadow Prices | | | | | | | |
| 36 | (or Lagrange Multipliers) in the Sensitivity Report will give us a | | | | | | | |
| 37 | shadow probability distribution over the states such that all assets | | | | | | | |
| 38 | have the same shadow-expected return ratio as risk-free bonds. | | | | | | | |

FIGURE **8.10**
**A simple example of arbitrage pricing theory**

states. The shortfall from this goal in each state is computed in cells B11:E11 by entering the formula

=$G$10–B9

in cell B11, and then copying B11 to B11:E11. Thus, a positive shortfall in cells B11:E11 denotes a failure to achieve the goal in some state. With this spreadsheet formulation, we can now ask Solver to find the highest goal that can be achieved in all states by an investment strategy.

In the Solver dialog box, let us tell Solver to maximize the target cell G10 by changing cells G5:G7, G10 subject to the constraints B11:E11 <= 0. (It is important to include G10 among the changing cells as shown. Solver will accept multiple ranges separated by commas in the changing-cells box.) To take advantage of the special linear structure of this problem, which will enable Solver to analyze this problem more accurately, let us also go to the

Myerson_C08.qxd 07-09-2004 09:45 Page 320

**Microsoft Excel Sensitivity Report**

**Changing Cells**

| Cell | Name | Final Value | Reduced Cost | Objective Coefficient | Allowable Increase | Allowable Decrease |
|---|---|---|---|---|---|---|
| $G$5 | Stock A Invest now | 0 | 0 | 0 | 0.05227 | 0.00370 |
| $G$6 | Stock B Invest now | 0 | 0 | 0 | 0.08846 | 0.00455 |
| $G$7 | Stock C Invest now | 0 | 0 | 0 | 0.05610 | 0.02500 |
| $G$10 | Goal | 0 | 0 | 1 | 1E+30 | 1 |

**Constraints**

| Cell | Name | Final Value | Shadow Price | Constraint R.H. Side | Allowable Increase | Allowable Decrease |
|---|---|---|---|---|---|---|
| $B$11 | Shortfall State 1 | 0 | 0.01504 | 0 | 1E+30 | 1E+30 |
| $C$11 | Shortfall State 2 | 0 | 0.43609 | 0 | 1E+30 | 1E+30 |
| $D$11 | Shortfall State 3 | 0 | 0.34586 | 0 | 1E+30 | 1E+30 |
| $E$11 | Shortfall State 4 | 0 | 0.20301 | 0 | 1E+30 | 1E+30 |

FIGURE **8.11** **Solver Sensitivity Report for the problem in Figure 8.10**

Solver Options dialog box and check the Assume Linear Model option (then OK). Then we can go back to the basic Solver dialog box and click the Solve button. When Solver finishes and announces that it has found a solution, we should also select the Sensitivity option in the Reports box, which causes Solver to add a Sensitivity-Report page in our workbook.

For this example, Solver will report the value 0 in cell G10 as the maximum return that can be guaranteed with no net investment. (Solver may report a value of G10 slightly different from 0, such as 1.1E–06, which denotes 1.1∗10^(–6) = 0.0000011, but this tiny deviation from 0 is just due to roundoff error.) Getting this Solver output tells us that an arbitrage strategy for guaranteeing a positive return in all states with no net investment does not exist. Thus, the main result of arbitrage pricing theory tells us that there must exist some probability distribution that makes each stock have an expected return ratio equal to the return ratio of risk-free bonds. But where can we find this probability distribution?

The answer is in the Sensitivity Report that Solver added to the workbook when we selected the Sensitivity option when Solver finished the optimization. Solver's Sensitivity Report, as shown in Figure 8.11, includes a lot of numbers, but the only important numbers for our purposes are those listed for the constraints under the heading *Shadow Price* or *Lagrange Multiplier*. (Solver uses the term *Shadow Price* when the Assume Linear Model option has been selected, and it uses the equivalent term *Lagrange Multiplier* otherwise.) In general, the shadow price of a constraint is a measure of the rate at which the optimal value of the target cell would increase if we started increasing the value on the right-hand side of the constraint. But Solver's shadow prices have a special interpretation here in this problem of searching for arbitrage opportunities. This problem includes one constraint for each of the possible states, and the shadow prices of these constraints are nonnegative numbers that sum to 1. So the shadow

Myerson_C08.qxd 07-09-2004 09:45 Page 321

prices generated by Solver for this problem can be interpreted as a *shadow probability distribution* over the set of possible states.

This shadow probability distribution from the Sensitivity Report has been copied (and pasted-special transposed) to the range B18:E18 in Figure 8.10. Then the formula

```
=SUMPRODUCT($B$18:$E$18,B5:E5)
```

in cell B20 computes the expected return ratio for Stock A under this shadow probability distribution. Copying this formula to cells B20:B22, we find that all three stocks have shadow expected return ratios equal to the return ratio on risk-free bonds in this market.

This result is completely general. If you change the returns listed in cells B5:E7 of Figure 8.10 or the risk-free return ratio listed in cell A1, then one of two cases will hold. *Case 1* is that the Solver will find an optimal solution with 0 as the best goal that can be guaranteed in cell G10, in which case the shadow prices in the Sensitivity Report will form a shadow probability distribution with which all stocks offer the same expected return ratio as risk-free bonds. *Case 2* is that Solver may report "the set cell values do not converge," which means that arbitrage opportunities are possible and returns that exceed any positive goal can be guaranteed.

> **NOTE** (If you want to see what these arbitrage strategies look like, you can add the constraint G10 <= 1000, which tells Solver to try to stop when it finds a way to guarantee returns higher than $1000 with no initial investment. To get an example where such arbitrage opportunities exist, you can change the value of cell B5 in Figure 8.10 from 0.95 to 1.19 or higher, leaving all other parameters the same.)

Now look back at our asset pricing model in Figure 8.9. The current value of the asset, as computed in cell I9, is an expected value of the asset next year divided by the return ratio for risk-free bonds, but where the expected value is computed using the market-adjusted probabilities in cells G15:G515. So when these market-adjusted probabilities are used to compute the asset's expected return ratio (its future values in C15:C515 divided by its current value in I9), the resulting market-adjusted expected return ratio is exactly the risk-free return ratio in cell A10. Thus, the market-adjusted probabilities in cells G15:G515 of Figure 8.9 are the shadow probabilities that make this asset pricing model a special case of arbitrage pricing theory.

But Figure 8.9 also teaches us more about how such market-adjusted shadow probabilities may differ from traditional probabilities (which are defined by relative frequency, and so would all be equal to 1/501 in the analysis of 501 simulation outcomes). The difference is that higher market-adjusted probabilities are assigned to the bad outcomes where the market portfolio does poorly, and lower market-adjusted probabilities are assigned to the good outcomes where the market portfolio does well. This overweighting of bad outcomes occurs because risk-averse investors should be more

Myerson_C08.qxd 07-09-2004 09:45 Page 322

sensitive to income changes in the states where their overall portfolio has done badly.

## 8.7 Summary

Individuals with constant risk tolerance can maximize the sum of their certainty equivalents by sharing risks linearly in proportion to their individual risk tolerances. If the members of an investment partnership were not using such an optimal sharing rule, then one partner could propose such an optimal sharing rule, together with some fixed payments among the partners, with the result that this partner does better and no one else does worse (in terms of their certainty equivalents). When partners share a risky gamble in this optimal way, the sum of their certainty equivalents of their separate shares is the same as what the certainty equivalent of the whole gamble would be to a single individual whose risk tolerance was the sum of the partners' risk tolerances. In this sense, a partnership should evaluate gambles and choose among them just as if it were a single person with a risk tolerance equal to the sum of the partners' individual risk tolerances.

This simple theory of sharing risks in proportion to risk tolerances is based on an assumption that the way that profits are divided among the partners does not affect the probability distribution of total profits earned. But this assumption may fail, and so this simple theory of optimal risk sharing may not apply when the partnership's profits depend on the efforts of an individual whose work cannot be perfectly monitored by the other partners. His temptation to shirk instead of work is an example of moral-hazard problems that can be caused by systems of risk sharing and insurance. In such cases, this individual may have to bear a larger share of the risks to give him more motivation for diligent efforts on behalf of the partnership. We learned how to compute optimal linear and nonlinear sharing rules that satisfy the moral-hazard incentive constraints for such an agent.

We then moved from the study of partnerships to publicly held corporations with a stock market that includes many investors with constant risk tolerance. Among two financial assets that have the same marginal probability distribution of values next year, the one that has lower correlation with the overall market portfolio should have a higher value now because it offers better insurance against the risks of the diversified market portfolio that investors hold. We developed a spreadsheet model for computing an asset's current value in the stock market from the joint distribution of the future values of this asset and the market portfolio. This model yields asset values such that when expected values are computed using some market-adjusted probabilities for the possible outcomes, the expected growth ratio of every asset's value is equal to the return ratio on risk-free bonds or bank accounts. This market-adjusted probability distribution puts

Myerson_C08.qxd 07-09-2004 09:45 Page 323

relatively more weight on bad outcomes where the market portfolio does poorly because risk-averse investors are relatively more sensitive to income changes in such bad outcomes. The property of all assets having equal expected return ratios in terms of some shadow probability distribution is a general characteristic of any asset pricing system that does not create arbitrage opportunities.

## Exercises

**1.** G. Washington owns a real-estate development that will pay returns drawn from a Normal distribution, with expected value \$50,000 and standard deviation \$12,000. Washington has constant risk tolerance, and his risk-tolerance index is \$8000. J. Madison's risk-tolerance index is \$5000. (The development does not require any significant unobservable effort from Washington himself.)

**(a)** What is Washington's certainty equivalent for the real-estate development if he keeps the development for himself?

**(b)** Washington is planning to sell Madison a share in the development. To maximize their total certainty-equivalent values of their shares, what share should Washington sell to Madison? What is the minimum price that Washington should accept for selling this share? What is the maximum price that Madison would be willing to pay for this share?

**(c)** B. Arnold owns a real-estate development upriver that will pay returns drawn from a Normal distribution with expected value \$64,000 and standard deviation \$20,000. Arnold has offered to exchange his upriver development for Washington's development. Washington expects to take Madison on as a partner in his real-estate deals in any case. Should Washington accept Arnold's offer to exchange developments? Explain your answer.

**2.** *Case:* The Wilson Estate

Daniel Wilson and Rebecca Wilson Tisler have just inherited their mother's real-estate holdings, which consist of scattered commercial properties, some of which have substantial mortgages. Without their mother to manage the properties, they plan to sell the properties and split the returns equally.

A sudden offer by P. J. Cooney to buy all of their mother's properties for \$140,000 over the cost of repaying all mortgages has caused disagreement between the brother and sister, however. Daniel is eager to accept the offer, but Rebecca feels strongly that they should try to do better by selling the various properties separately over the coming year. You have been asked to help resolve their conflict.

After lengthy discussions with Daniel and Rebecca, you have found that they essentially agree about their prospective risks if they turn down Cooney's offer. Appraisals of the properties by independent real-estate agents have suggested that the expected returns from selling the properties separately is \$200,000 over the cost of repaying all mortgages, although their actual realized returns could be substantially above or below that amount. So both Daniel and Rebecca agree that a Normal distribution with this mean of \$200,000 and a standard deviation of \$75,000 can accurately describe their beliefs about their combined returns from selling the properties separately after rejecting Cooney's offer.

Where the siblings differ is in their risk tolerance, however. Daniel, with children about to enter college, is substantially less risk-tolerant than Rebecca. By asking them about their willingness to take hypothetical risks, you have assessed Daniel's risk-tolerance index to be \$20,000, and you have assessed Rebecca's risk-tolerance index to be \$45,000. It seems reasonable to assume that they each have constant risk tolerance.

**(a)** What is Daniel's certainty equivalent for a 50% share of the returns from selling the properties after rejecting Cooney's offer? What is Rebecca's certainty equivalent for a 50% share of the returns from selling the properties after rejecting Cooney's offer?

**(b)** Consider an alternative plan in which Rebecca pays her brother some amount of money now to buy a larger share of the returns from the properties. To increase the sum of their certainty equivalents as much as possible, what share

Myerson_C08.qxd 07-09-2004 09:45 Page 324

should Rebecca take? How much money would she have to pay Daniel now to buy this increased share, so that his certainty equivalent for his remaining share plus her payment to him should be equal to $70,000 (the amount that he would get if they sold to Cooney now and divided the money equally)?

**(c)** If Rebecca paid her brother and increased her share of the returns from selling the properties as you described in part (b), then what would be her certainty equivalent for her share of the returns minus her payment to her brother?

**(d)** Make a chart showing the (inverse) cumulative distribution of the net returns (including the payment to or from the other sibling) that Daniel and Rebecca will each get under your plan from part (b).

**3.** If the manager is diligent, then the gross profits that our company will earn from a new project will be a random variable drawn from a Normal distribution with mean $900,000 and standard deviation $200,000. But if the manager shirks, then these profits will be reduced by 20%. We cannot observe whether the manager is diligent or shirking, but we will be able to observe the gross profits that are generated. The manager has constant risk tolerance of $50,000. Shirking would be worth an additional $40,000 in compensation to the manager. The manager's pay may depend on the gross profit earned, but this pay cannot be less than $80,000 in any case, and the manager would quit now if the certainty equivalent of his compensation plan was less than $200,000. We want to maximize our company's expected net profit, after subtracting the amount that we pay the manager.

**(a)** Let us consider linear compensation plans, adjusted to the minimum wage of $80,000 where necessary. That is, suppose the wage we pay the manager will be some fixed constant plus a fixed fraction of the gross profits, or $80,000, whichever is larger. Find the compensation plan of this linear form that maximizes our expected net profit. Under this plan, what is the highest gross profit for which the manager gets only $80,000? What fraction of profits over this amount do we pay to the manager?

**(b)** If we consider nonlinear compensation plans, how much can you increase our expected net profit?

**(c)** Suppose instead that we compensate the manager in such a way that he does not quit but he shirks. What is the highest expected net profit that our company can earn from this project with a shirking manager?

**(d)** How would your optimal linear compensation plan in part (a) change if shirking reduced gross profits by 25%, but everything else is the same?

**4.** Consider again Part C of the Superior Semiconductor case, as described at the beginning of Chapter 4 and analyzed in Section 4.6. Let us assume, for simplicity, that all costs and revenues of this T-regulator product will accrue two years from now, but the interest rate on risk-free bonds is essentially zero. The value in two years of $1 that is invested now in the well-diversified market portfolio will be a Generalized Lognormal random variable with quartile boundary points $0.80, $1.10, and $1.50. Suppose that in the T-regulator project the development costs, the event of successful development, and the number of entering competitors are believed to be independent of the returns to the market portfolio. But the total value of the market for the T-regulator product in this case will be dependent on the same macroeconomic forces that will determine the returns to the market portfolio in the stock market. To be specific, suppose that Superior Semiconductor's marketing manager says that if she were told that the stock market portfolio would decrease to $0.80 in two years per dollar invested now, then she would revise her median value of the whole T-regulator market to $85 million.

**(a)** Assuming that investors in the stock market understood all the facts just described and in Part C of the Superior Semiconductor case, would you recommend that Superior Semiconductor develop the new T-regulator product? How would the total value of Superior Semiconductor's stock change if the company announced that it was developing the T-regulator device?

**(b)** How would your answer change if the marketing manager instead assessed a conditional median of $95 million for value of the T-regulator market given the stock market portfolio decreasing to $0.80 in two years?

Myerson_C09.qxd 07-09-2004 09:46 Page 325

# 9

# Dynamic Models of Growth and Arrivals

Dynamic models are models that describe the way that things change over time. In this chapter, we consider a variety of techniques for working with dynamic models. After reviewing the concept of net present value, we consider some simple models for forecasting how values may change and grow over time. Basic formulas of Brownian motion are introduced to make models of growth processes over shorter time intervals. Then, focusing on repeated investment processes where growth ratios are drawn independently across all periods from the same probability distribution, we show how a log-optimal investment strategy may offer the best likely performance of investments in the long run.

Next, we consider arrival processes, like the arrivals of customers into a shop, where the times between arrivals are independent Exponential random variables. The Exponential distribution has a convenient no-memory property for waiting times that simplifies the modeling of such arrival processes. Using this property of Exponential waiting times, we construct a simple model of a queuing system, to analyze the expected costs and benefits of adding an extra server in a shop where customers arrive at random. The role of the Poisson and Gamma probability distributions in such arrival processes is also introduced.

Finally, we consider the length of time that will be needed to complete a project that consists of many tasks, where some tasks may require others to be completed first, and the time required to complete each task is uncertain. To help manage such projects, we develop a model to estimate the probability that decreasing the time for any one task would decrease the final completion time for the whole project.

Myerson_C09.qxd 07-09-2004 09:46 Page 326

## 9.1 Net Present Value

Suppose that you had an investment opportunity that is expected to yield $10,000 of income every year for the next 10 years. Over the 10 years, you could expect to earn $100,000 from this investment. Should such an investment be worth $100,000 to you today? The answer is No, it should be worth less than $100,000, but how much less depends on the interest rate that you can get on your savings.

Suppose that by putting your money in the bank, you could earn 5% interest per year (after all compounding is taken into account). If you put $100,000 in the bank today, then at the end of the first year, you would have $105,000. You could then pay yourself $10,000 out of this bank account and still have $95,000 in the bank at the start of the second year. So at the end of the second year, you would have $95,000*(1.05) = $99,750, and you could then pay yourself another $10,000 and begin the third year with $89,750 in the bank. If you continue earning 5% interest on your remaining bank balance each year, then you go on paying yourself $10,000 at the end of the year in the coming decade, and you would end up with an extra $37,110 in the bank. So you should not pay $100,000 now for an investment that offers the same $10,000 a year for 10 years but leaves you nothing at the end! The actual value of this investment to you now should be the amount of money that if put in the bank now, would just enable you to take out the same $10,000 a year for 10 years and have nothing left at the end. Excel provides function NPV, which computes this value, where NPV stands for *net present value*. With the interest rate at 0.05, the investment paying $10,000 per year for 10 years has a net present value of $77,217.

The use of the NPV function is illustrated in Figure 9.1 for a stream of incomes that have been randomly generated from a Normal distribution with mean $10,000 and standard deviation $1000. The values in cells B5:B14 of Figure 9.1 represent the incomes to be earned in each of the next 10 years, at the year's end. With the interest rate 0.05 entered into cell A2, the net present value is computed in cell D3 by the formula

```
=NPV(A2,B5:B14)
```

The bank balance at the end of the first year, after paying the first year's income in cell B5, is computed in cell D5 by the formula

```
=D3*(1+$A$2)–B5
```

Similarly, the bank balance at the end of each subsequent year, after paying the income that is specified for this year in the B column, is computed in cells D6:D14 by entering the formula

```
=D5*(1+$A$2)–B6
```

into cell D6 and then copying D6 to D6:D14. Notice that with the initial deposit that is specified by the NPV function, the final balance after paying

Myerson_C09.qxd 07-09-2004 09:46 Page 327

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | Annual interest rate | | To get Incomes from bank: | | | | |
| 2 | 0.05 | | | Deposit NPV now | | NetPresentValue | |
| 3 | | | | 77571.22 | | 77571.22 | |
| 4 | FutureYear | Income | | Bank balance | | 1/(1+i)^Year | |
| 5 | 1 | 10738.63 | | 70711.15 | | 0.9524 | |
| 6 | 2 | 10620.81 | | 63625.89 | | 0.9070 | |
| 7 | 3 | 9882.02 | | 56925.17 | | 0.8638 | |
| 8 | 4 | 10108.55 | | 49662.88 | | 0.8227 | |
| 9 | 5 | 9042.25 | | 43103.77 | | 0.7835 | |
| 10 | 6 | 11177.71 | | 34081.25 | | 0.7462 | |
| 11 | 7 | 8275.17 | | 27510.14 | | 0.7107 | |
| 12 | 8 | 9827.17 | | 19058.48 | | 0.6768 | |
| 13 | 9 | 11336.29 | | 8675.11 | | 0.6446 | |
| 14 | 10 | 9108.86 | | 0.00 | | 0.6139 | |
| 15 | | | | | | | |
| 16 | Letting money accumulate in bank, we have at end of each year: | | | | | | |
| 17 | FutureYear | From Income | | From NPV deposited now | | | |
| 18 | 1 | 10738.63 | | 81449.78 | | | |
| 19 | 2 | 21896.37 | | 85522.27 | | | |
| 20 | 3 | 32873.21 | | 89798.38 | | | |
| 21 | 4 | 44625.42 | | 94288.30 | | | |
| 22 | 5 | 55898.94 | | 99002.71 | | | |
| 23 | 6 | 69871.60 | | 103952.85 | | | |
| 24 | 7 | 81640.35 | | 109150.49 | | | |
| 25 | 8 | 95549.54 | | 114608.02 | | | |
| 26 | 9 | 111663.31 | | 120338.42 | | | |
| 27 | 10 | 126355.34 | | 126355.34 | | | |
| 28 | and at any year thereafter | | | | | | |
| 29 | FutureYear | From Income | | From NPV deposited now | | | |
| 30 | 15 | 161264.99 | | 161264.99 | | | |
| 31 | | | | | | | |
| 32 | FORMULAS | | | | | | |
| 33 | B5. =NORMINV(RAND(),10000,1000) | | | | B18. =B5 | | |
| 34 | B5 copied to B5:B14 | | | | D18. =D3*(1+$A$2) | | |
| 35 | D3. =NPV(A2,B5:B14) | | | | B19. =B18*(1+$A$2)+B6 | | |
| 36 | D5. =D3*(1+$A$2)-B5 | | | | D19. =D18*(1+$A$2) | | |
| 37 | D6. =D5*(1+$A$2)-B6 | | | | B19:D19 copied to B19:D27 | | |
| 38 | D6 copied to D6:D14 | | | | B30. =B27*(1+$A$2)^(A30-A27) | | |
| 39 | F5. =1/(1+$A$2)^A5 | | | | D30. =D3*(1+$A$2)^A30 | | |
| 40 | F5 copied to F5:F14 | | | | | | |
| 41 | F3. =SUMPRODUCT(B5:B14,F5:F14) | | | | | | |

FIGURE **9.1**
**Net present value of a future income stream**

out the last income is exactly \$0. So the net present value NPV(rate, incomes) is the amount of money that if deposited now at the given interest rate, would be just enough to pay ourselves the given stream of future incomes.

The formula for computing the net present value is shown in the F column of Figure 9.1. The essential idea is that when banks pay the interest rate $i$ per year, a future payment of \$1 to be delivered $t$ years from now can be effectively purchased now by depositing $\$1/(1 + i)\text{^}t$ in the bank now. These current prices for buying dollars at the end of future years are computed in cells F5:F14. Then the net present value of the whole stream of incomes in B5:B14 is computed in cell F3 by the formula =SUMPRODUCT(B5:B14,F5:F14). Notice that the NPV function in cell D3 yields the same net present value.

Suppose now that we will let our incomes over some interval of future years accumulate in a bank account that pays an interest rate $i$. Then, at any year $t$ after the end of the interval when these incomes are earned, the

Myerson_C09.qxd 07-09-2004 09:46 Page 328

total amount in our bank account will be the same as we could get from simply depositing the net present value of the future income stream today. This fact is illustrated in rows 16 to 30 of Figure 9.1. In this case, we have an income stream over the next 10 years that has the net present value $NPV = \$77{,}571.22$ at the interest rate $i = 0.05$. So for any year $t$ where $t > 10$, the accumulated value of our bank account from these incomes would have the value $NPV*(1 + i)\text{^}t$. For the case of $t = 15$, Figure 9.1 shows that this formula in cell D30 yields the same value as the accumulated bank balance from the income stream in cell B30.

So for any two income streams, the one with the higher net present value will ultimately provide the greater monetary value, if we want to spend these incomes at any time after the intervals over which they are earned and can deposit them in the bank until that time. For this reason, economists recommend that future income streams should be evaluated according to their net present value, evaluated at an appropriate interest rate.

## 9.2 Forecasting Models

Whenever we make a projected forecast from past data, we must assume that in some ways the future will be like the past. So we generally assume that there are various quantities that have been and will be drawn independently from the same probability distribution every year. Such random variables are called *stationary* because, although they may take different values each year, the distribution from which they are drawn stays fixed over time. For each of these stationary random variables, we can statistically estimate the parameters of its probability distribution from our past data, and we can then forecast its future values by computer simulations drawn from this estimated distribution. The crucial modeling question is, which quantities should be stable enough to be projected in this way?

Consider again the data for Funds 1 and 2 from the 1980s that we saw in Chapter 5. Suppose that (from a position in 1990) we were trying to make a forecast of the values of these funds in the next few years, based on this given data. One (bad) assumption that we might make is that the value of each fund is such a stationary random variable, independently drawn each year from the same distribution. In our past data, the end-of-year values of Fund 1 had a mean of \$127 and a standard deviation of \$52. So we might project future end-of-year values of Fund 1 as independent draws from a Normal or Lognormal distribution with mean \$127 and standard deviation \$52. But most people would consider this model to be absurd, because the funds' values in the data show a clear upward trend. The values below the mean all occurred in the first 7 years of our data, and the values above the mean all occurred in the latter 4 years, so we may anticipate that values in the near future will be closer to the recent high values than the

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | End-of-year price data for two stock funds | | | | | | | |
| 2 | | Fund 1 | Fund 2 | | Statistics for fund values | | | |
| 3 | 1979 | 64.76 | 47.34 | | Fund 1 | Fund 2 | | |
| 4 | 1980 | 78.70 | 60.61 | | 127 | 99 | Mean | |
| 5 | 1981 | 85.44 | 72.61 | | 52 | 40 | Stdev | |
| 6 | 1982 | 78.18 | 60.41 | | | | | |
| 7 | 1983 | 107.45 | 89.36 | FORMULAS | | | | |
| 8 | 1984 | 108.01 | 85.63 | E4. =AVERAGE(B3:B13) | | | | |
| 9 | 1985 | 123.79 | 104.11 | E5. =STDEV(B3:B13) | | | | |
| 10 | 1986 | 155.85 | 119.87 | E4:E5 copied to F4:F5 | | | | |
| 11 | 1987 | 195.31 | 140.39 | B17. =LN(B4/B3) | | | | |
| 12 | 1988 | 180.95 | 134.12 | B17 copied to B17:B26 | | | | |
| 13 | 1989 | 216.23 | 175.28 | B28. =AVERAGE(B17:B26) | | | | |
| 14 | | | | B29. =STDEV(B17:B26) | | | | |
| 15 | Logarithmic growth rates | | | B28:B29 copied to C28:C29 | | | | |
| 16 | | Fund 1 | Fund 2 | B30. =CORREL(B17:B26,C17:C26) | | | | |
| 17 | 1980 | 0.19 | 0.25 | B34. =B13 | | | | |
| 18 | 1981 | 0.08 | 0.18 | C34. =C13 | | | | |
| 19 | 1982 | -0.09 | -0.18 | D35:E35. {=CORAND($B$30)} | | | | |
| 20 | 1983 | 0.32 | 0.39 | B35. =B34*EXP(NORMINV(D35,B$28,B$29)) | | | | |
| 21 | 1984 | 0.01 | -0.04 | C35. =C34*EXP(NORMINV(E35,C$28,C$29)) | | | | |
| 22 | 1985 | 0.14 | 0.20 | B35:E35 copied to B36:E40 | | | | |
| 23 | 1986 | 0.23 | 0.14 | | | | | |
| 24 | 1987 | 0.23 | 0.16 | | | | | |
| 25 | 1988 | -0.08 | -0.05 | | | | | |
| 26 | 1989 | 0.18 | 0.27 | | | | | |
| 27 | Growth rate statistics | | | | | | | |
| 28 | Means | 0.121 | 0.131 | | | | | |
| 29 | Stdevns | 0.137 | 0.173 | | | | | |
| 30 | Correl | 0.9021 | | | | | | |
| 31 | | | | | | | | |
| 32 | Simulation Model | | | | | | | |
| 33 | Year | Fund 1 | Fund 2 | | | | | |
| 34 | *1989* | *216.23* | *175.28* | (corands) | | | | |
| 35 | 1990 | 297.35 | 255.26 | 0.9260 | 0.9221 | | | |
| 36 | 1991 | 383.59 | 332.08 | 0.8364 | 0.7782 | | | |
| 37 | 1992 | 388.25 | 298.51 | 0.2140 | 0.0844 | | | |
| 38 | 1993 | 455.70 | 383.64 | 0.6139 | 0.7566 | | | |
| 39 | 1994 | 458.52 | 447.40 | 0.2016 | 0.5527 | | | |

FIGURE **9.2**
**Forecasting model for two stock funds**

older low values. Thus, the funds' values do not appear to have been all drawn from one stationary distribution over the 11 years shown in our data.

But if we compute these fund's logarithmic growth rates, then we find a more stable pattern, which is shown in cells B17:C26 of Figure 9.2. (Recall from Section 4.2 that a logarithmic growth rate is computed from the growth ratio by the natural logarithm function LN.) For example, the annual growth rates of Fund 1 have a mean of 0.121 and a standard deviation of 0.137 in this data, and growth rates above and below the mean can be found both early and late in the data series. So, as there is no obvious difference between the growth rates earlier and later in the series, we could reasonably assume that Fund 1's annual logarithmic growth rate is a stationary random variable, drawn independently each year from a Normal distribution with mean 0.121 and standard deviation 0.137.

Although we may assume that the funds' growth rates are independent across years, the high correlation between the two funds' annual growth

Myerson_C09.qxd 07-09-2004 09:46 Page 330

rates in this data (0.9021 in cell B30 of Figure 9.2) implies that we should not assume that the two funds' growth rates are independent of each other in any one year. So we may reasonably assume that the annual logarithmic growth rates of Funds 1 and 2 have a joint probability distribution that is Multivariate Normal, with the means, standard deviations, and correlation that we have found in our historical data (shown in cells B28:C29 and C30 of Figure 9.2). Under this assumption, the growth ratios and resulting values of the two Funds become Lognormal random variables, as shown in cells B35:C39 of Figure 9.2. (Recall from Section 4.2 that the growth ratio is computed from the logarithmic growth rate by the EXP function.)

**NOTE** (Notice that the distribution assumption used here is slightly different from the distributional assumption that we used when we analyzed the same data in Chapter 5. For short-term one-period portfolio models like those in Chapter 5, there are some technical advantages to assuming that growth ratios are Multivariate Normal random variables, as was discussed in Section 5.6. But for the multiperiod dynamic models that we are considering here, it is generally better to assume instead that logarithmic growth rates are Multivariate Normal random variables, which makes the growth ratios into Lognormal random variables, as was discussed in Section 4.3. Because of the small standard deviations here, there is very little practical difference between the Normal and Lognormal assumptions for this example.)

The forecasting model in Figure 9.2 is built purely by statistical analysis of past data. But most applications should involve a combination of statistical analysis and subjective assessment. For example, consider the *Global Sports* case that appears as Exercise 5 at the end of Chapter 4. A simulation model of this case is shown in Figure 9.3.

This case concerns the publication of a new magazine, but it could be about any new product. We are given data about the price at which this new product will be sold ($25 per unit), the variable cost of production ($8 per unit), and the fixed costs of producing any positive amount of this new product ($4 million in the first year, $3 million in the second year, and $2 million in each year thereafter). Cells B11:C19 in Figure 9.3 contain given estimates of the mean and standard deviation of the logarithmic growth rate of sales for each year after the first year, which are based on historical data (in columns I:O) about the growth rates of other similar products during their first decade after introduction. That is, in the sales data for the other magazines that have been previously introduced by this publishing company, the logarithmic growth rates of sales in the second year have had mean 0.50 and standard deviation 0.30, the growth rates in the third year have had mean 0.30 and standard deviation 0.20, and so on. The simulated growth rates for this proposed *Global Sports* magazine in cells D11:D19 of Figure 9.3 are based on the assumption that, in each year, the growth rate of sales of any new magazine will be drawn independently out of a fixed

Myerson_C09.qxd 07-09-2004 09:46 Page 331

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | GLOBAL SPORTS case | | | | FixedCost/Year ($millions) | | | | Sales data for six previously | | | | | | |
| 2 | SellingPrice($) | | VariableCost($) | | 4 | year 1 | | | introduced magazines: | | | | | | |
| 3 | 25 | | 8 | | 3 | year 2 | | | | Subscriptions sold (1000s) | | | | | |
| 4 | | | | | 2 | later years | | | Year | Mag1 | Mag2 | Mag3 | Mag4 | Mag5 | Mag6 |
| 5 | | | | | | | | | 1 | 449 | 230 | 35 | 143 | 25 | 303 |
| 6 | Subjectively assessed quartiles for quantity sold in year 5 | | | | | | | | 2 | 661 | 494 | 62 | 349 | 33 | 333 |
| 7 | 50000 | 100000 | 200000 | | | | | | 3 | 852 | 609 | 62 | 599 | 55 | 447 |
| 8 | | | | | | | | | 4 | 881 | 692 | 73 | 834 | 61 | 688 |
| 9 | Year | LogGrowRate of Sales | | | #Sold | Profit ($million) | | | 5 | 1003 | 765 | 87 | 988 | 61 | 705 |
| 10 | 1 | E | Stdev | Sim'd | 59664 | -2.986 | | | 6 | 1001 | 719 | 88 | 967 | 63 | 746 |
| 11 | 2 | 0.50 | 0.30 | 0.067 | 63802 | -1.915 | | | 7 | 976 | 684 | 87 | 1002 | 66 | 738 |
| 12 | 3 | 0.30 | 0.20 | 0.488 | 103961 | -0.233 | | | 8 | 925 | 672 | 88 | 991 | 67 | 789 |
| 13 | 4 | 0.20 | 0.15 | 0.117 | 116866 | -0.013 | | | 9 | 929 | 700 | 88 | 923 | 69 | 783 |
| 14 | 5 | 0.10 | 0.07 | 0.068 | 125051 | 0.126 | | | 10 | 925 | 649 | 89 | 933 | 72 | 798 |
| 15 | 6 | 0.00 | 0.04 | 0.036 | 129588 | 0.203 | | | | | | | | | |
| 16 | 7 | 0.00 | 0.04 | -0.025 | 126346 | 0.148 | | | | Logarithmic growth rates | | | | | |
| 17 | 8 | 0.00 | 0.04 | 0.003 | 126688 | 0.154 | | | Year | Mag1 | Mag2 | Mag3 | Mag4 | Mag5 | Mag6 |
| 18 | 9 | 0.00 | 0.04 | 0.013 | 128287 | 0.181 | | | 2 | 0.39 | 0.76 | 0.57 | 0.89 | 0.29 | 0.1 |
| 19 | 10 | 0.00 | 0.04 | 0.029 | 132044 | 0.245 | | | 3 | 0.25 | 0.21 | -0 | 0.54 | 0.5 | 0.29 |
| 20 | | | | | | | | | 4 | 0.03 | 0.13 | 0.17 | 0.33 | 0.11 | 0.43 |
| 21 | Interest rate | | NPV of profits ($million) | | | | | | 5 | 0.13 | 0.1 | 0.17 | 0.17 | 0.01 | 0.03 |
| 22 | 0.1 | | -3.970 | if continue all 10 years | | | | | 6 | -0 | -0.1 | 0.01 | -0 | 0.02 | 0.06 |
| 23 | | | -2.714 | if stop at year 1 | | | | | 7 | -0 | -0 | -0 | 0.04 | 0.06 | -0 |
| 24 | FORMULAS | | (assuming profits accrue at end of year) | | | | | | 8 | -0.1 | -0 | 0.01 | -0 | 0 | 0.07 |
| 25 | J18. =LN(J6/J5) | | | | | | | | 9 | 0 | 0.04 | -0 | -0.1 | 0.03 | -0 |
| 26 | J18 copied to J18:O26 | | | | E13. =E14/EXP(D14) | | | | 10 | -0 | -0.1 | 0.01 | 0.01 | 0.04 | 0.02 |
| 27 | B11. =AVERAGE(J18:O18) | | | | E13 copied to E10:E12 | | | | | | | | | | |
| 28 | C11. =STDEV(J18:O18) | | | | F10. =($A$3-$C$3)*E10/1000000-E2 | | | | | | | | | | |
| 29 | D11. =NORMINV(RAND(),B11,C11) | | | | F11. =($A$3-$C$3)*E11/1000000-E3 | | | | | | | | | | |
| 30 | B11:D11 copied to B11:D19 | | | | F12. =($A$3-$C$3)*E12/1000000-$E$4 | | | | | | | | | | |
| 31 | E14. =GENLINV(RAND(),A7,B7,C7) | | | | F12 copied to F12:F19 | | | | | | | | | | |
| 32 | E15. =E14*EXP(D15) | | | | C22. =NPV(A22,F10:F19) | | | | | | | | | | |
| 33 | E15 copied to E15:E19 | | | | C23. =NPV(A22,F10) | | | | | | | | | | |

FIGURE **9.3** ***Global Sports* case**

Normal probability distribution, for which the mean and standard deviation can be estimated from this past data. That is, we are assuming that the annual growth rates for this new *Global Sports* magazine would be similar to the growth rates of all other magazines for each year after introduction.

Although this assumption of similarity of growth rates may be reasonable, it could be very unreasonable to assume that the first-year sales for this proposed new magazine should be similar to the first-year sales of all other magazines that this company has introduced. If this company has introduced some outstanding new magazines in the past, these past successes should not lead managers to assume that any new magazine proposal will be just as outstanding. The proposed theme of this new *Global Sports* magazine is different from any other magazine that this company has ever introduced. Somewhere in the analysis, the decision makers at this company should take account of whether this theme seems likely to appeal to many potential readers or few. To do so, they could ask the relevant experts on the marketing staff to subjectively estimate a probability distribution for *Global Sport's* total sales in any one year, taking account of the unique features of this proposed magazine and the various types of people

Myerson_C09.qxd 07-09-2004 09:46 Page 332

who might be interested in it. Once the level of sales in any one year has been estimated in this way, then the sales in other years can be generated by the assumed growth rates.

In this case, it is assumed that the experts on the marketing staff are most comfortable thinking about the potential sales of this magazine in its fifth year of publication. (This later date may enable the marketing experts to ignore uncertainty about how long the magazine may take to catch on in the target populations where it may have appeal.) We are told that the marketing experts have subjectively assessed quartile boundary points of 50,000, 100,000, and 200,000 for the total number of subscriptions that may be sold in the fifth year. These subjectively assessed quartiles are applied with the Generalized Lognormal distribution to yield simulated year 5 demand for the magazine in cell E14 of Figure 9.3. Then sales after year 5 are simulated in cells E15:E19 by multiplying the previous year's sales by the growth ratio (which is EXP of the logarithmic growth rate). Similarly, sales before year 5 are simulated in cells E10:E13 by dividing the following year's sales by that year's growth ratio. Thus, cells E10:E19 in Figure 9.3 contain simulated sales quantities for each year such that the growth rates in each year are random variables with distributions estimated from past data about other magazines. But the overall level of sales uses the subjectively assessed quartiles for year 5 that take account of the unique potential of this proposed magazine.

## 9.3 Forecasting Example: Goeing Case

Let us consider now another, more complex example. Suppose that a major aerospace manufacturer called Goeing Aircraft has recently had net sales and operating income as shown in cells A2:C20 of Figure 9.4. As viewed at the end of 1997, Goeing's most recent annual net sales were $6804 million and Goeing's most recent annual operating income was $704 million, as shown in cells B20:C20. Suppose that Goeing's decision makers at this time (1997) want a model to forecast the operating income over the next 10 years. They may be particularly interested in the net present value of their future income, discounted (say) at a 5% interest rate. But to add another question for this case, suppose also Goeing is planning to borrow money to buy back some of its own stock, and this new debt will commit Goeing to fixed annual payments over the next 10 years. The chairman of the board wants to have a probability of 0.90 that Goeing's annual operating income will never be less than these annual debt payments during any of the next 10 years. We want to estimate the largest annual debt payment that Goeing could undertake and still meet the chairman's criterion.

In this section, we consider three forecasting models to generate such estimates, all based on the same data but making different assumptions about what patterns from this data can be extrapolated into the future.

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | GOEING AIRCRAFT PAST DATA ($MILLIONS) | | | | | | | | |
| 2 | Year | Sales | Income | | Income | | | | |
| 3 | 1980 | 2439 | 207 | | LogGrowthRate | | | | |
| 4 | 1981 | 2828 | 262 | | 0.236 | | | | |
| 5 | 1982 | 2936 | 215 | | -0.198 | | | | |
| 6 | 1983 | 2711 | 172 | | -0.223 | | | | |
| 7 | 1984 | 3339 | 209 | | 0.195 | | | | |
| 8 | 1985 | 3106 | 216 | | 0.033 | | | | |
| 9 | 1986 | 4091 | 291 | | 0.298 | | | | |
| 10 | 1987 | 4902 | 328 | | 0.120 | | | | |
| 11 | 1988 | 4607 | 213 | | -0.432 | | | | |
| 12 | 1989 | 5089 | 305 | | 0.359 | | | | |
| 13 | 1990 | 6083 | 362 | | 0.171 | | | | |
| 14 | 1991 | 8279 | 662 | | 0.604 | | | | |
| 15 | 1992 | 8794 | 834 | | 0.231 | | | | |
| 16 | 1993 | 9055 | 900 | | 0.076 | | | | |
| 17 | 1994 | 7631 | 815 | | -0.099 | | | | |
| 18 | 1995 | 6577 | 688 | | -0.169 | | | | |
| 19 | 1996 | 5855 | 581 | | -0.169 | | | | |
| 20 | 1997 | 6804 | 704 | | 0.192 | | | | |
| 21 | | | | | | | | | |
| 22 | | | | Mean | 0.072 | | | | |
| 23 | | | | StDev | 0.259 | | | | |
| 24 | | | | | | | | | |
| 25 | FORECAST MODEL 1 (extrapolating Income growth) | | | | | | | | |
| 26 | Year | Income | | | | | | | |
| 27 | *1997* | *704* | *current* | | | | | | |
| 28 | 1998 | 816 | | FORMULAS | | | | | |
| 29 | 1999 | 770 | | E4. =LN(C4/C3) | | | | | |
| 30 | 2000 | 938 | | E4 copied to E4:E20 | | | | | |
| 31 | 2001 | 826 | | E22. =AVERAGE(E4:E20) | | | | | |
| 32 | 2002 | 1637 | | E23. =STDEV(E4:E20) | | | | | |
| 33 | 2003 | 1131 | | B28. =B27*EXP(NORMINV(RAND(),$E$22,$E$23)) | | | | | |
| 34 | 2004 | 996 | | B28 copied to B28:B37 | | | | | |
| 35 | 2005 | 1375 | | B42. =NPV(0.05,B28:B37) | | | | | |
| 36 | 2006 | 1144 | | C42. =MIN(B28:B37) | | | | | |
| 37 | 2007 | 1062 | | B39. =AVERAGE(B43:B543) | | | | | |
| 38 | | | | B40. =PERCENTILE(B43:B543,0.1) | | | | | |
| 39 | E | 9975 | 640 | B39:B40 copied to C39:C40 | | | | | |
| 40 | 10%ile | 4515 | 361 | | | | | | |
| 41 | NPV(.05) | | MinIncome | | | | | | |
| 42 | SimTab: | 8118 | 770 | | | | | | |
| 43 | 0 | 5806 | 583 | | | | | | |
| 44 | 0.002 | 8380 | 556 | | | | | | |
| 45 | 0.004 | 13755 | 659 | | | | | | |
| 46 | 0.006 | 13959 | 683 | | | | | | |

FIGURE **9.4**
**Goeing case, with Forecast Model 1 (extrapolating income growth only)**

In our first and simplest model, shown in Figure 9.4, we ignore the sales data and try to extrapolate the given income data in cells C3:C20. Figures 9.5 and 9.6 show two more sophisticated models that take account of the given sales data in cells B3:B20 as well as the income data.

The given annual incomes from 1980 to 1997 have mean $442 million and standard deviation $258 million. But it would be very unreasonable to assume that annual incomes are stationary independent draws from a Normal distribution with this mean and standard deviation, because the incomes below this mean were all in the first 11 years of the data, and the incomes above this mean were all in the last 7 years. So instead, our first forecasting model extrapolates the annual logarithmic growth rates of

Myerson_C09.qxd 07-09-2004 09:46 Page 334

income, as calculated in cells E4:E20 of Figure 9.4. The annual income growth rates from 1981 to 1997 have mean 0.072 (calculated in cell E22) and standard deviation 0.259 (in cell E23), and we can find growths above and below this mean both early and late in the historical data. Starting with the most recent annual income in cell B27, projected incomes for the next 10 years are simulated in cells B28:B37 of Figure 9.4 by entering the formula

=B27*EXP(NORMINV(RAND( ),$E$22,$E$23))

into cell B28, and then copying B28 to B28:B37. In this model, annual income is assumed to have logarithmic growth rates that are stationary random variables, drawn independently each year from a Normal distribution with the mean and standard deviation that has been observed in our historical data. With this assumption, the projected annual incomes in cells B28:B37 are Lognormal random variables. Then the net present value of these simulated future incomes, at a 5% interest rate, is computed in cell B42 by the formula

=NPV(0.05,B28:B37)

The minimum annual income in this simulated forecast is computed in cell C42 by the formula

=MIN(B28:B37)

The spreadsheet in Figure 9.4 includes a simulation table in cells A42:C543 containing net present values and minimum incomes from 501 simulations of this model. The expected value and 0.1 percentile of the net present values are estimated from this simulation data in cells B39 and B40 by the formulas

=AVERAGE(B43:B543) and =PERCENTILE(B43:B543,0.1)

Similarly, estimates of the expected value and 0.10 percentile of the minimum annual income over the coming decade are computed in cells C39 and C40 (by copying the formulas from B39:B40). This analysis suggests that the net present value of Goeing's incomes over the coming decade may have an expected value of $9975 million (in cell B39), but Goeing's annual debt burden should not exceed $361 million (in cell C40) if we want a probability of 0.1 of having any year where the annual income is less than the debt burden.

Figure 9.5 shows a second forecasting model that is based on the same historical data in cells A2:C20 but instead extrapolates the observed growth of sales from this data. The annual logarithmic growth rates of sales from 1981 to 1997 have mean 0.060 (calculated in cell E22 of Figure 9.5) and standard deviation 0.149 (in cell E23), and we can find sales growth rates above and below this mean both early and late in the historical data. Starting with the most recent annual sales in cell B27, projected sales for

| | A | B | C | D | E | F | G | H | I |
|---|---|---|---|---|---|---|---|---|---|
| 1 | GOEING AIRCRAFT PAST DATA ($MILLIONS) | | | | | | | | |
| 2 | Year | Sales | Income | | Sales | Income/Sales | | | |
| 3 | 1980 | 2439 | 207 | LogGrowthRate | | ratio | | | |
| 4 | 1981 | 2828 | 262 | | 0.148 | 0.0926 | | | |
| 5 | 1982 | 2936 | 215 | | 0.037 | 0.0732 | | | |
| 6 | 1983 | 2711 | 172 | | -0.080 | 0.0634 | | | |
| 7 | 1984 | 3339 | 209 | | 0.208 | 0.0626 | | | |
| 8 | 1985 | 3106 | 216 | | -0.072 | 0.0695 | | | |
| 9 | 1986 | 4091 | 291 | | 0.275 | 0.0711 | | | |
| 10 | 1987 | 4902 | 328 | | 0.181 | 0.0669 | | | |
| 11 | 1988 | 4607 | 213 | | -0.062 | 0.0462 | | | |
| 12 | 1989 | 5089 | 305 | | 0.100 | 0.0599 | | | |
| 13 | 1990 | 6083 | 362 | | 0.178 | 0.0595 | | | |
| 14 | 1991 | 8279 | 662 | | 0.308 | 0.0800 | | | |
| 15 | 1992 | 8794 | 834 | | 0.060 | 0.0948 | | | |
| 16 | 1993 | 9055 | 900 | | 0.029 | 0.0994 | | | |
| 17 | 1994 | 7631 | 815 | | -0.171 | 0.1068 | | | |
| 18 | 1995 | 6577 | 688 | | -0.149 | 0.1046 | | | |
| 19 | 1996 | 5855 | 581 | | -0.116 | 0.0992 | | | |
| 20 | 1997 | 6804 | 704 | | 0.150 | 0.1035 | | | |
| 21 | | | | | | | | | |
| 22 | | | | Mean | 0.060 | 0.0796 | | | |
| 23 | | | | StDev | 0.149 | 0.0193 | | | |
| 24 | | | 1YearAutocorrel | | 0.1394 | 0.8279 | | | |
| 25 | FORECAST MODEL 2 (extrapolating Sales growth and Income/Sales) | | | | | | | | |
| 26 | Year | Sales | Inc/S | Income | | | | | |
| 27 | *1997* | *6804* | *.1035* | *704* | *current* | | | | |
| 28 | 1998 | 6847 | 0.0999 | 684 | | | | | |
| 29 | 1999 | 7741 | 0.0714 | 553 | | | | | |
| 30 | 2000 | 9322 | 0.0896 | 835 | | | | | |
| 31 | 2001 | 10086 | 0.0693 | 699 | FORMULAS | | | | |
| 32 | 2002 | 12059 | 0.0689 | 831 | E4. =LN(B4/B3) | | | | |
| 33 | 2003 | 11572 | 0.1196 | 1384 | F4. =C4/B4 | | | | |
| 34 | 2004 | 13085 | 0.0746 | 976 | E4:F4 copied to E4:F20 | | | | |
| 35 | 2005 | 17613 | 0.0809 | 1426 | E22. =AVERAGE(E4:E20) | | | | |
| 36 | 2006 | 15578 | 0.0639 | 995 | E23. =STDEV(E4:E20) | | | | |
| 37 | 2007 | 16642 | 0.0956 | 1591 | E24. =CORREL(E4:E19,E5:E20) | | | | |
| 38 | | | | | E22:E24 copied to F22:F24 | | | | |
| 39 | E | 6212 | 424 | | B28. =B27*EXP(NORMINV(RAND(),$E$22,$E$23)) | | | | |
| 40 | 10%ile | 4014 | 259 | | C28. =NORMINV(RAND(),$F$22,$F$23) | | | | |
| 41 | NPV(.05) | | MinIncome | | D28. =B28*C28 | | | | |
| 42 | SimTable | 7410 | 553 | | B28:D28 copied to B28:D37 | | | | |
| 43 | 0 | 5426 | 310 | | B42. =NPV(0.05,D28:D37) | | | | |
| 44 | 0.002 | 4696 | 182 | | C42. =MIN(D28:D37) | | | | |
| 45 | 0.004 | 6073 | 472 | | B39. =AVERAGE(B43:B543) | | | | |
| 46 | 0.006 | 4312 | 323 | | B40. =PERCENTILE(B43:B543,0.1) | | | | |
| 47 | 0.008 | 6816 | 328 | | B39:B40 copied to C39:C40 | | | | |

FIGURE **9.5**
**Goeing case, Forecast Model 2 (extrapolating sales growth and income/sales)**

the next 10 years are simulated in cells B28:B37 of Figure 9.5 by entering the formula

=B27∗EXP(NORMINV(RAND( ),$E$22,$E$23))

into cell B28, and then copying B28 to B28:B37. Then to generate a forecast of future incomes from this forecast of future sales, we use the fact that the ratios of income over sales in the years from 1981 to 1997 had a mean of 0.0796 (calculated in cell F22 of Figure 9.5) and a standard deviation of 0.0193 (in cell F23). So we may simulate the ratios of income over sales in

Myerson_C09.qxd 07-09-2004 09:46 Page 336

future years as independent Normal random variables with this mean and standard deviation by entering the formula

=NORMINV(RAND( ),$F$22,$F$23)

into cell C28 and copying C28 to C28:C37. Then our simulated forecast of annual incomes in the next decade can be computed in cells D28:D37 of Figure 9.5 by entering the formula

=B28*C28

(sales multiplied by the income/sales ratio) into cell D28 and then copying it to cells D28:D37.

The spreadsheet in Figure 9.5 includes a simulation table that contains net present values and minimum incomes from 501 simulations of this second model. This simulation table suggests that the net present value of Goeing's incomes over the coming decade may have an expected value of $6212 million, but Goeing's annual debt burden should not exceed $259 million if we want a probability of 0.1 of having any year where the annual income is less than the debt burden. Notice that these estimates are lower than those that we got from our first model in Figure 9.4. This difference may be due to the fact that the second model is based on an extrapolation of the observed growth rate of sales, which is lower than the observed growth rate of incomes that was the basis for our extrapolations in the first model. By using this lower projected growth rate, the second model generates more conservative estimates of Goeing's future incomes.

The second model in Figure 9.5 uses the assumption that the sales growth rates and the income/sales ratios are independent across different years. But do these assumptions seem justified in the historical data? In cell F24 of Figure 9.5, the formula

=CORREL(F4:F19,F5:F20)

computes the correlation of the income/sales ratios from 1981 to 1996 with the income/sales ratios of the subsequent year. This statistic may be called the *autocorrelation* of the observed series of values. If every year's data were independent of the previous year's data, then we would expect to find an autocorrelation near zero. Because we find a relatively high autocorrelation of 0.8279 for the observed series of income/sales ratios here, we may want to investigate a third model, in which the income/sales ratios do not change so much from one year to another. Indeed, notice that the income/sales ratios in cells F4:F20 of Figure 9.5 seem to drift gradually over time, with relatively small changes from one year to the next. On the other hand, the autocorrelation of the observed sales growth rates (0.1394 in cell E24) is not so large, and so we may reasonably maintain the assumption that the sales growth rates are independent across different years.

Figure 9.6 shows such a third model. Cells G2:G8 here contain the array formula

{=REGRESSN(F4:F19,F5:F20)}

| | A | B | C | D | E | F | G | H | I | J |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | GOEING AIRCRAFT PAST DATA ($MILLIONS) | | | | | | (I/S) | (I/S-previous) | | |
| 2 | Year | Sales | Income | | Sales | Inc/S | Regressn Array | | | |
| 3 | 1980 | 2439 | 207 | LogGrowthRate | | ratio | X(1) Coefficient | | | |
| 4 | 1981 | 2828 | 262 | | 0.148 | 0.0926 | 0.8602 | | | Residuals |
| 5 | 1982 | 2936 | 215 | | 0.037 | 0.0732 | Intercept | | | -0.018 |
| 6 | 1983 | 2711 | 172 | | -0.080 | 0.0634 | 0.0116 | | | -0.011 |
| 7 | 1984 | 3339 | 209 | | 0.208 | 0.0626 | Std Error | | | -0.004 |
| 8 | 1985 | 3106 | 216 | | -0.072 | 0.0695 | 0.0114 | | | 0.004 |
| 9 | 1986 | 4091 | 291 | | 0.275 | 0.0711 | | | | 0.000 |
| 10 | 1987 | 4902 | 328 | | 0.181 | 0.0669 | | | | -0.006 |
| 11 | 1988 | 4607 | 213 | | -0.062 | 0.0462 | Statnry I/S | | | -0.023 |
| 12 | 1989 | 5089 | 305 | | 0.100 | 0.0599 | 0.083 | | | 0.009 |
| 13 | 1990 | 6083 | 362 | | 0.178 | 0.0595 | | | | -0.004 |
| 14 | 1991 | 8279 | 662 | | 0.308 | 0.0800 | | | | 0.017 |
| 15 | 1992 | 8794 | 834 | | 0.060 | 0.0948 | | | | 0.014 |
| 16 | 1993 | 9055 | 900 | | 0.029 | 0.0994 | | | | 0.006 |
| 17 | 1994 | 7631 | 815 | | -0.171 | 0.1068 | | | | 0.010 |
| 18 | 1995 | 6577 | 688 | | -0.149 | 0.1046 | | | | 0.001 |
| 19 | 1996 | 5855 | 581 | | -0.116 | 0.0992 | | | | -0.002 |
| 20 | 1997 | 6804 | 704 | | 0.150 | 0.1035 | | | | 0.007 |
| 21 | | | | | | | | | | |
| 22 | | | | Mean | 0.060 | 0.0796 | | Residuals' correls: | | |
| 23 | | | | StDev | 0.149 | 0.0193 | | with Sales | | 0.2154 |
| 24 | | | 1YearAutocorrel | | 0.1394 | 0.8279 | | Autocorrel | | 0.2927 |
| 25 | FORECAST MODEL 3 (Sales growth and autoregressive Income/Sales) | | | | | | | | | |
| 26 | Year | Sales | Inc/S | Income | FORMULAS | | | | | |
| 27 | *1997* | *6804* | *.1035* | *704* | E4. =LN(B4/B3) | | | F4. =C4/B4 | | |
| 28 | 1998 | 8073 | 0.1132 | 914 | E4:F4 copied to E4:F20 | | | | | |
| 29 | 1999 | 8018 | 0.0867 | 695 | E22. =AVERAGE(E4:E20) | | | | | |
| 30 | 2000 | 9343 | 0.0774 | 723 | E23. =STDEV(E4:E20) | | | | | |
| 31 | 2001 | 11304 | 0.0901 | 1018 | E24. =CORREL(E4:E19,E5:E20) | | | | | |
| 32 | 2002 | 12687 | 0.1008 | 1278 | E22:E24 copied to F22:F24 | | | | | |
| 33 | 2003 | 12542 | 0.0862 | 1081 | G2:G8. {=REGRESSN(F4:F19,F5:F20)} | | | | | |
| 34 | 2004 | 12798 | 0.0945 | 1210 | G12. =G6/(1-G4) | | | | | |
| 35 | 2005 | 12765 | 0.1055 | 1347 | J5. =F5-$G$6-F4*$G$4 | | | | | |
| 36 | 2006 | 9991 | 0.1115 | 1113 | J5 copied to J5:J20 | | | | | |
| 37 | 2007 | 10969 | 0.0981 | 1076 | J23. =CORREL(J5:J20,E5:E20) | | | | | |
| 38 | | | | | J24. =CORREL(J6:J20,J5:J19) | | | | | |
| 39 | E | 7114 | 610 | | B28. =B27*EXP(NORMINV(RAND(),$E$22,$E$23)) | | | | | |
| 40 | 10%ile | 4490 | 386 | | C28. =NORMINV(RAND(),$G$6+C27*$G$4,$G$8) | | | | | |
| 41 | NPV(.05) | | MinIncome | | D28. =B28*C28 | | | | | |
| 42 | SimTab | 7920 | 695 | | B28:D28 copied to B28:D37 | | | | | |
| 43 | 0 | 6134 | 635 | | B42. =NPV(0.05,D28:D37) | | | | | |
| 44 | 0.002 | 5302 | 431 | | C42. =MIN(D28:D37) | | | | | |
| 45 | 0.004 | 6708 | 625 | | B39. =AVERAGE(B43:B543) | | | | | |
| 46 | 0.006 | 7285 | 666 | | B40. =PERCENTILE(B43:B543,0.1) | | | | | |
| 47 | 0.008 | 6122 | 548 | | B39:B40 copied to C39:C40 | | | | | |

FIGURE **9.6**
**Goeing case, Forecast Model 3 (sales growth, autoregressive income/sales)**

Thus in G2:G8 we have a regression analysis to explain the income/sales ratios from years 1982 to 1997 (in F5:F20) based on the previous years' data (in F4:F19). (Recall the introduction to regression in Sections 6.7– 6.8.) The results of this regression suggest that if this year's income/sales ratio is $R(t)$, then next year's ratio, $R(t + 1)$, can be predicted as a Normal random variable with conditional mean and standard deviation

$$E(R(t+1)|R(t)) = \$G\$6 + R(t)*\$G\$4 = 0.0116 + R(t)*0.8602$$
$$Stdev(R(t+1)|R(t)) = \$G\$8 = 0.0114$$

Myerson_C09.qxd 07-09-2004 09:46 Page 338

So with the most recent income/sales ratio in cell C27, we can forecast the income/sales ratios for 10 future years in cells C28:C37 of Figure 9.6 by entering the formula

=NORMINV(RAND( ),$G$6+C27*$G$4,$G$8)

into cell C28, and then copying C28 to C28:C37.

Cells J5:J20 in Figure 9.6 compute the residual errors that we get from applying this autoregressive model to our given income/sales ratios between 1982 and 1997. That is, for each year in this period, the *residual error* in J5:J10 is the difference between the actual income/sales ratio and the ratio that would have been forecasted for that year from the previous year by our regression model in G2:G8. The low correlations found in J23 and J24 tend to validate this model's assumption that the errors around this expected value are independent random variables each year. (Cell J23 computes the correlation of the residual errors in J5:J20 with the sales growth rates in E5:E20, and cell J24 computes the correlation of the residual errors in J6:J20 with the previous year's residual error.)

In such an autoregressive time-series model, it may be useful to identify the value that would make next year's expected value the same as this year's. For this model of the income/sales ratio, such a *stationary value* can be computed by the formula

$$R = 0.0116 + R*0.8602$$

which has the solution $R = 0.083$ shown in cell G12 of Figure 9.6. In a year when the income/sales ratio is greater than this stationary value 0.083, we expect next year's ratio to be less than this year's. But if this year's income/sales ratio were less than this stationary value, we would expect next year's ratio to be greater than this year's, moving back toward the stationary value.

Future sales values are forecasted in cells B28:B37 of Figure 9.6 by the same logarithmic growth model that was used in Figure 9.5. Then the forecasted incomes for the next 10 years are computed in cells D28:D37 of Figure 9.6 by multiplying the forecasted sales times the forecasted income/sales ratios (again, just as in Figure 9.5).

The spreadsheet in Figure 9.6 includes a simulation table that contains net present values and minimum incomes from 501 simulations of this third model. This simulation table suggests that the net present value of Goeing's incomes over the coming decade may have an expected value of $7114 million, and that Goeing's annual debt burden should not exceed $386 million if we want a probability of 0.1 of having any year where the annual income is less than the debt burden. Notice that these estimates are higher than those that we got from our second model in Figure 9.5 (although the expected net present value is still not as high as in the first model in Figure 9.4). This difference may be due to the fact that the third model in Figure 9.6 takes account of the fact that the decade after 1997 is beginning with a relatively high income/sales ratio (0.1035). In this

Myerson_C09.qxd 07-09-2004 09:46 Page 339

autoregressive model, such a high ratio today can make us more optimistic about similarly high ratios in the near future.

These three different forecasting models for this Goeing Aircraft example have been described in order to illustrate different ways of using the same data. Each of these models is more sophisticated than the one before, taking more account of more structural patterns that can be found in the data, but none of these models can be called "wrong." In advanced statistics courses, you can learn many powerful techniques for evaluating how well different models may fit a given set of data. But in any real decision problem, the ultimate test for a probabilistic forecasting model must be whether it seems a reasonable description of the decision maker's uncertainty and beliefs given the available data.

## 9.4 Brownian-Motion Growth Models

Figure 9.7 shows end-of-year prices for a stock over 19 years. The observed logarithmic growth rates for this stock price are computed for each year in C6:C23 of Figure 9.7, and it is found (in C1:C2) that the mean and standard deviation of these annual growth rates are 0.1349 and 0.2993, respectively. These statistics could be used to make a forecasting model of end-of-year prices for this stock, assuming that each year's logarithmic growth rate will be another independent Normal random variable with this mean and standard deviation. In such a model, the price for the stock at the end of each year would be the price at the end of the previous year multiplied by a Lognormal random variable with log-mean 0.1349 and log-standard-deviation 0.2993. (Recall the definition of the log-mean and log-standard-deviation parameters from Section 4.3.)

But suppose that we need to make a model of the stock's prices at the end of every month over the next two years, and we only have this annual data. To give an application where such a model might be needed, suppose that somebody has offered to sell us an exotic derivative asset that, two years from now, will pay an amount equal to the highest end-of-month closing price per share this stock achieves over the next 24 months. To evaluate such a derivative asset, we need to think about how the stock might vary from month to month, even though we only have data about how the stock has varied from year to year. Of course, the stock price is likely to change less over a month than over a year, but our model will require a precise quantitative estimate of how much less. A good method for constructing such a model is provided by the theory of Brownian motion.

Brownian motion is named after the botanist Robert Brown (1773–1858), who saw through a microscope particles of dust moving randomly around in water. Brownian motion is a mathematical model developed to describe the movement of such particles. Mathematicians

Myerson_C09.qxd 07-09-2004 09:46 Page 340

| | A | B | C | D | E | F |
|---|---|---|---|---|---|---|
| 1 | | | 0.1349 | Log-drift | | |
| 2 | | | 0.2993 | Log-volatility | | |
| 3 | Stock price data | | | | | |
| 4 | Year | Price per share | | | | |
| 5 | 1978 | 4.70 | Log-growth rate | FORMULAS FROM RANGE A1:C24 | | |
| 6 | 1979 | 5.00 | 0.062 | C6. =LN(B6/B5) | | |
| 7 | 1980 | 6.54 | 0.268 | C6 copied to C6:C23 | | |
| 8 | 1981 | 3.33 | -0.674 | C1. =AVERAGE(C5:C23) | | |
| 9 | 1982 | 5.02 | 0.409 | C2. =STDEV(C5:C23) | | |
| 10 | 1983 | 6.48 | 0.256 | | | |
| 11 | 1984 | 8.39 | 0.258 | | | |
| 12 | 1985 | 11.61 | 0.325 | | | |
| 13 | 1986 | 11.36 | -0.022 | | | |
| 14 | 1987 | 8.22 | -0.323 | | | |
| 15 | 1988 | 13.47 | 0.494 | | | |
| 16 | 1989 | 19.79 | 0.385 | | | |
| 17 | 1990 | 22.69 | 0.137 | | | |
| 18 | 1991 | 23.88 | 0.051 | | | |
| 19 | 1992 | 20.06 | -0.174 | | | |
| 20 | 1993 | 21.63 | 0.075 | | | |
| 21 | 1994 | 23.50 | 0.083 | | | |
| 22 | 1995 | 39.19 | 0.511 | | | |
| 23 | 1996 | 53.25 | 0.307 | | | |
| 24 | | | | | | |
| 25 | | | | | | |
| 26 | | | | | | |
| 27 | | | | | | |
| 28 | | | | | | |
| 29 | | | | | | |
| 30 | | | | | | |
| 31 | | | | | | |
| 32 | | | | | | |
| 33 | | | | | | |
| 34 | | | | | | |
| 35 | | | | | | |
| 36 | | | | | | |
| 37 | | | | | | |
| 38 | | | | | | |
| 39 | | | | | | |

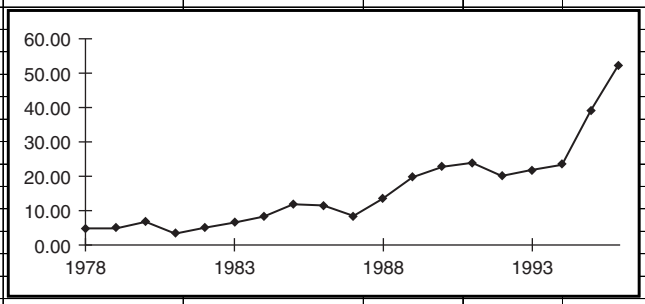


FIGURE **9.7**
**Historical data on annual end-of-year prices for a stock**

understand Brownian motion as the limit of simple *random walk* models where, in each short interval of time, the particle makes a small step either to the left or to the right, so that the numerical value that represents the particle's position will either increase or decrease by some small amount.

In a Brownian-motion model, we have some process randomly generating, at each point in time, a value that we may interpret as measuring the position of some particle which is wandering along a line. During any interval of time, the difference between the final value and the initial value is the *Brownian motion* over that time interval. The Brownian motions over separate time intervals are independent random variables. As the sum of many small steps, the Brownian motion over any time interval will be a Normal random variable. The mean of the Brownian motion is proportional to the length of the time interval, but the standard deviation is proportional to the square root of the length of the time interval, and the

Myerson_C09.qxd 07-09-2004 09:46 Page 341

constants of proportionality are called the *drift* and the *volatility*, respectively. That is, if the drift is $D$ and the volatility is $V$ then, during any time interval of length $T$ (measured in years or some other time unit), the net change in the particle's position will be a Normal random variable with mean $D*T$ and standard deviation $V*(T\text{^}0.5)$.

If we assumed that the price per share of stock changed over time according to a Brownian motion, then we would be assuming that the difference between the final price and the initial price during any year will be independent of the initial price at the beginning of the year. But if the price wandered down below \$1, then a \$1 price decrease thereafter would necessarily have probability 0, because negative prices are impossible. So we cannot realistically assume that prices change according to a Brownian motion.

But it can be reasonable to assume that, during any year, the ratio of the final price divided by the initial price will be independent of the initial price at the beginning of the year. This assumption is satisfied in a probability model where the natural logarithm of the price (rather than the price itself) changes according to a Brownian motion. In such a model, we may say that the price is changing over time according to *log-Brownian motion*. In such a log-Brownian motion, the Brownian drift and volatility of the logarithm of the price may be called the *log-drift* and the *log-volatility* of the price.

Recall from Section 4.2 that the EXP function is the inverse of the natural logarithm function LN, and the EXP function converts addition to multiplication. So over some time interval, if the natural logarithm of a price changes by the addition of some quantity $x$, then the price itself over that interval is being multiplied by the quantity EXP($x$). Thus, if a stock's price changes over time according to a log-Brownian motion with log-drift $D$ and log-volatility $V$, then, over any time interval of length $T$, the price at the end of the interval will be equal to the price at the beginning of the interval multiplied by a Lognormal random variable of the form EXP(**X**), where **X** is a Normal random variable with mean $D*T$ and standard deviation $V*(T\text{^}0.5)$.

To apply such a log-Brownian motion model to our stock price in this example, we use the annual data to estimate the log-drift and log-volatility of the stock price. Using years as our unit of time, the time interval between data points in Figure 9.7 is $T = 1$. So we can estimate the log-drift $D$ and log-volatility $V$ from the data in Figure 9.7 by the equations

$$D = D*1 = 0.1349, \qquad V = V*(1\text{^}0.5) = 0.2993$$

because the annual logarithmic growth rates have mean 0.1349 and standard deviation 0.2993 in this data.

For the monthly model, we simply change the time interval from one year to

$$T = 1/12 = 0.0833 \text{ years}$$

Myerson_C09.qxd 07-09-2004 09:46 Page 342

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | LOG-BROWNIAN MODEL OF STOCK SHARE PRICE AND TWO DERIVATIVE ASSETS | | | | | | | |
| 2 | Given parameters of share-price process (estimated from annual data) | | | | | | | |
| 3 | 0.1349 | Log-drift (annual) | | | | | | |
| 4 | 0.2993 | Log-volatility (annual) | | | | | | |
| 5 | Calculated parameters for model with Period = 1 month | | | | | | | |
| 6 | 0.0833 | Period length (in years) | | | | | | |
| 7 | 0.0112 | Mean logarthmic growth rate per period | | | | | | |
| 8 | 0.0864 | Stdev of logarthmic growth rate per period | | | | | | |
| 9 | | | | | | | | |
| 10 | Years | Share value | | Simulated final value (at 2 years) of an option | | | | |
| 11 | 0 | 53.25 | | to buy one share for $60 in 2 years. | | | | |
| 12 | 0.08333 | 50.98 | | 3.59 | | | | |
| 13 | 0.16667 | 52.40 | | | | | | |
| 14 | 0.25 | 53.81 | | Simulated final value of a derivative paying, | | | | |
| 15 | 0.33333 | 55.52 | | in 2 years, the highest monthly closing price | | | | |
| 16 | 0.41667 | 65.27 | | per share of the stock over the next 24 months. | | | | |
| 17 | 0.5 | 68.44 | | 85.25 | | | | |
| 18 | 0.58333 | 66.49 | | | | | | |
| 19 | 0.66667 | 64.03 | | | | | | |
| 20 | 0.75 | 72.40 | | | | | | |
| 21 | 0.83333 | 71.23 | | | | | | |
| 22 | 0.91667 | 74.13 | | | | | | |
| 23 | 1 | 76.65 | | | | | | |
| 24 | 1.08333 | 68.85 | | | | | | |
| 25 | 1.16667 | 81.48 | | | | | | |
| 26 | 1.25 | 85.25 | | | | | | |
| 27 | 1.33333 | 82.78 | | | | | | |
| 28 | 1.41667 | 72.10 | | | | | | |
| 29 | 1.5 | 68.94 | | | | | | |
| 30 | 1.58333 | 61.57 | | | | | | |
| 31 | 1.66667 | 58.82 | | | | | | |
| 32 | 1.75 | 57.39 | | | | | | |
| 33 | 1.83333 | 62.33 | | | | | | |
| 34 | 1.91667 | 62.01 | | | | | | |
| 35 | 2 | 63.59 | | | | | | |
| 36 | | | | | | | | |
| 37 | | | | | | | | |
| 38 | FORMULAS | | | | | | | |
| 39 | A6. =1/12 | | | | | | | |
| 40 | A7. =A3*A6 | | | | | | | |
| 41 | A8. =A4*(A6^0.5) | | | | | | | |
| 42 | A12. =A11+$A$6 | | | | | | | |
| 43 | B12. =B11*EXP(NORMINV(RAND(),$A$7,$A$8)) | | | | | | | |
| 44 | A12:B12 copied to A12:B35 | | | | | | | |
| 45 | D12. =MAX(B35-60,0) | | | | | | | |
| 46 | D17. =MAX(B12:B35) | | | | | | | |

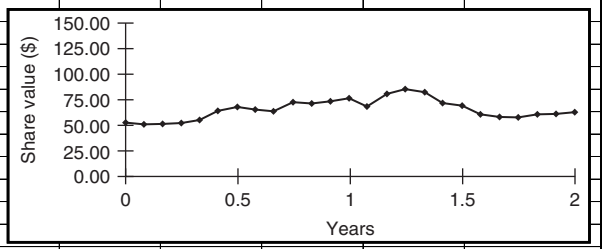


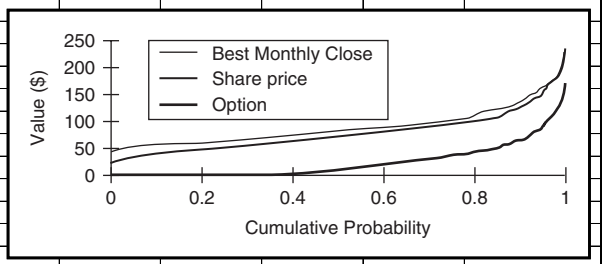


FIGURE **9.8**
**Forecast model of monthly stock prices and two derivative assets**

keeping the annual log-drift and log-volatility that we have estimated from the data. With this time interval, the log-mean $m$ and log-standard-deviation $s$ of the stock's growth ratio from one month to the next are

$$m = D*T = 0.1349*0.0833 = 0.0112$$

$$s = V*(T\text{^}0.5) = 0.2993*(0.0833\text{^}0.5) = 0.0864$$

as computed in cells A7 and A8 of Figure 9.8. Then, entering the most recent price of 53.25 into cell B11 of Figure 9.8, the simulated forecast for the next month's price is constructed in cell B12 by the formula

=B11*EXP(NORMINV(RAND( ),$A$7,$A$8))

Myerson_C09.qxd 07-09-2004 09:46 Page 343

Copying this formula from B12 to cells B12:B35 yields our forecasting model for the end-of-month values of this stock's price over the next 24 months. The highest end-of-month price over these 24 months can then be simulated in cell D17 by the formula =MAX(B12:B35).

Cell D12 in Figure 9.8 simulates the value of a more conventional derivative: an option to buy a share of this stock in 24 months for a price of $60. At the exercise date (in 24 months), this option will have value $0 if the stock price is less than $60, and otherwise the value of this option will be the difference between the actual stock price and the exercise price ($60). Thus, with cell B35 simulating the stock price at the end of the 24 months, the value of the option can be simulated in cell D12 by the formula

=MAX(B35–60,0)

Figure 9.8 shows an estimate of the cumulative distribution for this option's value, as well as the stock's value after 24 months and the highest end-of-month stock price over this period, based on data from 400 simulations of cells D12, B35, and D17 in this spreadsheet.

## 9.5 Log-Optimal Investment Strategies

Consider again the repeated investment process that was introduced in Section 4.3. In each round of this process, the decision maker decides how much of his money he wants to invest. Then a fair coin is tossed, and if it is Heads, he gets back 2 times the amount that he invested, but if it is Tails then he gets back 1/4 of the amount he invested. So if he starts the round with $W$ dollars and he invests $x$ dollars, then his wealth at the end of the round is either $W - x + 2*x$ or $W - x + x/4$, each with probability 1/2. That is, his wealth at the start of the next round will be either $W + x$ if this round's coin toss is Heads or $W - 0.75*x$ if it is Tails.

Consider what happens if, at each round of this repeated investment process, we advise the decision maker to use the strategy that maximizes his expected wealth. Given the initial wealth $W$, investing $x$ now will yield expected wealth

$$0.5*(W + x) + 0.5*(W - 0.75*x) = W + 0.125*x$$

The decision maker cannot invest more than he has, so we have the constraint $x \le W$. Subject to this constraint, this expected value is maximized by choosing $x = W$, that is, by investing 100% of his wealth in each round, which yields an expected growth ratio of 1.125 for his wealth in each round. If the decision maker starts with $1000 and applies this strategy of betting all his wealth in 100 independent rounds, his expected return at the end of 100 rounds is

$$\$1000*(1.125\text{^}100) = \$130{,}392{,}390$$

Myerson_C09.qxd 07-09-2004 09:46 Page 344

But a strategy of betting 100% of one's wealth in this investment process is very unlikely to yield such a happy outcome after 100 rounds. The final wealth will be

$$\$1000 * \mathbf{R}_1 * \mathbf{R}_2 * \mathbf{R}_3 * \cdots * \mathbf{R}_{99} * \mathbf{R}_{100}$$

where each $\mathbf{R}_i$ is an independent random variable that is equally likely to be 2 or 1/4, denoting the return ratio in the $i$th round of the process. With very high probability, about half of these $\mathbf{R}_i$ random variables will equal 2 and about half will equal 1/4, in which case his final wealth will decline to a tiny fraction of a cent, because

$$\$1000*(2\text{^}50)*(1/4\text{^}50) = \$1000*(1/2)\text{^}50 = 0.00000000000089$$

To end up with more wealth than the initial $1000, there must be at least 67 Heads among the 100 coin tosses, and the probability of such a favorable outcome is very small, less than 1/2000. With probability 0.9995, this strategy will leave the decision maker with less money than if he had never gambled anything in this process. The high expected value that we calculated comes almost entirely from the tiny probability of winning vast sums of money in the unlikely event that we get more than 80 Heads in 100 tosses.

But there are other strategies for investing in this process that can give us a much higher probability of coming out ahead. To keep things simple, let us consider strategies in which the amounts invested in each round are always the same proportion of wealth that the decision maker has at the start of the round. If the decision maker invests a fraction $k$ of his wealth at each round, keeping the fraction $1 - k$ safely in cash, then his wealth at the end of 100 rounds will be

$$\$1000*(k*\mathbf{R}_1 + 1 - k)*(k*\mathbf{R}_2 + 1 - k)* \cdots *(k*\mathbf{R}_{99} + 1 - k)*(k*\mathbf{R}_{100} + 1 - k)$$

when he starts with $1000 in the first round. Here $\mathbf{R}_1, \ldots \mathbf{R}_{100}$ are independent random variables, each equally likely to be 2 or 0.25. The spreadsheet shown in Figure 9.9 simulates the results of four such strategies, where the invested fraction $k$ is 0.02, 0.15, 0.50, and 1.

In Figure 9.9, cells A6:A105 contain the formula

```
=IF(RAND( )<0.5,2,0.25)
```

to simulate $\mathbf{R}_1, \ldots, \mathbf{R}_{100}$, the random return ratios at each round. With the fraction 0.02 in cell B2 and the initial wealth 1000 in cell B5, the results of investing 2% of wealth at the first round are simulated in cell B6 by the formula

```
=B5*(B$2*$A6+(1–B$2)*1)
```

(Each dollar from B5 that is invested in the first round of the process will yield A6 dollars for the second round, and each dollar held as cash will yield 1 dollar for second round.) Copying cell B6 to B6:B105 gives us a

| | A | B | C | D | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | | Fraction invested: | | | | | FORMULAS | | | | |
| 2 | | 0.02 | 0.15 | 0.5 | 1 | | A6. =IF(RAND()<0.5,2,0.25) | | | | |
| 3 | | | | | | | A6 copied to A6:A105 | | | | |
| 4 | Bet | Simulated changes in wealth | | | | | B6. =B5*(B$2*$A6+(1-B$2)*1) | | | | |
| 5 | pays | 1000 | 1000 | 1000 | 1000 | | B6 copied to B6:E105 | | | | |
| 6 | 0.25 | 985 | 888 | 625 | 250 | | H21. =B105 | H21 copied to H21:K21 | | | |
| 7 | 0.25 | 970 | 788 | 391 | 63 | | H16. =AVERAGE(H22:H522) | | | | |
| 8 | 0.25 | 956 | 699 | 244 | 16 | | H17. =PERCENTILE(H22:H522,0.25) | | | | |
| 9 | 0.25 | 941 | 620 | 153 | 4 | | H18. =PERCENTILE(H22:H522,0.5) | | | | |
| 10 | 0.25 | 927 | 551 | 95 | 1 | | H19. =PERCENTILE(H22:H522,0.75) | | | | |
| 11 | 2.00 | 946 | 633 | 143 | 2 | | H16:H19 copied to H16:K19 | | | | |
| 12 | 0.25 | 932 | 562 | 89 | 0 | | | | | | |
| 13 | 2.00 | 950 | 646 | 134 | 1 | | | Fraction invested: | | | |
| 14 | 2.00 | 969 | 743 | 201 | 2 | | | 0.02 | 0.15 | 0.5 | 1 |
| 15 | 0.25 | 955 | 660 | 126 | 0 | | | FinalWealth statistics | | | |
| 16 | 2.00 | 974 | 759 | 189 | 1 | | Mean | 1278 | 6118 | 62184 | 0 |
| 17 | 0.25 | 959 | 673 | 118 | 0 | | 25%ile | 1138 | 1276 | 3 | 0 |
| 18 | 0.25 | 945 | 597 | 74 | 0 | | 50%ile | 1264 | 2775 | 40 | 0 |
| 19 | 0.25 | 931 | 530 | 46 | 0 | | 75%ile | 1404 | 6038 | 549 | 0 |
| 20 | 0.25 | 917 | 471 | 29 | 0 | | | FinalWealths from B105:E105 | | | |
| 21 | 0.25 | 903 | 418 | 18 | 0 | | SimTable | 1404 | 6038 | 549 | 0 |
| 22 | 2.00 | 921 | 480 | 27 | 0 | | 0 | 1454 | 7824 | 1316 | 0 |
| 23 | 2.00 | 939 | 552 | 40 | 0 | | 0.002 | 1309 | 3596 | 95 | 0 |
| 24 | 2.00 | 958 | 635 | 61 | 0 | | 0.004 | 1454 | 7824 | 1316 | 0 |
| 25 | 0.25 | 944 | 564 | 38 | 0 | | 0.006 | 1505 | 10138 | 3160 | 0 |
| 26 | 0.25 | 930 | 500 | 24 | 0 | | 0.008 | 1614 | 17022 | 18199 | 0 |
| 27 | 2.00 | 948 | 575 | 36 | 0 | | 0.01 | 1309 | 3596 | 95 | 0 |
| 28 | 2.00 | 967 | 662 | 53 | 0 | | 0.012 | 956 | 349 | 0 | 0 |
| 29 | 2.00 | 987 | 761 | 80 | 0 | | 0.014 | 1404 | 6038 | 549 | 0 |
| 30 | 0.25 | 972 | 675 | 50 | 0 | | 0.016 | 990 | 452 | 0 | 0 |
| 31 | 2.00 | 991 | 777 | 75 | 0 | | 0.018 | 1099 | 984 | 1 | 0 |
| 32 | 0.25 | 976 | 689 | 47 | 0 | | 0.02 | 1264 | 2775 | 40 | 0 |

FIGURE **9.9**
**Comparison of four investment strategies for a repeated investment process**

simulation of the strategy of investing 2% of the available wealth at every round. Then copying B6:B105 into columns C, D, and E, we can simulate the results of investing the other fixed fractions (15%, 50%, and 100%) of wealth that have been entered into cells C2, D2, and E2. So the final wealth after 100 rounds from each strategy are contained in cells B105:E105 of this spreadsheet. These cells are not shown in Figure 9.9, but they are echoed in cells H21:K21, and a table with data from 501 simulations is stored below these cells. From this simulation data, cells H16:K19 compute the mean and quartiles of the final wealths generated by these four investment strategies.

Notice that the strategy of investing 15% at each round yields higher values of all three quartile boundary points than the other strategies that are considered in this spreadsheet. Furthermore, the final returns from the 15% strategy that we see in cells I21:I32 of Figure 9.9 were better than the other three strategies in all but one of the 12 simulations from which we have data in H21:K32. So we see that in the long run, a strategy of betting about 15% of wealth seems likely to yield better results than strategies of investing much smaller or much larger fractions of wealth in this gamble. Let us now do some more general analysis to see why this strategy performs

Myerson_C09.qxd 07-09-2004 09:46 Page 346

so well and to find another investment strategy (investing 16.67%) that would be likely to do even better in the long run.

Suppose that the decision maker will invest in this process for $T$ rounds and will follow a strategy of investing some fixed fraction $k$ of his wealth at every round, while holding the fraction $1 - k$ safely as cash. Then let $\mathbf{q}_i$ denote the random growth ratio for the decision maker's wealth in the $i$th round of this process, so we have here

$$\mathbf{q}_i = k*\mathbf{R}_i + 1 - k$$

where $\mathbf{R}_i$ is equally likely to be 2 or 1/4 in each round $i$. Let $W_0$ denote the decision maker's initial wealth at the start of the first round. Then his final wealth after $T$ rounds will be

$$\mathbf{W}_T = \mathrm{W}_0 * \mathbf{q}_1 * \mathbf{q}_2 * \cdots * \mathbf{q}_{T-1} * \mathbf{q}_T$$

Taking the logarithm converts multiplication to addition, so we get

$$\mathrm{LN}(\mathbf{W}_T) = \mathrm{LN}(\mathrm{W}_0) + \mathrm{LN}(\mathbf{q}_1) + \mathrm{LN}(\mathbf{q}_2) + \cdots + \mathrm{LN}(\mathbf{q}_T)$$

Now let $\mathbf{G}$ denote the average logarithmic growth rate per round over $n$ rounds when the fraction $k$ is invested at every round:

$$\mathbf{G} = (\mathrm{LN}(\mathbf{q}_1) + \mathrm{LN}(\mathbf{q}_2) + \cdots + \mathrm{LN}(\mathbf{q}_T))/T$$

Then the final wealth $\mathbf{W}_T$ satisfies

$$\mathrm{LN}(\mathbf{W}_T) = \mathrm{LN}(\mathrm{W}_0) + \mathbf{G}*T$$

and so, applying the EXP function to both sides

$$\mathbf{W}_T = W_0*\mathrm{EXP}(\mathbf{G}*T)$$

But $\mathbf{G}$ is the average of many random variables that are drawn independently from the same probability distribution, and so we can apply the law of large numbers and the Central Limit Theorem to $\mathbf{G}$. When $T$ is large, the law of large numbers implies that this average logarithmic growth rate $\mathbf{G}$ is very likely to be very close to $E(\mathrm{LN}(\mathbf{q}_i))$, the expected logarithmic growth rate in any round.

This expected logarithmic growth rate $E(\mathrm{LN}(\mathbf{q}_i))$ depends on the fraction $k$ of wealth that is invested in the risky gamble at each round. So let $m(k)$ here denote the expected logarithmic growth rate per round when the fraction $k$ of wealth is invested in the gamble:

$$m(k) = E(\mathrm{LN}(\mathbf{q}_i)) = E(\mathrm{LN}(k*\mathbf{R}_i + 1 - k))$$

For this example, we have

$$m(k) = E(\mathrm{LN}(\mathbf{q}_i)) = 0.5*\mathrm{LN}(k*2 + 1 - k) + 0.5*\mathrm{LN}(k*0.25 + 1 - k)$$

When $T$ is large, the average logarithmic growth rate $\mathbf{G}$ is very likely to be very close to $m(k)$, and so the final wealth $\mathbf{W}_T$ is likely to be close to $W_0*\mathrm{EXP}(m(k)*T)$. Thus, the best likely outcomes occur when we choose $k$ so as to maximize the expected logarithmic growth rate $m(k)$.

Myerson_C09.qxd 07-09-2004 09:46 Page 347

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | | | FractionInvested | | | | | |
| 2 | | | 0.166667 | | | | | |
| 3 | Prob'y | BetPays | RtnRatio | LogRateReturn | | | | |
| 4 | 0.5 | 2.00 | 1.17 | 0.1542 | | | | |
| 5 | 0.5 | 0.25 | 0.87 | -0.1335 | | | | |
| 6 | | | | | | | | |
| 7 | | E(LogRateReturn) | | 0.0103 | | | | |
| 8 | | Stdev(LogRateReturn) | | 0.1438 | | | | |
| 9 | | | | | | | | |
| 10 | Number periods repeated | | | 100 | | | | |
| 11 | | E(LogFinalReturn) | | 1.030964 | | | | |
| 12 | | Stdev(LogFinalReturn) | | 1.4384 | | | | |
| 13 | InitialWealth | | 1000 | | | | 7861 | E(FinalReturn) |
| 14 | Lognormal approximation of FinalWealth | | | | | | | |
| 15 | | %rank | Percentile | | | | | |
| 16 | | 0.05 | 263 | | | | | |
| 17 | | 0.25 | 1063 | | | | | |
| 18 | | 0.50 | 2804 | | | | | |
| 19 | | 0.75 | 7398 | | | | | |
| 20 | | 0.95 | 29872 | | | | | |
| 21 | | | | | | | | |
| 22 | FORMULAS | | | | | | | |
| 23 | C4. =$C$2*B4+(1-$C$2)*1 | | | | C35. =C2 | | | |
| 24 | D4. =LN(C4) | | | | B37. =DISCRINV(RAND(),B4:B5,A4:A5) | | | |
| 25 | C4:D4 copied to C5:D5 | | | | D38. =LN($C$35*B38+(1-$C$35)*1) | | | |
| 26 | D7. =SUMPRODUCT(D4:D5,A4:A5) | | | | D38 copied to D38:D538 | | | |
| 27 | D8. =STDEVPR(D4:D5,A4:A5) | | | | D35. =AVERAGE(D38:D538) | | | |
| 28 | D11. =D7*D10 | | | | D36. =STDEV(D38:D538) | | | |
| 29 | D12. =D8*(D10^0.5) | | | | | | | |
| 30 | F13. =C13*(SUMPRODUCT(C4:C5,A4:A5)^D10) | | | | | | | |
| 31 | C16. =$C$13*EXP(NORMINV(B16,$D$11,$D$12)) | | | | C16 copied to C16:C20 | | | |
| 32 | SOLVER: max D7 (or D35) by changing C2. | | | | | | | |
| 33 | | | | | | | | |
| 34 | From simulation | | FracInv | LogRateReturn statistics | | | | |
| 35 | | | 0.166667 | 0.0072 | Mean | | | |
| 36 | | Bet pays | | 0.1440 | Stdev | | | |
| 37 | SimTabl | 0.25 | | LogRateReturn | | | | |
| 38 | 0 | 0.25 | | -0.1335 | | | | |
| 39 | 0.002 | 2.00 | | 0.1542 | | | | |
| 40 | 0.004 | 2.00 | | 0.1542 | | | | |
| 41 | 0.006 | 2.00 | | 0.1542 | | | | |
| 42 | 0.008 | 2.00 | | 0.1542 | | | | |
| 43 | 0.01 | 0.25 | | -0.1335 | | | | |

FIGURE **9.10**
**Computing the log-optimal strategy for the repeated investment process**

Figure 9.10 shows a spreadsheet for this maximization. The possible returns from the gamble, per dollar invested, at each round are listed in cells B4:B5, with corresponding probabilities in A4:A5. Cell C2 contains the fraction $k$ that is to be invested at each round. With this fraction invested in the risky gamble, the corresponding growth ratios for the decision maker's wealth in each possible outcome are to be computed in C4:C5 by entering the formula

$$=\$C\$2*B4+(1-\$C\$2)*1$$

into cell C4 and then copying it to C5. The logarithmic growth rates are computed in D4:D5 by entering the formula

$$=LN(C4)$$

Myerson_C09.qxd 07-09-2004 09:46 Page 348

into cell D4 and then copying it to D5. Then the expected logarithmic growth rate per round $m(k)$ is computed in cell D7 by the formula

=SUMPRODUCT(D4:D5,A4:A5)

Using this spreadsheet, you can verify that when the fraction invested is $k = 0.15$ (entered into cell C2), the expected logarithmic growth rate is $m(0.15) = 0.0102$ (calculated in cell D7). But changing the fraction invested to 0.02 or 0.5 or 1 would significantly reduce the expected logarithmic growth rate, to 0.0023 or –0.0323 or –0.3466, respectively. So we can explain the superior long-run performance of the 15% investment strategy in Figure 9.9 by the fact that the 15% investment strategy yields a higher expected logarithmic growth rate than the other three investment strategies that were considered in Figure 9.9.

But what fraction should be invested to achieve the best likely performance in the long run? This problem can be answered by letting Solver maximize this expected logarithmic growth rate in cell D7 by changing the fraction invested in cell C2 of Figure 9.10. The result of this maximization is the value $k = 0.1667$ in cell C2. That is, the maximum likely growth can be achieved by investing 16.67% of total wealth at each round, which yields an expected logarithmic growth rate of 0.0103 per round.

This strategy, which maximizes the expected logarithmic growth rate among all strategies of investing fixed fractions of wealth, is called the *log-optimal* investment strategy. If two people invest in this process for many rounds, one using this log-optimal investment strategy and the other investing some other fixed fraction of wealth at each round, then the probability of the log-optimal strategy yielding more wealth than the other strategy will increase toward 1 as the number of rounds becomes large.

We can also apply the Central Limit Theorem (from Section 2.6) to get a better approximation to the probability distribution of the final wealth, $\mathbf{W}_T$, for any strategy of investing any fixed fraction of wealth $k$. Let $s(k)$ denote the standard deviation of the logarithmic growth rate in any round when the fraction $k$ is invested in the gamble:

$$s(k) = \mathit{Stdev}(\mathrm{LN}(\mathbf{q}_i)) = \mathit{Stdev}(\mathrm{LN}(k*\mathbf{R}_i + 1 - k))$$

Cell D8 in Figure 9.10 computes this standard deviation $s(k)$ for the value of $k$ that has been entered into cell D2 by the formula

=STDEVPR(D4:D5,A4:A5)

By the Central Limit Theorem, when the number of rounds $T$ is large and the fraction of wealth $k$ is invested at each round, the average logarithmic growth rate $\mathbf{G}$ per round will be approximately Normal with mean $m(k)$ and standard deviation $s(k)/(T\text{^}0.5)$. So the growth ratio $\mathbf{W}_T/\mathrm{W}_0 = \mathrm{EXP}(\mathbf{G}*T)$ over $T$ rounds of this investment strategy is approximately Lognormal with log-mean $m(k)*T$ and log-standard-deviation $(s(k)/(T\text{^}0.5))*T = s(k)*(T\text{^}0.5)$.

Myerson_C09.qxd 07-09-2004 09:46 Page 349

Thus, for any number $p$ between 0 and 1, the final wealth $\mathbf{W}_T$ has approximately probability $p$ of being less than the value

$$W_0 * \text{EXP}(\text{NORMINV}(p, m(k)*T, s(k)*(\text{T}\text{\^{}}0.5)))$$
$$= W_0 * \text{EXP}(m(k)*T + \text{NORMINV}(p, 0, 1)*s(k)*(T\text{\^{}}0.5))$$

This formula can be used to estimate percentile values for the final wealth distribution, as shown in cells C16:C20 of Figure 9.10. When $T$ is large, $T$ becomes much larger than its square root, $T$^0.5. Thus, for any $p$ between 0 and 1, the log-optimal strategy $k$ that maximizes $m(k)$ will have the highest $p$ percentile values in the long run.

## 9.6 Exponential Arrival Models

Let us now consider the problem of how to make a good probability model of the random arrival of customers in a shop over some period of time. Customers arrive at random points in time, and the time that we may have to wait from any one customer's arrival until the next customer's arrival will be a random variable. In a model of such arrival processes, it is often very useful to assume that these waiting times between customer arrivals are random variables drawn from an Exponential distribution. In this section, we show why this assumption of Exponential waiting times is so useful and what it implies about other important quantities in the arrival processes.

Recall from Section 4.8 that Exponential probability distributions have one parameter: the mean. The standard deviation of an Exponential random variable is equal to its mean. An Exponential random variable is always a positive number, like a Lognormal random variable. But an Exponential probability distribution always has its highest probability density near 0, unlike the Normal and Lognormal distributions that have higher density near the mean. For any positive number μ, an Exponential random variable with mean μ can be simulated by the Simtools formula EXPOINV(RAND( ), μ), which is equivalent to the Excel formula –μ*LN(1 – RAND( )).

We may say that customer arrivals have the *no-memory property* if, at any point in time, our belief about how long we may have to wait until the next customer's arrival does not depend on how long it has been since the last customer arrived. To state the no-memory property more precisely, let $\mathbf{X}$ denote the random waiting time (say, measured in minutes) from the last customer's arrival until the next customer's arrival. Let $m$ and $n$ be any two positive numbers. At the moment when the last customer arrived, the probability that we would have to wait more than $m$ minutes until the next customer arrives is $P(\mathbf{X} > m)$. But if it is now $n$ minutes since the last customer

| | A | B | C | D | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | NO-MEMORY PROPERTY OF EXPONENTIAL WAITING TIMES | | | | | | | | | | |
| 2 | Expected waiting time | | | | | | | | | | |
| 3 | 10 | | | | | | | | | | |
| 4 | | | | Time | | | | | | | |
| 5 | | | | 2 | Interval | | | | | | |
| 6 | | | | 5 | Restart | | | | | | |
| 7 | | | | 7 | Restart+Interval | | | | | | |
| 8 | | | | | | | | | | | |
| 9 | | | | P(W>2) | P(W>7 \| W>5) | | Count(W>5) | | | | |
| 10 | | | | 0.820 | 0.810 | | 300 | | | | |
| 11 | | | | | | Making Exponential random variables: | | | | | |
| 12 | Random waiting times | | | | | | | 0.6009 | (rand) | | |
| 13 | SimTab | 1.7315 | | >2? | >7? \| >5 | | | 9.1848 | using EXPOINV | | |
| 14 | 0 | 4.595 | | 1 | .. | | | 9.1848 | using GAMINV | | |
| 15 | 0.002 | 0.4029 | | 0 | .. | | | 9.1848 | using LN | | |
| 16 | 0.004 | 2.703 | | 1 | .. | | | | | | |
| 17 | 0.006 | 4.3078 | | 1 | .. | FORMULAS | | | | | |
| 18 | 0.008 | 23.856 | | 1 | 1 | D7. =D6+D5 | | | | | |
| 19 | 0.01 | 4.1477 | | 1 | .. | B13. =EXPOINV(RAND(),A3) | | | | | |
| 20 | 0.012 | 8.8153 | | 1 | 1 | D14. =IF(B14>$D$5,1,0) | | | | | |
| 21 | 0.014 | 5.2047 | | 1 | 0 | E14. =IF(B14>$D$6,IF(B14>$D$7,1,0),"..") | | | | | |
| 22 | 0.016 | 12.787 | | 1 | 1 | D14:E14 copied to D14:E514 | | | | | |
| 23 | 0.018 | 13.641 | | 1 | 1 | D10. =AVERAGE(D14:D514) | | | | | |
| 24 | 0.02 | 0.1784 | | 0 | .. | E10. =AVERAGE(E14:E514) | | | | | |
| 25 | 0.022 | 18.774 | | 1 | 1 | G10. =COUNT(E14:E514) | | | | | |
| 26 | 0.024 | 3.7272 | | 1 | .. | H12. =RAND() | | | | | |
| 27 | 0.026 | 1.7513 | | 0 | .. | H13. =EXPOINV(H12,A3) | | | | | |
| 28 | 0.028 | 41.333 | | 1 | 1 | H14. =GAMINV(H12,A3,A3) | | | | | |
| 29 | 0.03 | 13.929 | | 1 | 1 | H15. =-A3*LN(1-H12) | | | | | |

FIGURE **9.11**
**No-memory property of Exponential waiting times**

arrived, then we know now that $\mathbf{X} > n$ and so the probability that we will have to wait $m$ minutes more until the next customer arrives is $P(\mathbf{X} > n + m | \mathbf{X} > n)$. The no-memory property asserts that

$$P(\mathbf{X} > m) = P(\mathbf{X} > n + m | \mathbf{X} > n)$$

This no-memory property holds when we assume that the times between successive customer arrivals are Exponential random variables, and it does not hold for any other probability distribution. This convenient property is the reason why the Exponential distribution is so commonly used to describe waiting times in such arrival processes.

Figure 9.11 contains a model to illustrate the no-memory property of the Exponential distribution. The mean waiting time for this example is $\mu = 10$ minutes, and this mean has been entered into cell A3. The random waiting time $\mathbf{X}$ is simulated in cell B13 by the formula

=EXPOINV(RAND( ),A3)

and the outcomes of 501 simulations of this Exponential random variable are contained in cells B14:B514. From this data, the spreadsheet estimates $P(\mathbf{X} > m)$ and $P(\mathbf{X} > n + m | X > n)$ for the values of $m = 2$ and $n = 5$. Cell D10 indicates that waiting time was more than 2 minutes in 82% of these 501 simulations, cell G10 indicates that the waiting time was more than 5 minutes in 300 simulations, and cell E10 indicates that the waiting time was

| | A | B | C | D | E | F | G | H | I | J | K |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | MeanTimeBetweenArrivals | | | Time until arrival# | | | | 8 | Gamma | mean | stdev |
| 2 | 10 | | | | | | 80.03 | | 62.68 | 80 | 28.28 |
| 3 | | | | | | | | | | | |
| 4 | | | #Arrivals before time | | | | | 60 | Poisson | mean | stdev |
| 5 | Arrival# | ExplWait | TotalTime | | | | 7 | | 5 | 6 | 2.449 |
| 6 | 1 | 6.18 | 6.18 | 1 | | | | | | | |
| 7 | 2 | 9.65 | 15.83 | 1 | | | | | | | |
| 8 | 3 | 4.91 | 20.74 | 1 | G2 and I2 should have same distributn. | | | | | | |
| 9 | 4 | 11.09 | 31.83 | 1 | G5 and I5 should have same distributn. | | | | | | |
| 10 | 5 | 7.50 | 39.34 | 1 | | | | | | | |
| 11 | 6 | 13.27 | 52.61 | 1 | | FORMULAS | | | | | |
| 12 | 7 | 2.22 | 54.82 | 1 | | B6. =EXPOINV(RAND(),$A$2) | | | | | |
| 13 | 8 | 25.21 | 80.03 | 0 | | C6. =SUM($B$6:B6) | | | | | |
| 14 | 9 | 15.97 | 96.01 | 0 | | D6. =IF(C6<$H$4,1,0) | | | | | |
| 15 | 10 | 0.99 | 96.99 | 0 | | B6:D6 copied down column | | | | | |
| 16 | 11 | 14.66 | 111.66 | 0 | | G2. =VLOOKUP(H1,A6:C35,3) | | | | | |
| 17 | 12 | 1.70 | 113.36 | 0 | | I2. =GAMINV(RAND(),J2,K2) | | | | | |
| 18 | 13 | 9.74 | 123.10 | 0 | | J2. =A2*H1 | | | | | |
| 19 | 14 | 3.26 | 126.36 | 0 | | K2. =A2*(H1^0.5) | | | | | |
| 20 | 15 | 13.38 | 139.74 | 0 | | G5. =SUM(D6:D35) | | | | | |
| 21 | 16 | 0.99 | 140.73 | 0 | | G6. =IF(C35<$H$4,"MORE???","") | | | | | |
| 22 | 17 | 29.04 | 169.77 | 0 | | I5. =POISINV(RAND(),J5) | | | | | |
| 23 | 18 | 1.36 | 171.13 | 0 | | J5. =H4/A2 | | | | | |
| 24 | 19 | 1.20 | 172.33 | 0 | | K5. =J5^0.5 | | | | | |
| 25 | 20 | 0.60 | 172.92 | 0 | | | | | | | |
| 26 | 21 | 0.00 | 172.92 | 0 | | | | | | | |
| 27 | 22 | 23.95 | 196.88 | 0 | | | | | | | |
| 28 | 23 | 1.45 | 198.33 | 0 | | | | | | | |
| 29 | 24 | 36.26 | 234.59 | 0 | | | | | | | |
| 30 | 25 | 11.85 | 246.44 | 0 | | | | | | | |
| 31 | 26 | 13.20 | 259.64 | 0 | | | | | | | |
| 32 | 27 | 30.12 | 289.76 | 0 | | | | | | | |
| 33 | 28 | 2.86 | 292.63 | 0 | | | | | | | |
| 34 | 29 | 11.76 | 304.39 | 0 | | | | | | | |
| 35 | 30 | 15.16 | 319.55 | 0 | | | | | | | |

FIGURE **9.12**
**Poisson and Gamma random variables from Exponential waiting times**

more than 7 minutes in 81% of these 300 simulations. So this simulation data suggests that $P(\mathbf{X} > 2)$ and $P(\mathbf{X} > 5 + 2|\mathbf{X} > 5)$ are both close to 0.82. In fact, the correct probabilities are $P(\mathbf{X} > 2) = 0.8187 = P(\mathbf{X} > 5 + 2|\mathbf{X} > 5)$. That is, when waiting times have an Exponential distribution with mean 10 minutes, the probability of having to wait another 2 minutes or more for the next customer is always 0.8187, regardless of how long we have already waited since the last customer.

The no-memory property is not necessarily appropriate to every arrival process. Consider a bus stop where buses arrive every 30 minutes exactly. If the last bus just left, then we would expect to wait 30 minutes for the next bus. But if we have already waited 28 minutes, then we should be confident that the bus will arrive within the next 2 minutes. So the no-memory property does not apply to buses that arrive with great regularity. The no-memory property of the Exponential distribution is appropriate for arrivals that are extremely erratic, so the probability of a new arrival in the next minute does not depend on when arrivals happened before.

Figure 9.12 shows a model of the first 30 arrivals in a shop where the waiting times between successive customers' arrivals are independent

Myerson_C09.qxd 07-09-2004 09:46 Page 352

Exponential random variables with mean 10 minutes. The mean time between arrivals has been entered into cell A2. The formula

=EXPOINV(RAND( ),$A$2)

has been entered in cell B6 to denote the time of the first customer's arrival, starting (say) from "time 0" when the shop opens in the morning. Then cell B6 has been copied to cells B7:B35 to simulate the waiting times between each successive pair of customer's arrivals. The formula

=SUM($B$6:B6)

has been entered into cell C6 and copied to cells C6:C35 to compute the times (counting from "time 0") when each of the first 30 customers arrive.

When the times between successive customers' arrivals are independent Exponential random variables with mean μ, the time until the *k*th arrival is a Gamma random variable with mean μ∗*k* and standard deviation μ ∗(*k*^0.5). (Recall from Section 4.8 that Gamma probability distributions can be parameterized by their mean and standard deviation, and a Gamma random variable with mean μ and standard deviation σ can be constructed by the Simtools formula GAMINV(RAND( ) μ, σ, ).) In Figure 9.12, for example, the time of the eighth customer's arrival is simulated in cell C13, which is echoed in cell G2 (when the number 8 is in cell H1). Because the expected waiting time between successive arrivals is 10 minutes in this model, the time until the eighth arrival is a Gamma random variable with mean 10∗8 = 80 minutes and standard deviation 10∗(8^0.5) = 28.28 minutes. A random variable with the same Gamma distribution can be constructed using the Simtools function GAMINV, as shown in cell I2. (Cells I2 and C13 do not have the same value in Figure 9.12 because they are generated from different random seeds, but a large simulation table should yield cumulative distributions that are very close for these two random variables.)

The number of arrivals in the first 60 minutes is computed in cell G5 of Figure 9.12. To do this computation, the number 60 has been entered into cell H4 and the formula

=IF(C6<$H$4,1,0)

has been entered into cell D6 and copied to D6:D35. Then the formula

=SUM(D6:D35)

in cell G5 counts the number of arrival times in C6:C35 that are less than 60.

When the times between successive customers' arrivals are independent Exponential random variables with mean μ, the number of arrivals in any time period of length *T* is a Poisson random variable with mean *T*/μ. Recall from Section 4.8.9 that Poisson probability distributions are characterized by one parameter, the mean (because the standard deviation of a Poisson random variable is the square root of its mean). The value of a Poisson

Myerson_C09.qxd 07-09-2004 09:46 Page 353

random variable with mean μ is always a nonnegative integer and can be simulated with Simtools by the formula POISINV(RAND( ), μ). If the expected waiting time is 10 minutes between arrivals, then the expected number of arrivals in 60 minutes is 60/10 = 6, which is computed in cell J5 of Figure 9.12. So a Poisson random variable with this mean can also be simulated by the formula

=POISINV(RAND( ),J5)

as shown in cell I5 of Figure 9.12.

In general, the expected arrival rate per unit time is the reciprocal of the expected time between successive arrivals. That is, if μ denotes the expected waiting time between arrivals and λ denotes the expected rate of arrivals per unit time, then λ and μ satisfy the equation

$$\lambda = 1/\mu$$

In this example, the expected time between arrivals μ is 10 minutes (or 1/6 hour), and so the expected arrival rate λ is 1/10 per minute (or 6 arrivals per hour).

We could just as well have described these arrivals as a *Poisson process* with this expected rate of arrivals. In any Poisson process with expected arrival rate λ, the number of arrivals in any time period of length $T$ is a Poisson random variable with mean $\lambda * T$, and the numbers of arrivals in separate periods of time are independent of each other. In such a Poisson process, the times between successive arrivals are independent Exponential random variables with a mean $1/\lambda$. Thus, we can say that cells C6:C35 in Figure 9.12 contain arrival times for a Poisson process that has an expected arrival rate of 0.1 per minute (or 6 per hour).

Now let us consider a Poisson process over a short time interval of length $t$, where $t$ is small enough that $\lambda * t \le 0.1$. Then there is very small probability of two or more arrivals in any period of length $t$, and so the number of arrivals in this short subinterval may be well approximated by a random variable that is either 1, with probability $\lambda * t$, or 0 with probability $1 - \lambda * t$. For example, consider Figure 9.13, where cell A1 contains the value 0.1 and cell B8 contains a RAND( ). Then the Poisson random variable =POISINV(B8,$A$1) in cell C8 and the {0, 1} random variable =IF(B8<1–$A$1,0,1) in cell D8 will have the same value more than 99% of the time. Notice that this approximation keeps the same expected number of arrivals, because $\lambda * t * 1 + (1 - \lambda * t) * 0 = \lambda * t$.

The model in Figure 9.13 simulates 20 minutes of arrivals in a Poisson process with expected arrival rate λ = 0.1 per minute. Over this interval of length $T = 20$ minutes, the expected number of arrivals is $\lambda * T = 0.1 * 20 = 2$, which is computed in cell A2. So the number of arrivals over these 20 minutes can be simulated in cell A2 by the formula =POISINV(RAND( ),2). But the number of arrivals in any subinterval of length $t = 1$ minute is a Poisson random variable with mean $\lambda * t = 0.1 * 1 = 0.1$, which can be simulated by the formula =POISINV(RAND( ),0.1). Twenty such independent Poisson

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | 0.1 | Expected arrival rate per minute | | | | | | |
| 2 | 2 | Expected arrivals in 20 minutes | | | | | | |
| 3 | 2 | Poisson simulation of arrivals during a 20-minute interval | | | | | | |
| 4 | 5 | Sum of Poisson arrivals in 20 one-minute intervals | | | | | | |
| 5 | 5 | Sum of 0-or-1 arrivals in 20 one-minute intervals (Binomial) | | | | | | |
| 6 | | | | | | | | |
| 7 | Minute | (rand) | Poisson | 0-or-1 approximation | | | | |
| 8 | 1 | 0.4333 | 0 | 0 | | | | |
| 9 | 2 | 0.205 | 0 | 0 | FORMULAS | | | |
| 10 | 3 | 0.0836 | 0 | 0 | A2. =A1*COUNT(A8:A27) | | | |
| 11 | 4 | 0.4781 | 0 | 0 | A3. =POISINV(RAND(),A2) | | | |
| 12 | 5 | 0.8029 | 0 | 0 | B8. =RAND() | | | |
| 13 | 6 | 0.9178 | 1 | 1 | C8. =POISINV(B8,$A$1) | | | |
| 14 | 7 | 0.6538 | 0 | 0 | D8. =IF(B8<1-$A$1,0,1) | | | |
| 15 | 8 | 0.7628 | 0 | 0 | B8:D8 copied to B8:D27 | | | |
| 16 | 9 | 0.3264 | 0 | 0 | A4. =SUM(C8:C27) | | | |
| 17 | 10 | 0.7363 | 0 | 0 | A5. =SUM(D8:D27) | | | |
| 18 | 11 | 0.658 | 0 | 0 | C44. =A4 | | | |
| 19 | 12 | 0.0147 | 0 | 0 | D44. =A5 | | | |
| 20 | 13 | 0.4474 | 0 | 0 | E44. =A3 | | | |
| 21 | 14 | 0.9407 | 1 | 1 | G41. =COUNT(B45:B545) | | | |
| 22 | 15 | 0.6443 | 0 | 0 | G43. =MAX(C45:E545) | | | |
| 23 | 16 | 0.9205 | 1 | 1 | C40. =AVERAGE(C45:C545) | | | |
| 24 | 17 | 0.9877 | 1 | 1 | C41. =STDEV(C45:C545) | | | |
| 25 | 18 | 0.65 | 0 | 0 | C40:C41 copied to C40:E41 | | | |
| 26 | 19 | 0.1124 | 0 | 0 | C30. =COUNTIF(C$44:C$544,$B30)/$G$40 | | | |
| 27 | 20 | 0.9855 | 1 | 1 | C30 copied to C30:E38 | | | |
| 28 | | | | | | | | |
| 29 | | Value | Relative frequency in SimTable below | | | | | |
| 30 | | 0 | 0.156 | 0.134 | 0.132 | | | |
| 31 | | 1 | 0.253 | 0.261 | 0.234 | | | |
| 32 | | 2 | 0.265 | 0.291 | 0.281 | | | |
| 33 | | 3 | 0.186 | 0.196 | 0.176 | | | |
| 34 | | 4 | 0.090 | 0.072 | 0.122 | | | |
| 35 | | 5 | 0.028 | 0.026 | 0.042 | | | |
| 36 | | 6 | 0.016 | 0.014 | 0.010 | | | |
| 37 | | 7 | 0.006 | 0.006 | 0.004 | | | |
| 38 | | 8 | 0.000 | 0.000 | 0.000 | | | |
| 39 | | | | | | | | |
| 40 | | Mean | 1.978 | 1.974 | 2.108 | | Sample size | |
| 41 | | Stdev | 1.443 | 1.375 | 1.433 | | 501 | |
| 42 | | | Sums | | Independent | | Maximum value | |
| 43 | | | Poisson | Binomial | Poisson | | 7 | |
| 44 | | SimTable | 5 | 5 | 2 | | | |
| 45 | | 0 | 0 | 0 | 4 | | | |
| 46 | | 0.002 | 2 | 2 | 2 | | | |
| 47 | | 0.004 | 1 | 1 | 1 | | | |

FIGURE **9.13**
**Approximating a Poisson random variable by a Binomial random variable**

random variables with mean 0.1 appear in cells C8:C27 of Figure 9.13, and their sum in cell A4 is also a Poisson random variable with mean $20*0.1 = 2$. So cells A4 and A3 have exactly the same probability distribution and can equally well simulate the number of arrivals in the 20-minute interval.

Then cells D8:D27 in Figure 9.13 contain random variables that equal 0 or 1 with probabilities $1 - \lambda*t = 0.9$ and $\lambda*t = 0.1$, respectively, and cell A5 contains the sum of these 20 {0, 1} random variables. So cell A5 is a Binomial random variable with parameters $n = 20$ and $p = 0.1$. (Recall that a Binomial random variable with parameters $n$ and $p$ can be interpreted as the number of successes in $n$ independent trials, where $p$ is the probability

| | A | B | C | D | E | F | G |
|---|---|---|---|---|---|---|---|
| 1 | Customer type | 1 | 2 | 3 | | NextOfAnyType | |
| 2 | E(WaitTime) | 10 | 15 | 20 | | 4.61538 | E(Wait) |
| 3 | E(ArrivalRate) | 0.1 | 0.06667 | 0.05 | | 0.21667 | E(Rate) |
| 4 | Sim'd ArrivalTime | 13.2374 | 9.27027 | 17.2905 | | 9.27027 | Sim'd time |
| 5 | NextType? | 0 | 1 | 0 | | | |
| 6 | Proby(NextType) | 0.46154 | 0.30769 | 0.23077 | | | |
| 7 | | | | | | | |
| 8 | FORMULAS | | | | | | |
| 9 | B3. =1/B2 | | B23. =B5 | | | | |
| 10 | B3 copied to B3:D3 | | B23 copied to B23:D23 | | | | |
| 11 | F3. =SUM(B3:D3) | | F23. =F4 | | | | |
| 12 | F2. =1/F3 | | B21. =AVERAGE(B24:B524) | | | | |
| 13 | B4. =EXPOINV(RAND(),B2) | | B21 copied to B21:D21,F21 | | | | |
| 14 | B4 copied to B4:D4 | | | | | | |
| 15 | F4. =MIN(B4:D4) | | | | | | |
| 16 | B5. =IF(B4=$F$4,1,0) | | | | | | |
| 17 | B5 copied to B5:D5 | | | | | | |
| 18 | B6. =B3/$F$3 | | | | | | |
| 19 | B6 copied to B6:D6 | | | | | | |
| 20 | | | | | | | |
| 21 | Average | 0.46108 | 0.30339 | 0.23553 | | 4.54644 | |
| 22 | | Next 1? | Next 2? | Next 3? | | TimeNext | |
| 23 | SimTable | 0 | 1 | 0 | | 9.27027 | |
| 24 | 0 | 1 | 0 | 0 | | 6.13836 | |
| 25 | 0.002 | 0 | 0 | 1 | | 8.9201 | |
| 26 | 0.004 | 1 | 0 | 0 | | 7.00113 | |
| 27 | 0.006 | 0 | 0 | 1 | | 4.62873 | |
| 28 | 0.008 | 0 | 0 | 1 | | 4.74802 | |
| 29 | 0.01 | 0 | 0 | 1 | | 1.70336 | |

FIGURE **9.14**
**Arrivals of different types, with different Exponential waiting times**

of success in each trial.) But each cell in D8:D27 will be equal to the adjacent cell in C8:C27 with high probability (because the same RAND( )s in cells B8:B27 are used to generate both the small-mean Poisson random variables in cells C8:C27 and the {0, 1} random variables in cells D8:D27). The mean of the Binomial random variable in cell A5 is $n*p = 20*0.1 = 2$, which is the same as the mean of the Poisson random variables in cells A3 and A4. In general, any Poisson probability distribution can be closely approximated by a Binomial probability distribution with parameters $n$ and $p$ such that the mean of the Binomial ($n*p$) is the same as the mean of the Poisson and the probability parameter $p$ is small (say, $p \le 0.1$).

Let us now consider situations where there may be several different types of customers arriving at the shop, and the waiting time until the next arrival of each type $i$ is always an Exponential random variable with some given mean $\mu_i$, independently of all previous arrivals. Figure 9.14 illustrates such a situation with three types of customers. In this example, the expected waiting time between arrivals is 10 minutes for type 1 customers, 15 minutes for type 2 customers, and 20 minutes for type 3 customers, as listed in cells B2:D2. The expected arrival rates for each type of customer are the reciprocals of these expected waiting times: 1/10 per minute for type 1 customers, 1/15 per minute for type 2 customers, and 1/20 per minute for type 3 customers, as computed in cells B3:D3 (with the formula =1/B2 in cell B3 copied to B3:D3). To simulate the waiting times until the

Myerson_C09.qxd 07-09-2004 09:46 Page 356

next arrival of each type of customer, cells B4:D4 contain independent Exponential random variables with the means given in B2:D2. Cell B4 contains the formula =EXPOINV(RAND( ),B2), and B4 has been copied to B4:D4.

But now suppose that we are interested in the next arrival of any customer, regardless of type. The time until the next arrival of any type is the minimum of the separate waiting times for the various types, as computed in cell F4 of Figure 9.14 by the formula =MIN(B4:D4). There are two important facts that you should know about this next arrival of any type in such a model where each type has Exponential waiting times.

First, the time until the next arrival of any type is also an Exponential random variable, and its mean is the reciprocal of the sum of the expected arrival rates of the separate types. That is, the expected arrival rate of all types together is the sum of the expected arrival rates of the separate types, which is computed for this example by the formula =SUM(B3:D3) in cell F3 of Figure 9.14. The expected waiting time for the next arrival of any type is then the reciprocal of this overall expected arrival rate, as computed in cell F2 by the formula =1/F3.

Second, for each type *i*, the probability of the next type to arrive being type *i* is the expected arrival rate of type *i* divided by the sum of all types' expected arrival rates. So for this example, the probability of the next arrival being a type 1 customer can be computed in cell B6 of Figure 9.14 by the formula =B3/$F$3. Copying B6 to B6:D6 yields the analogous probabilities for the other two types. The probabilities in B6:D6 and the expected waiting time in F2 can be verified by simulation analysis. (See row 21 and below in Figure 9.14.)

## 9.7 Queuing Models

Probability models with Exponential waiting times and Poisson arrivals can be used to analyze sophisticated queuing problems in shops and service centers. To illustrate such analysis, let us consider an example that is modeled in Figure 9.15. Our analysis of this example will also introduce the Iterative Process command in Simtools.

In this example, we are concerned with a shop that sells electronics equipment. A fundamental question that the owners of this shop want to analyze is how many people should be employed to serve customers in this shop. The answer depends on many parameters: the expected rate at which customers arrive, the expected time it takes to serve a customer, the cost of having an idle server in the shop, the goodwill cost of irritating customers by making them wait for a server, the expected length of time that customers are willing to wait, and the cost of losing sales when impatient customers leave without being served. Our basic assumptions about these parameters in this example are summarized in cells A2:A8 of Figure 9.15.

Myerson_C09.qxd 07-09-2004 09:46 Page 357

| | A | B | C | D | E | F | G | H |
|---|---|---|---|---|---|---|---|---|
| 1 | PARAMETERS | | | | | | | |
| 2 | 6 | Arrival rate (customers per hour) | | | | | | |
| 3 | 3 | Service rate (customers per hour) for each server | | | | | | |
| 4 | 3 | Servers | | | | | | |
| 5 | 0.25 | E(TimeWillingToWait) (hrs) for a customer in the queue | | | | | | |
| 6 | 5 | Q'd customer waiting costs ($/hr) | | | | | | |
| 7 | 10 | Idle server costs ($/hr) | | | | | | |
| 8 | 30 | Cost of losing a customer ($) | | | | | | |
| 9 | | | | | | | | |
| 10 | | #Customers in System | | | | In service | In queue | |
| 11 | | STATE | 6 | | | 3 | 3 | |
| 12 | | | | | | | | |
| 13 | | Transitions | # change | Rate | Proby | | | |
| 14 | | Arrival | 1 | 6 | 0.222222 | | | |
| 15 | | Service | -1 | 9 | 0.333333 | | | |
| 16 | | Loss | -1 | 12 | 0.444444 | | | |
| 17 | | | | Sum of rates | | | | |
| 18 | | | | 27 | | | | |
| 19 | 0.037037 | E(Waiting time until next transition) | | | | | | |
| 20 | 0.06693 | Simulated waiting time until next transition (Duration) | | | | | | |
| 21 | Loss | Next transition | | | | | | |
| 22 | | #Customers in System | | | | | | |
| 23 | | UPDATE | 5 | | | | | |
| 24 | | | | Avg costs/hr in simulatn | | | | |
| 25 | | Total time simulated | | Queue | Idle | Lost | TotalCost/hour | |
| 26 | | | 85.48929 | 0.861829 | 13.12337 | 19.65159 | 33.63679 | |
| 27 | | | | Costs: | | | | |
| 28 | OUTPUT | #InSys | Duratn | Queue | Idle | Lost | | |
| 29 | Iteration | 6 | 0.06693 | 1.003956 | 0 | 30 | | |
| 30 | 1 | 0 | 0.000173 | 0 | 0.005188 | 0 | | |
| 31 | 2 | 1 | 0.008106 | 0 | 0.162128 | 0 | | |
| 1027 | 998 | 5 | 0.014818 | 0.148181 | 0 | 0 | | |
| 1028 | 999 | 6 | 0.011069 | 0.166035 | 0 | 0 | | |
| 1029 | 1000 | 5 | 0.018966 | 0.189662 | 0 | 0 | | |
| 1030 | ITERATIVE PROCESS: state=C11, update=C23, output=A29:F1029 | | | | | | | |
| 1031 | FORMULAS | | | | | | | |
| 1032 | F11. =MIN(C11,A4) | | | | B29. =C11 | | | |
| 1033 | G11. =C11-F11 | | | | C29. =A20 | | | |
| 1034 | D14. =A2 | | | | D29. =G11*A6*C29 | | | |
| 1035 | D15. =A3*F11 | | | | E29. =A7*(A4-F11)*C29 | | | |
| 1036 | D16. =G11/A5 | | | | F29. =IF(A21=B16,A8,0) | | | |
| 1037 | D18. =SUM(D14:D16) | | | | C26. =SUM(C30:C1029) | | | |
| 1038 | E14. =D14/$D$18 | | | | D26. =SUM(D30:D1029)/$C$26 | | | |
| 1039 | E14 copied to E14:E16 | | | | D26 copied to D26:F26 | | | |
| 1040 | A19. =1/D18 | | | | G26. =SUM(D26:F26) | | | |
| 1041 | A20. =EXPOINV(RAND(),A19) | | | | | | | |
| 1042 | A21. =DISCRINV(RAND(),B14:B16,E14:E16) | | | | | | | |
| 1043 | C23. =C11+SUMIF(B14:B16,A21,C14:C16) | | | | | | | |

FIGURE **9.15**
**A queuing model simulated as an iterative process**

Customers arrive in this shop according to a Poisson process with a mean arrival rate of 6 customers per hour, as shown in cell A2. That is, the time between each pair of successive customer arrivals is an independent Exponential random variable with mean 1/6 hours (or 10 minutes). The length of time that it takes a server to serve any one customer is assumed to be an Exponential random variable with mean 1/3 hour. Thus, as shown in cell A3 of Figure 9.15, each server can serve customers at a mean rate of three customers per hour in any period of time when the server is not idle (due to a lack of customers in the shop).

Myerson_C09.qxd 07-09-2004 09:46 Page 358

The number of servers in the shop, which has been entered in cell A4, is the decision variable in this model. The analysis in Figure 9.15 estimates the expected cost when three servers are employed in the shop, but the critical question for further analysis is whether the expected cost could be decreased by employing a smaller or larger number of servers.

Notice that three servers have the capacity to serve customers at an expected rate of 3∗3 = 9 customers per hour. So if customers arrived with perfect regularity every 1/6 hour and if there was no variance in service times around the expected value of 1/3 hour, then these three servers would be able to serve everybody without any waiting. Indeed, if everything were so perfectly regular, then two servers would be enough to serve all customers without a queue growing in the shop. But we are assuming that the customers' interarrival times and service times are random variables and are not perfectly regular. So there will be periods when customers arrive more frequently than expected or when service times are longer than expected, and then a queue of customers waiting for service may begin to form, even with three servers at work in the shop. There will be other periods when customers arrive less frequently and servers are idle in the shop.

To account for loss of goodwill when customers have to wait for service, cell A6 in Figure 9.15 specifies a cost of $5 per hour for each customer who is waiting in the queue to be served. But frustrated customers in the queue may also decide to leave when they are not willing to wait any longer. We assume here that the expected time that any customer is willing to wait in the queue is an Exponential random variable with mean 0.25 hours, as shown in cell A5. When such a customer leaves the queue without being served, the cost to the shop is assumed to be $30 in lost profits (that could have been earned from selling to this customer), as specified in cell A8.

Cell A7 in Figure 9.15 specifies that the cost of an idle server is $10 per hour. In this analysis, we will account for the cost of paying a server only when idle, because we are assuming that the labor cost of serving a customer has already been taken into account in the estimated profit per sale in cell A8.

The simulation model in Figure 9.15 treats the shop as an *iterative process* (also called a *Markov process*) that has a *state* which changes at different points in time. At any point in time, the state of this shop can be summarized by the number of customers who are in the shop, which is indicated in our simulation model by the value of cell C11. Among these customers, the number who are being served is indicated in cell F11, which contains the formula

```
=MIN(C11,A4)
```

When there are more customers than servers in the shop, then the excess customers are in the queue, as indicated in cell G11 by the formula

```
=C11–F11
```

The state will change at some time in the near future, when a customer arrives or departs from the shop. In the language of such iterative processes,

Myerson_C09.qxd 07-09-2004 09:46 Page 359

such a change of state is called a *transition*. Our model must simulate our uncertainty about how long we may have to wait for the next transition and about what this next transition will be. Here, there are three possible transitions: a new customer may arrive, a customer in service may depart after completing a purchase, or a customer in the queue may be lost after reaching the limit of patience. These three possible transitions are listed in cells B14:B16 of Figure 9.15. For each possible transition, the net change in the state that it causes is listed in the range C14:C16 (+1 for an arrival, –1 for a completed service or a loss from the queue).

Cells D14:D16 in Figure 9.15 compute rates for each of these possible transitions, which we must now explain. Cell A2 specified that the expected rate of customer arrivals is always six per hour, and cell D14 just echoes this parameter with the formula

=A2

The rate of departures due to completions of service depend on how many customers are being served. If there were just one server who was always kept busy serving one customer after another, then the expected rate of departures of served customers would be three per hour, as specified by cell A3. But if two servers were always kept busy serving one customer after another, then the expected departure rate of served customers would be twice the rate that one server generates. More generally, the expected departure rate of served customers is the expected rate for one server (in A3) multiplied by the number of working servers. But the current number of working servers is specified in cell F11 for this simulation model, and so the current expected rate of departures due to completed service is computed in cell D15 by the formula

=A3*F11

For each customer in the queue, the time until quitting is assumed to have an Exponential distribution with the mean shown in cell A5 (0.25). The reciprocal of this expected waiting time is the expected departure rate when there is one customer in the queue. That is, if we could always keep exactly one customer in the queue (say, by sending in a new customer whenever one departed), then expected rate of loss from the queue would be 1/A5 = 1/0.25 = 4 per hour. But when there are more customers in the queue, any one of them is equally likely to quit at any time (by the no-memory property of the Exponential distribution), and so the expected rate of loss from the queue is proportional to the number in the queue, which is listed in cell G11 for this simulation model. Thus, the expected departure rate of losses from the queue is computed in cell D16 by the formula

=G11/A5

To interpret these expected rates, consider some very short time period, and let $\varepsilon$ denote the length of this period. (You may think of $\varepsilon$ as 0.001 hours, which is 3.6 seconds.) The probability of any transition in this short period is approximately proportional to the length of time $\varepsilon$ that we are

Myerson_C09.qxd 07-09-2004 09:46 Page 360

considering, and the rates shown in cells D14:D16 of Figure 9.15 are the constants of this proportionality relationship. In Figure 9.15, the values of these three cells (6, 9, and 12) tell us that if there are now six customers in the shop (three in service and three in the queue), then during the next $\varepsilon$ fraction of an hour, the probability of a new arrival is approximately $6*\varepsilon$, the probability of a customer completing service is close to $9*\varepsilon$, and the probability of a loss from the queue is close to $12*\varepsilon$. (These approximations are good when the period $\varepsilon$ is short enough that the probability of two transitions within the period is negligible.)

Now recall the two lessons from Figure 9.14 in the previous section, but let us modify the terminology to speak of "transitions" instead of "arrivals" (because some of the transitions here are actually departures): (1) Where the time until the next transition of each type is an independent Exponential random variable, the waiting time until the next transition of any type is an Exponential random variable with mean equal to the reciprocal of the sum of the expected rates for these different types. (2) The probability of any particular type of transition being the next to occur is its expected rate divided by the sum of the expected rates for all types of transitions. In Figure 9.15, the sum of the expected rates for all types of transitions is computed in cell D18 by the formula

```
=SUM(D14:D16)
```

So the expected waiting time until the next transition is computed in cell A19 by the formula

```
=1/D18
```

and the waiting time until the next transition can be simulated in cell A20 by the formula

```
=EXPOINV(RAND( ),A19)
```

The probability of the next transition being a new arrival is computed in cell E14 by the formula

```
=D14/$D$18
```

Then copying E14 to E15 and E16 similarly gives us the probabilities of the next transition being a completion of service or a loss from the queue. Then the type of transition that occurs next can be simulated in cell A21 by the formula

```
=DISCRINV(RAND( ),B14:B16,E14:E16)
```

In this simulation model, the new number of customers in the shop after the next transition can be computed in cell C23 by the formula

```
=C11+SUMIF(B14:B16,A21,C14:C16)
```

In this formula, Excel's SUMIF function searches in the range B14:B16 for cells that contain the same value as cell A21, and it returns the sum of the

Myerson_C09.qxd 07-09-2004 09:46 Page 361

values of the corresponding cells in the range C14:C16. So SUMIF here will simply return the "change" number (1 or –1) in C14:C16 that corresponds to the transition in B14:B16 that matches the simulated next transition in A21.

To simulate the shop after the next transition, we could iteratively copy cell C23 and paste-special its value into cell C11. Iteratively repeating this operation would enable us to simulate the shop through as many transitions as we like. Simtools contains a special procedure called *Iterative Process* for doing this kind of simulation.

The Iterative Process procedure asks us to specify three ranges: the State range, the Update range, and the Output range. The State range is a cell or rectangular range of cells that describe the current state of the process. In this example, the State range is just cell C11, which represents the number of customers currently in the shop. The Update range, which must be the same size as the State range, should contain formulas that compute the next simulated state of the process after the next transition. In this example, the Update range is cell C23, which simulates the number of customers who will be in the shop after the next transition. For each iteration, Iterative Process will copy the Update range and paste-special its values onto the State range.

The Iterative Process's Output range stores the information that we want to record from each iteration of the process. The format of the Output range is similar to the Simulation Tables and the column-input Data Tables that have been used throughout this book. The top row of the Output range should contain the outputs that we want to record in the table below, except the top-left cell of the Output range should be unused. Then at each iteration, the values from this top Output row will be recorded in the next row down in the Output range, continuing down until the bottom row of the Output range has been filled. Iteration numbers will be recorded in the left column of the Output range. The data in Figure 9.15 was generated using the Output range A29:F1029, which contains five output formulas in cells B29:F29 and the recorded values of these five outputs from 1000 iterations in the 1000 rows of B30:F1029. (Rows 32 to 1026 are hidden in Figure 9.15.)

In the top Output row B29:F29 of Figure 9.15, cell B29 shows the current state of the shop, echoed from above by the formula =C11. Cell C29 shows the time (in hours) until the next transition, using the formula =A20. The costs incurred at this iteration are shown in cells D29:F29. Cell D29 computes the cost of customers kept waiting in the queue by the formula

```
=G11*A6*C29
```

(because each of the G11 customers in the queue generates a goodwill cost of A6 dollars per hour during the C29 hours of this iteration). Cell E29 computes the cost of idle servers by the formula

```
=A7*(A4–F11)*C29
```

Myerson_C09.qxd 07-09-2004 09:46 Page 362

(because each of the A4 – F11 idle servers costs A7 dollars per hour during the C29 hours of this iteration). Cell F29 computes the cost of a lost customer (if any) at this iteration by the formula

=IF(A21=B16,A8,0)

(because A8 dollars is lost only when the next transition in A21 matches the "Loss" in B16).

The results from 1000 iterations of this Iterative Process are summarized in cells C26:G26 of Figure 9.15. These 1000 iterations simulate a total period of over 85 hours, as computed in cell C26 by the formula =SUM(C30:C1029). From this simulation data, cell G26 shows an average total cost per hour of $33.63, in which the cost of queued customers is $0.86 per hour, the cost of idle servers is $13.12 per hour, and the cost of lost sales from impatient customers who depart unserved is $19.65 per hour.

The high cost of lost sales suggests that it may be worthwhile to hire another server to keep queue lengths smaller and thus reduce the expected number of departures from the queue. Indeed, when the number of servers in cell A4 of Figure 9.15 is increased to 4, the results of large simulations of this model indicate that expected total costs may decrease to about $28 per hour (including about $20 of idle server costs and about $8 of lost sales).

## 9.8 A Simple Inventory Model

Inventory management is another important application of iterative processes. To introduce some basic techniques for modeling inventory processes, let us consider a simple example about a shop that sells television sets.

The basic parameters that characterize this example are listed in cells A2:A7 of Figure 9.16. Customers wanting to buy a television come into the shop as a Poisson arrival process with a mean rate of 6 per day (cell A2). The shop sells televisions at the price of $400 each (cell A4), and the shop buys its televisions from a wholesale supplier for $250 each (cell A3). To cover the supplier's costs of sending a delivery truck out to this shop, the shop must pay a fixed shipping charge of $500 each time that it orders a new shipment of televisions (cell A6). When there are no televisions in the shop's inventory, then customers will leave and buy from another store instead, so the shop will incur a big cost of lost sales if it runs out of inventory too often. The shop has a storage area where it can keep many televisions in inventory, but the cost of holding such inventories is accounted to be $0.60 per television per day in inventory (cell A5). (This inventory holding cost may include the opportunity cost of storage space, insurance costs, interest on the prepaid cost of the televisions, etc.) When the shop calls its supplier to order a new batch of televisions, the new order will take either 4 or 5 days to be delivered, each with probability 1/2.

| | A | B | C | D | E | F | G | H | I | J | K | L | M | N | O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | PARAMETERS | | | | | | | | | | | FORMULAS | | | |
| 2 | 6 | Expected rate of demand (units per day) | | | | | | | | J9. | | =A5*(G10+G9)/2+A6*A2/(G10-G9) | | | |
| 3 | 250 | Cost per unit | | | | | | | | J10. | | =(2*A6*A2/A5)^0.5 | | | |
| 4 | 400 | Selling price per unit | | | | | | | | | | I14. | =POISINV(RAND(),A2) | | |
| 5 | 0.6 | Inventory holding cost per unit per day | | | | | | | | | | G15. | =MIN(G14,I14) | | |
| 6 | 500 | Fixed cost of ordering a new shipment | | | | | | | | | | G16. | =G14-G15+SUM(B14:F14) | | |
| 7 | Shipment delivery time may be 4 or 5 days, each proby=1/2. | | | | | | | | | | | G17. | =IF(G16<=G9,G10-G16,0) | | |
| 8 | INVENTORY POLICY | | | | | | | With instant delivery: | | | | B20. | =G17 | | |
| 9 | When Inventory+OnOrder falls to | | | | | | 30 | E(CostRate) | | 78 | | C20. | =B14 | | |
| 10 | | | | then order up to | | | 130 | | EOQ | 100 | | C20 copied to C20:E20 | | | |
| 11 | SIMULATION MODEL | | | | | | | | | | | F20. | =IF(RAND()<0.5,0,E14) | | |
| 12 | Units on order, by days traveled | | | | | | Inventory | | | | | G20. | =G14-G15+F14+(E14-F20) | | |
| 13 | | 0 | 1 | 2 | 3 | 4 | on hand | | Demand | | | B27. | =B14 | | |
| 14 | STATE | 0 | 0 | 0 | 0 | 0 | 72 | | 4 | | | B27 copied to B27:G27 | | | |
| 15 | | | | | | | 4 | Sales | | | | H27. | =IF(G17>0,A6,0) | | |
| 16 | | | | | | | 68 | Inventory+OnOrder | | | | I27. | =(I14-G15)*(A4-A3) | | |
| 17 | | | | | | | 0 | New order (if any) | | | | G23. | =A5*AVERAGE(G28:G2027) | | |
| 18 | | On order, by days traveled | | | | | Invntry | | | | | H23. | =AVERAGE(H28:H2027) | | |
| 19 | | 0 | 1 | 2 | 3 | 4 | on hand | | | | | I23. | =AVERAGE(I28:I2027) | | |
| 20 | UPDATE | 0 | 0 | 0 | 0 | 0 | 68 | | | | | J23. | =SUM(G23:I23) | | |
| 21 | | | | | | | | | | | | | | | |
| 22 | | | | | | | Holding | Shippng | Lost | Total | | | | | |
| 23 | | | | Average costs/day | | | 33.83 | 28.25 | 20.78 | 82.85 | | | | | |
| 24 | | | | | | | | | | | | | | | |
| 25 | | On order, by days traveled | | | | | Invntry | Costs of: | | | | | | | |
| 26 | OUTPUT | 0 | 1 | 2 | 3 | 4 | on hand | Shippng | LostSales | | | ITERATIVE PROCESS: | | | |
| 27 | Iteratio | 0 | 0 | 0 | 0 | 0 | 72 | 0 | 0 | | | state=B14:G14, | | | |
| 28 | 1 | 100 | 0 | 0 | 0 | 0 | 30 | 0 | 0 | | | update=B20:G20, | | | |
| 29 | 2 | 0 | 100 | 0 | 0 | 0 | 25 | 0 | 0 | | | output=A27:I2027. | | | |
| 30 | 3 | 0 | 0 | 100 | 0 | 0 | 20 | 0 | 0 | | | | | | |
| 31 | 4 | 0 | 0 | 0 | 100 | 0 | 11 | 0 | 0 | | | | | | |
| 32 | 5 | 0 | 0 | 0 | 0 | 100 | 6 | 0 | 300 | | | | | | |
| 33 | 6 | 0 | 0 | 0 | 0 | 0 | 100 | 0 | 0 | | | | | | |
| 34 | 7 | 0 | 0 | 0 | 0 | 0 | 92 | 0 | 0 | | | | | | |

FIGURE **9.16** **A simple inventory model**

An *inventory policy* for this shop is a rule that specifies when the shop should order more televisions and how many it should order. We will assume here that the shop's manager checks inventory at the end of each day. Then for some pair of numbers ($s$, $S$), the manager could implement any policy of the form: "whenever you have $s$ or fewer televisions, then order enough to bring the stock up to $S$ televisions." The spreadsheet in Figure 9.16 is set up to analyze a policy of this form with $s = 30$ in cell G9 and $S = 130$ in cell G10. Under this policy, if the manager found 26 televisions in inventory at the end of a day when no televisions are on order, then the manager would place a new order for $130 - 26 = 104$ televisions, because $26 \le 30$.

With a profit margin of $400 - 250 = \$150$ per television and an expected demand rate of 6 customers per day, the shop would expect a profit rate of $6*150 = \$900$ per day if there were no costs of delivery, no inventory costs, and no cost of lost customers when inventories run out. But actual profits will be less because of these costs that depend on how the shop manages its inventory. The purpose of this spreadsheet is to generate an estimate of these inventory-dependent costs for the inventory policy that is indicated in cells G9:G10. After evaluating this inventory policy by simulation analysis,

Myerson_C09.qxd 07-09-2004 09:46 Page 364

we should redo our analysis with other values of the parameters in G9:G10 to search for an inventory policy that yields the lowest expected costs.

But before we analyze this policy, let us see how it would work in a simpler situation where new orders could be placed at any time and would be delivered instantly. In this case, we would never have to worry about losing any customers. The only costs would be the inventory holding costs and the fixed costs of delivery. Inventory would gradually decrease from G10 = 130 to G9 = 30 and then would instantly jump back up to 130. In the long run, the expected inventory on any randomly sampled day would be the average of these upper and lower bounds G10 and G9, and so the expected inventory holding cost would be

A5*(G10 + G9)/2 = 0.6*(130 + 30)/2 dollars per day

New orders would always be for G10 – G9 = 130 – 30 = 100 televisions in this simplified case. With demand averaging 6 per day in a Poisson process, the probability of a new order occurring on any randomly sampled day would be 6/100, and so the expected shipping cost would be

A6*A2/(G10 – G9) = 500*6/100 dollars per day

Thus, with instant delivery at any time, the expected long-run cost per day of shipping and holding inventories could be computed by the formula

=A5*(G10+G9)/2 + A6*A2/(G10–G9)

which appears in cell J9 of Figure 9.16. If you ask Solver to adjust the policy parameters in cells G9:G10 to minimize this expected daily cost in J9, subject to the constraints G9:G10 >= 0, then Solver will choose G9 = 0 and G10 = 100, which represents the policy of ordering 100 units whenever inventory reaches 0. More generally, it can be shown with calculus that this instant-delivery cost rate in cell J9 can always be minimized by setting G9 equal to 0 and setting G10 equal to the value that is computed in cell J10 by the formula

=(2*A6*A2/A5)^0.5

This order quantity is called the *economic order quantity,* or *EOQ*.

Now with deliveries actually taking 4 or 5 days, the shop would lose too many valuable customers if it reordered only when inventory reached 0. So the shop should try a higher reorder level in cell G9, and so in Figure 9.16 we try letting G9 equal 30. But the EOQ is still a good guess for the amount that should be ordered when inventory reaches this G9 level, and so in Figure 9.16 we try letting G10 equal 130, which is G9 + J10.

To evaluate this policy, we simulate the shop by an Iterative Process model in which each iteration represents one day. At the beginning of each day, the state of the shop is described by variables listed in cells B14:G14 of Figure 9.16. Cell G14 here lists the number of televisions that are on hand in inventory. Cells B14:F14 list, respectively, the numbers of undelivered

Myerson_C09.qxd 07-09-2004 09:46 Page 365

televisions that are on order and that have traveled through the supplier's delivery system for 0 to 4 days.

The number of customers who will want to buy a television in the shop on this day is a Poisson random variable that is simulated in cell I14 of Figure 9.16 by the formula

=POISINV(RAND( ),A2)

Cell G15 computes the day's sales, which require both inventory and demand, by the formula

=MIN(G14,I14)

The number of televisions that the shop will have in inventory or undelivered past orders at the end of this day is computed in cell G16 by the formula

=G14–G15+SUM(B14:F14)

Then, according to our inventory policy in G9:G10, the new orders that the shop should place at the end of this day can be computed in cell G17 by the formula

=IF(G16<=G9,G10–G16,0)

The Update range for the Iterative Process in Figure 9.16 is B20:G20. The formula =G17 has been entered into cell B20 to represent the fact that any televisions that are ordered today will be "on order" tomorrow morning and will then have traveled 0 days. The formula =B14 has been entered into cell C20, and C20 has been copied to C20:E20, to represent the fact that any units on order which had traveled 0 or 1 or 2 days before this morning will have traveled one day more (1 or 2 or 3, respectively) tomorrow morning. Cell F20 simulates the number of units that will still be undelivered after 4 days of travel tomorrow by the formula

=IF(RAND( ) <0.5,0,E14)

because any units on order that had traveled 3 days before this morning will have a probability 1/2 of being delivered tomorrow. Tomorrow's inventory includes today's inventory (G14) less today's sales (G15) plus any units that had traveled 4 days this morning (F14) or had traveled 3 days this morning and are not still undelivered tomorrow (E14 – F20). So the inventory on hand tomorrow morning can be computed in cell G20 by the formula

=G14–G15+F14+(E14–F20)

At each iteration, the values from this Update range B20:G20 are copied back to the State range B14:G14 to simulate the next day of the shop.

The Output table in A17:I2027 of Figure 9.16 lists, from 2000 iterations of this process, the values of the outputs that are computed in cells B27:I27. These outputs in B27:I27 include the quantities on order and in inventory and the daily costs of shipping and lost sales. Cells B27:G27 simply echo

Myerson_C09.qxd 07-09-2004 09:46 Page 366

the State range B14:G14. The daily shipping costs are computed in cell H27 by the formula

=IF(G17>0,A6,0)

The daily costs of lost sales are computed in cell I27 by the formula

=(I14–G15)*(A4–A3)

because each of I14 – G15 lost customers would have generated A4 – A3 dollars of profit. Row 28 shows that the Iterative Process was started on a day when 30 units were available in inventory and 100 units had just been ordered. Then from this table of outputs from 2000 simulated days, cells G23:J23 compute that inventory holding costs averaged about $34 per day, shipping costs averaged about $28 per day, and costs of lost sales averaged about $21 per day, making a total average cost of about $83 per day. Lower average daily costs may be generated by trying other inventory policies in cells G9:G10, however.

## 9.9 Project Length and Critical Tasks

In this section, we consider the length of time that will be needed to complete a project that consists of many tasks, where some tasks may require others to be completed first and the time required to complete each task is uncertain. To illustrate how such problems can be handled, let us consider a case in which a construction company has a contract to build a major new bridge and new roads connecting to the nearby highways. This construction project can be separated into eight activities, which we may call "activity A" through "activity H." For each activity, the following table lists the mean and standard deviation of the possible duration times (in days) for the activity and the predecessors that must be completed before it can begin. The project will be finally completed when activities F and H are both finished.

| *Activity* | *Mean Time* | *Standard Deviation* | *Predecessors* |
|---|---|---|---|
| A (pour foundation) | 20 | 7 | — |
| B (build support towers) | 50* | 10* | A |
| C (assemble cables) | 60 | 12 | — |
| D (install cables) | 15 | 3 | B,C |
| E (excavate connecting road) | 65 | 30 | — |
| F (pave connecting road) | 35 | 15 | E |
| G (install bridge roadbed) | 30 | 5 | D |
| H (pave bridge) | 10 | 3 | G |

*Activity B's time could be reduced 20% by a more expensive subcontractor.

Myerson_C09.qxd 07-09-2004 09:46 Page 367

The contract for this project specifies a goal of completing the project within 130 days, and the construction company will have to pay a penalty of $100,000 per day for each day that the project continues beyond 130 days. There is no penalty or bonus associated with completing the project in fewer than 130 days, however.

The mean and standard deviations for activity B are marked by an asterisk (*) because the construction company is considering an alternative to hire a more expensive subcontractor to work on this activity. If the more expensive subcontractor is hired, then duration times for activity B will be reduced by 20% (to mean 40 and standard deviation 8).

In this example, we have the problem of estimating the length of time that will be needed to complete the project. This project consists of several separate activities, each of which may take an uncertain amount of time. Some of these activities have *predecessors*, that is, other activities that must be completed first. For example, activity D (installing the bridge's cables) can begin only after the completion of activities B and C (building the support towers and assembling the cables), which are the predecessors of activity D.

Figure 9.17 shows a spreadsheet model to describe this situation. To describe our beliefs about the activities' duration times, cells B3:B10 list the expected duration time (in days) for each activity, and the corresponding standard deviations are listed in cells C3:C10. The duration of an activity must be a nonnegative number, and so it may be natural to use the Lognormal or Gamma distribution. We have seen that the Gamma distribution may be appropriate to describe some duration times (the time until the *k*th arrival in a model where successive arrivals are separated by Exponential waiting times), and so each activity's duration is modeled here as a Gamma random variable. Simtools provides a function GAMINV that can be used to generate a Gamma random variable parameterized by its mean and standard deviation. With this function, the duration of activity A is simulated in cell D3 by the formula

```
=ROUND(GAMINV(RAND( ),B3,C3),0)
```

Copying D3 to cells D3:D10 gives us our simulated duration times for all eight activities. In these formulas, ROUND(•, 0) is used to round the random variable to the nearest integer (0 decimal places) so that the duration is a whole number of days.

(*Warning:* Do not confuse the Simtools function GAMINV with the similarly named Excel function GAMMAINV. Both functions can make Gamma random variables, but GAMINV takes the mean and standard deviation as its parameters after the RAND( ). GAMMAINV instead takes two other parameters called "alpha" and "beta" that are difficult to interpret. When in doubt, you should use the Insert:Function dialog box for a reminder of which parameters are used in each function.)

Let us number the days of work starting with day 0, the first day of work on the project. In the spreadsheet, the starting days (counting forward from

| | A | B | C | D | E | F | G | H | I | J | K | L |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | PROJECT LENGTH MODEL | | | | | | | With reduced B time | | | | |
| 2 | Task | Mean | Stdev | Duratn | Start | Finish | | Duratn | Start | Finish | | |
| 3 | A | 20 | 7 | 16 | 0 | 16 | | 16 | 0 | 16 | | |
| 4 | B | 50 | 10 | 49 | 16 | 65 | * | 39 | 16 | 55 | | |
| 5 | C | 60 | 12 | 54 | 0 | 54 | | 54 | 0 | 54 | | |
| 6 | D | 15 | 3 | 21 | 65 | 86 | | 21 | 55 | 76 | | |
| 7 | E | 65 | 30 | 30 | 0 | 30 | | 30 | 0 | 30 | | |
| 8 | F | 35 | 15 | 77 | 30 | 107 | | 77 | 30 | 107 | | |
| 9 | G | 30 | 5 | 36 | 86 | 122 | | 36 | 76 | 112 | | |
| 10 | H | 10 | 3 | 14 | 122 | 136 | | 14 | 112 | 126 | | |
| 11 | | | | Completion time | | | | Completion with B reduced | | | | |
| 12 | | | | | 136 | days | | | 126 | days | | |
| 13 | Penalty ($1000s) per day | | | | | | | Is B critical? | | | | |
| 14 | 100 | | | | | | | | 1 | | | |
| 15 | after | | | | | Penalty | | | | Penalty with B reduced | | |
| 16 | 130 | days | | | | 600 | | | | 0 | | |
| 17 | | | | | | | | | | | | |
| 18 | FORMULAS | | | | | | | | | | | |
| 19 | D3. =ROUND(GAMINV(RAND(),B3,C3),0) | | | | | | | | | | | |
| 20 | D3 copied to D3:D10 | | | | | | | | | | | |
| 21 | F3. =E3+D3 | | | | | CONSTRUCTION PROJECT SEQENCE DIAGRAM | | | | | | |
| 22 | F3 copied to F3:F10 | | | | | A → | | B | | | | |
| 23 | E4. =F3 | | | | | | | | | | | |
| 24 | E6. =MAX(F4,F5) | | | | | | | | | | | |
| 25 | E8. =F7 | | | Start → | | C → | | D → | G → | H | | |
| 26 | E9. =F6 | | | | | | | | | | | |
| 27 | E10. =F9 | | | | | | | | | | | Finish |
| 28 | H3. =D3 | | | | | E → | | F | | | | |
| 29 | H3 copied to H5:H10 | | | | | | | | | | | |
| 30 | H4. =D4-0.2*D4 | | | | | | | | | | | |
| 31 | E3:F10 copied to I3:J10 | | | | | | | | | | | |
| 32 | E12. =MAX(F3:F10) | | | | | | | | | | | |
| 33 | I12. =MAX(J3:J10) | | | | | | | | | | | |
| 34 | I14. =IF(I12<E12,1,0) | | | | | | | | | | | |
| 35 | F16. =MAX(E12-$A$16,0)*$A$14 | | | | | | | | | | | |
| 36 | J16. =MAX(I12-$A$16,0)*$A$14 | | | | | | | | | | | |

FIGURE **9.17**
**A project-length model**

the first day of work) of the various activities are listed in cells E3:E10, and the completion days of the various activities are listed in cells F3:F10. Work on activities A, C, and E can start on day 0, because they have no predecessors, so the value 0 is entered into cells E3, E5, and E7. The completion day of an activity equals its starting day plus its duration time, and so the formula

=D3+E3

is entered into cell F3, and F3 is copied to F3:F10. Each activity that has one or more predecessors can start only when these predecessors are all completed, and so the starting day of such an activity must be the maximum of the completion days of its predecessor. For example, we enter the formula

=MAX(F4,F5)

into cell E6 because activity D has activities B and C as predecessors. Similarly, cell E4 contains the formula =F3 because activity B has activity A as a predecessor, cell E8 contains the formula =F7 because activity F has

Myerson_C09.qxd 07-09-2004 09:46 Page 369

activity E as a predecessor, cell E9 contains the formula =F6 because activity G has activity D as a predecessor, and cell E10 contains the formula =F9 because activity H has activity G as a predecessor. The final completion time, when all activities are completed, is computed in cell E12 by the formula =MAX(F3:F10).

In such a project-length problem, we say that an activity is *critical* if reducing the duration of this activity would decrease the time needed to complete the entire project. That is, if an activity is not likely to be critical, then efforts to reduce its duration are not likely to have any impact on the completion time of the project. So identifying the likely critical activities can help a planner to focus on the activities where accelerating the work would improve final completion times.

To identify whether activity B is critical, we reconstruct the model with duration times exactly the same except that activity B's duration time is reduced, and we ask whether the final completion time of this reconstructed model is the same. To build this reconstructed model in the range H3:J10 of Figure 9.17, we first enter the formula

```
=D3
```

into cell H3, and we then copy H3 to the range H3:H10. But next we must reedit cell H4, which represents the duration time of activity B, so that its value will be less the value of cell D4. If we only wanted to evaluate the probability of activity B being critical, then we could let H4 be less than D4 by any positive amount; for example, we could enter the formula =D4–1 into cell H4. But in this case, we are asked to specifically consider the possibility of reducing activity B's duration by 20%, and so we enter the formula

```
=D4–0.2*D4
```

into cell H4. With these formulas, each cell in H3:H10 other than H4 exactly echoes the activity duration time from the corresponding cell in D3:D10 (it does not make another independent draw from the same distributions), and cell H4 takes a value strictly less than B's duration time in cell D4. Next we simply copy the range E3:F10 to I3:J10, and we copy cell E12 to I12. The result is that cell I12 displays the project's final completion time when the duration time of activity B is reduced but all other activity duration times are the same as in D3:D10. Then the formula

```
=IF(I12<E12,1,0)
```

in cell I14 has the value 1 if the activity B is critical, but it has the value 0 if activity B is not critical.

The late penalties with regular and shortened B durations are computed in cells F16 and J16 of the spreadsheet. To answer the questions in the case, we can make a large simulation table with the values of cells E12 and I12 (as well as I14, F16, and J16 for convenience) in the model-output row. In data from 1001 simulations (not shown in the Figure 9.17 here), activity B

Myerson_C09.qxd 07-09-2004 09:46 Page 370

was critical about 56% of the time, so we may estimate that the probability of activity B being critical is about 0.56. With regular B durations, the final completion time in these simulations had a mean value of approximately 133 days, with a standard deviation of 18 days, and the mean penalty cost was about $812,000. With shortened B durations but all other random variables held fixed in this simulation table, the final completion time had a mean of approximately 128 days, with a standard deviation of 19 days, and the mean penalty cost was about $600,000. The difference between these mean penalty costs (approximately $212,000) is our estimate of the expected value to the construction company of reducing activity B's duration time by 20%.

## 9.10 Summary

This chapter introduced a variety of techniques for analyzing and simulating the performance of dynamic processes in which things change over time.

We began with a discussion of the net present value of a series of incomes to be earned at different times in the future. With any given interest rate, the net present value is the amount of money that if deposited now in a bank account that will pay this interest rate, would just enable us to withdraw these income amounts at the designated times. Conversely, if these incomes are deposited into the bank when they are earned, then the amount of money that we will have to spend in the long run will be the same as if we had deposited this net present value today. Excel's NPV function can be used to compute the net present value of a stream of incomes.

Next we considered a variety of models for forecasting time series. We showed in the *Global Sports* case how a good forecasting model may combine some parameters that are statistically estimated from historical data and other parameters that are directly assessed by subjective judgements.

In building time-series forecasting models, we generally assume that certain values are stationary random variables drawn from the same probability distribution in every time period. For such stationary random variables, we can use past data to estimate the distribution out of which they will be drawn in the future. But as we showed in the Goeing Aircraft case, different people may get different time-series models for the same situation with the same data if they make different assumptions about which random variables are stationary. These stationarity assumptions must ultimately be based on the decision maker's subjective beliefs about what underlying factors are likely to be stable over time. Thus, there is not necessarily any one "correct" model, and we may want to compare analyses based on different models that are based on different stationarity assumptions.

Myerson_C09.qxd 07-09-2004 09:46 Page 371

The ideas of Brownian motion were introduced to show how we may construct dynamic time-series models that have different time periods from the data that we have been given. If some value (say, the price of a stock) evolves over time according to a log-Brownian motion with some log-drift $D$ and some log-volatility $V$, then this value's growth ratio over any time period of length $T$ is a Lognormal random variable with log-mean $D*T$ and log-standard-deviation $D*(T\text{^}0.5)$, and the growth ratios over separate time periods will be independent random variables.

We then considered a simple repeated investment process where a decision maker can repeatedly make investments in assets that have growth ratios drawn independently from the same probability distribution at each round of the process. We considered stationary strategies that invest the same fractions of wealth at every round. The stationary strategy that maximizes the expected logarithmic growth rate of the decision maker's wealth is called log-optimal. If it is implemented over many repetitions of this process, then, with high probability, the log-optimal strategy will generally yield higher long-run payoffs than any other stationary strategy.

Next we considered an arrival process where the times between arrivals are independent random variables drawn from an Exponential distribution with some mean $\mu$. Such Exponential distributions have a convenient no-memory property, which implies that the probability of an arrival in the next minute is always independent of when arrivals have occurred in the past. With such Exponential waiting times, the expected number of arrivals per unit of time is $\lambda = 1/\mu$, and the number of arrivals in any time interval of length $T$ is a Poisson random variable with mean $\lambda*T$. The numbers of arrivals in separate time intervals are then independent Poisson random variables. The time until the $k$th arrival in such a process is a Gamma random variable with mean $\mu*k$ and standard deviation $\mu*(k\text{^}0.5)$.

We considered the case of several different types of customers, where the times between the arrivals of customers of each type $i$ are independent Exponential random variables with some given mean $\mu_i$. Then the expected rate of arrivals of type $i$ customers is $\lambda_i = 1/\mu_i$, the waiting time until the next arrival (of any type) is an Exponential random variable with mean $1/(\sum_j \lambda_j)$, and the probability of the next arrival being a customer of type $i$ is $\lambda_i/(\sum_j \lambda_j)$.

We applied these basic ideas about Exponential/Poisson arrival processes to build a simple queuing model of a shop where customers arrive at random, service times are Exponential random variables, and customers who are waiting to be served may randomly decide to leave. This queuing system was modeled as an iterative (or Markov) process, where the state of the process at any point in time is the number of customers in the shop and each iteration represents an arrival or departure from the queue. We simulated this model over time by using Simtools's Iterative Process command, which iteratively copies an update range (that simulates the next value of

Myerson_C09.qxd 07-09-2004 09:46 Page 372

the process's state) and pastes its values onto the range that represents the current state of the process, while filling a selected output range with data from each iteration of the process. We also showed how the problem of managing inventories in a shop can be simulated as an iterative process, where each iteration represents a day, and the state of the process lists the quantities that are available in inventory or have been recently ordered for delivery.

Finally, we considered a project-length problem, in which a given project is composed of various activities or tasks of uncertain duration, where some of these activities cannot be started until certain other activities are completed. We showed how to estimate the probability that any specific activity will be critical, in the sense that reducing the duration of this activity would reduce the duration of the entire project.

Excel functions in this chapter that were not previously emphasized include NPV, SUMIF, and COUNTIF, as well as the Simtools functions EXPOINV, POISINV, and GAMINV. Simtools's Iterative Process command was also introduced.

## Exercises

**1.** Suppose that there are just two alternative investments available: cash and a risky stock fund. The value of this stock fund can vary substantially every month. From our analysis of past data, we have found that the monthly logarithmic growth rate for the stock fund has a mean of 0 and a standard deviation of 0.33. So let us suppose that in each of the 12 months of the year, the stock fund's growth ratio during the month will be an independent Lognormal random variable with log-mean 0 and log-standard-deviation 0.33. To keep things simple, suppose that cash earns no interest, and there is no penalty for buying or selling the stock fund at any time. Three investors named Alpha, Beta, and Gamma each start the year with an initial wealth of $10,000, but they follow different investment strategies, as follows.

*Alpha* has decided to keep $5000 in cash and will invest the rest of his wealth in the stock fund, holding it there for the next 12 months.

*Beta* is a more active investor and plans to readjust her investment at the beginning of each month so that half of her wealth is in cash and the other half is in the stock fund. So in the first month her portfolio looks like Alpha's, but thereafter she may transfer money between her cash box and the stock fund each month, depending on whether the stock fund went up or down in value during the previous month.

*Gamma* is the most active investor. At the beginning of each month, he will toss a fair coin to determine whether he feels lucky. If the coin comes up Heads, then he will hold all of his current wealth in the stock fund for the month. But if the coin comes up Tails, then he will hold all his current wealth in cash for the month. (Of course, the outcome of his coin toss each month will be independent of the subsequent growth of the stock fund.)

**(a)** Make a simulation model of the growth ratios of the stock fund over the 12 months of the year and the resulting performance of these three investors' portfolios. Make a simulation table that lists the final wealths of the three investors at the end of 12 months, including data from at least 1000 simulations. Use this data to answer the following questions.

**(b)** For each of the three investors, estimate the *expected value* and the *standard deviation* of his or her final wealth after 12 months.

**(c)** For each of the three investors, estimate the probability that his or her final wealth will be

Myerson_C09.qxd 07-09-2004 09:46 Page 373

greater than the initial $10,000 at the end of 12 months.

**(d)** For each of the three investors, estimate the three quartile boundary points (the values with cumulative probability 1/4, 1/2, and 3/4) for his or her final wealth at the end of 12 months.

**2.** Suppose that I invest a fraction $\theta$ of my wealth in the risky stock fund from Exercise 1, and I hold the other $1 - \theta$ fraction of my wealth in cash. If **X** denotes the return next month per dollar invested today, then each dollar today will yield $\theta * \mathbf{X} + (1 - \theta)$ for me next month. Recall that **X** is a Lognormal random variable with log-mean 0 and log-standard-deviation 0.33. Based on a sample of at least 1000 draws from the distribution of **X**, estimate my expected logarithmic growth rate $E(\mathrm{LN}(\theta * \mathbf{X} + (1 - \theta)))$ as a function of $\theta$. Numerically compute these estimates for $\theta = 0, 0.1, 0.2, \ldots, 0.9, 1.0$, and make a chart to summarize your results. What fraction should I invest if I want to maximize the expected logarithmic growth rate of my total wealth?

**3.** Suppose that you have some opportunity to bet on a strangely bent coin coming up Heads. If you decide to bet, then the stakes are $10. That is, you win $10 if the coin comes up Heads, but you lose $10 if the coin comes up Tails. You have the option to play this game up to 12 times, and you can decide whether to play again after seeing the result of each coin toss. Once you decide to stop playing, the coin will be put away, and you will not see it any more. Because you have never seen a bent coin like this before, you are not sure about how frequently it will tend to come up Heads in the long run. Suppose that your beliefs about this long-run frequency of coming up Heads, given everything that you know now, can be described by a Beta probability distribution mean of 0.4 and standard deviation of 0.25.

Based on at least 2000 simulations, estimate the *expected value* and *standard deviation* of your net total winnings, and the *probability* that your net total winnings would be positive, for each of the following strategies.

**(a)** *Strategy 1.* Play the game 12 times, no matter how many times you may win or lose.

**(b)** *Strategy 2.* Start playing, but then stop when your net winnings are positive. (That is, stop as soon as you have won more times than you have lost. If that never happens after 12 plays, then you will have to stop in any case.)

**(c)** *Strategy 3.* Start playing, but then stop when your net winnings are negative. (That is, stop as soon as you have lost more times than you have won. If that never happens after 12 plays, then you will have to stop in any case.)

Also use your data from 2000 simulations to answer the following questions about the results of the first coin toss:

**(d)** Estimate the probability that the bent coin will come up Heads on the first toss.

**(e)** If you learned that the bent coin came up Heads on the first toss, then, given this information, what would be the conditional expected value and the conditional standard deviation of the coin's long-run frequency of coming up Heads?

**4.** A couple has $60,000 that they want to invest for 10 years (anticipating their young children's future college expenses), and they have identified two mutual funds in which they plan to invest. In each year, the logarithmic rates of return for these funds are Multivariate Normal random variables. The logarithmic rate of return for Fund 1 each year has mean 0.07 and standard deviation 0.25. The logarithmic rate of return for Fund 2 each year has mean 0.06 and standard deviation 0.15. The correlation coefficient between the two funds' logarithmic rates of return each year is 0.65. The funds' rates of return in any one year are independent of their past and future rates of return. The couple wants to evaluate the risky returns from their investments 10 years from now using a constant risk tolerance of $50,000.

**(a)** First let us consider plans in which the couple invests some amount $\$X$ in Fund 1 now and invests all the rest of their available cash, $\$60{,}000 - X$, in Fund 2 now. To be specific, consider any amount $\$X$ between $0 and $60,000 that is a multiple of $10,000: $X = 0$; $X = 10{,}000$; $X = 20{,}000$; $X = 30{,}000$; $X = 40{,}000$; $X = 50{,}000$; and $X = 60{,}000$. Suppose that, after this initial investment, the couple will hold these investments for 10 years, reinvesting all returns in each fund. Based on data from at least 400 simulations of this situation, make a table and a chart showing, for each of these seven possible values of $X$: the *expected* total value of the couple's portfolio 10 years from now, the low amount that has *cumulative probability 0.1* for the value of their portfolio 10 years from now, and the couple's *certainty equivalent* of their risky portfolio 10 years from now (with constant risk tolerance $50,000).

**(b)** Next, let us consider plans in which the couple invests equally in the two funds but also puts some cash into risk-free short-term bonds that

pay 4% interest per year (compounded annually). That is, consider plans in which the couple invests some amount $\$Y$ in Fund 1 now, invests the same amount $\$Y$ in Fund 2 now, and invests the rest of their available cash, $\$60{,}000 - 2*Y$, in the 4% bonds. To be specific, consider any amount $Y$ between 0 and 30,000 that is a multiple of 5000: $Y = 0$; $Y = 5000$; $Y = 10{,}000$; $Y = 15{,}000$; $Y = 20{,}000$; $Y = 25{,}000$; $Y = 30{,}000$. Suppose that, after this initial investment, the couple will hold these investments for 10 years, reinvesting all returns in each fund. Make a table and a chart showing, for each of these seven possible values of $Y$, the couple's *certainty equivalent* of their risky portfolio 10 years from now (with risk tolerance $50,000).

**(c)** In parts (a) and (b), we assumed that all initial investments would be held with reinvestment for 10 years. Now let us consider a more complex investment plan, as follows. The couple initially invests $30,000 in Fund 1 and another $30,000 in Fund 2. At the end of each year thereafter, if the value of their investment in either fund has grown larger than $30,000, then they will sell the excess over $30,000 and invest this excess in 4% bonds; but if their investment in the fund is worth less than $30,000 at the end of the year, then they will simply hold this investment into the next year. Thus they will begin each year with $30,000 or less in each mutual fund, plus a growing balance in safe bonds that pay 4% annual interest. For this investment plan, estimate the *expected* total final value of the couple's portfolio in 10 years and the couple's *certainty equivalent* of their final portfolio in 10 years (with risk tolerance $50,000). Also, show the *cumulative risk profile* for their total final portfolio value 10 years from now. (Ignore any effects of taxes here.)

**5.** A construction project consists of the activities listed in the following table:

| *Activity* | *Predecessors* | *Duration (months):* | *Mean* | *Standard Deviation* |
|---|---|---|---|---|
| A (get funds) | — | | 6 | 0.6 |
| B (design) | A | | 8 | 1.3 |
| C (site prep) | A | | 2 | 0.2 |
| D (foundation) | B,C | | 2 | 0.3 |
| E (frame, roof) | D | | 3 | 1.0 |
| F (exterior work) | E | | 3 | 0.6 |
| G (interior work) | D | | 7 | 1.5 |
| H (landscape) | F,G | | 5 | 1.2 |

**(a)** Estimate the probability that the project will take less than 30 months to complete and the probability that the project will take more than 36 months to complete.

**(b)** Plot your estimated (inverse) cumulative distribution of completion times.

**(c)** Estimate the probability that activity B will be critical, the probability that activity C will be critical, and the probability that activity G will be critical.

**6.** The Ajax Manufacturing Company is introducing a new quad-speed widget.

Our marketing experts figure that sales revenue in the first year has a probability 0.25 of being less than $3.0 million, probability 0.50 of being less than $5.0 million, and probability 0.75 of being less than $8.4 million. You may assume that beliefs about first-year sales fit a Generalized Lognormal probability distribution.

We also anticipate that in each year after the first, the logarithmic growth rate of sales revenue will be independently drawn from a Normal distribution with mean 0.04 and standard deviation 0.1.

Our cost data suggest that the expected value of variable costs should be 70% of sales revenue, but this fraction may vary from year to year. To describe our uncertainty, you may assume that the ratio of variable costs to sales revenue will be an independent random variable each year drawn from a Beta distribution with mean 0.70 and standard deviation 0.05. Fixed costs of administration and advertising each year will be $1.5 million.

Ajax will evaluate this new product by the net present value (NPV) of profits that it generates over 10 years, using a 10% discount rate.

**(a)** Construct a spreadsheet model to simulate the first 10 years of production of this new product.

**(b)** Based on at least 1000 simulations, show the expected value, the standard deviation, and the (inverse) cumulative distribution of the NPV of profits.

Myerson_C09.qxd 07-09-2004 09:46 Page 375

**(c)** Using your simulation data, also estimate the conditionally expected value of the NPV of profits given that first-year sale revenue is more than $4 million.

**(d)** From your simulation data, make a scatterplot that shows the observed relationship between first-year sales revenue and the 10-year NPV of profits.

**7.** PG Corporation stock is currently selling for $140 per share. In recent years, annual logarithmic growth rates of PG stock have had an average of 0.18 and a standard deviaion of 0.14. So suppose that the future growth of PG stock value will be a log-Brownian motion with annual log-drift 0.18 and log-volatility of 0.14. Consider an option to buy one share of PG stock two years from now for $180.

**(a)** Asssuming the statistical probabilities described, estimate the probability that this option will have a positive value two years from now, and estimate the expected monetary value of this option two years from now. Make a chart showing the (inverse) cumulative probability distribution for the value of this stock option two years from now.

**(b)** The option may be exercised at any time in the next two years, but financial theorists have argued that such options should always be held to the latest exercise date. Assuming the statistical probabilities described, estimate the probability of the event that PG stock will be below $180 two years from now but PG stock will have closed above $180 at the end of at least one of the next 24 months.

**8.** The corner grocery store has four cash registers where customers can be checked out after shopping, but there are only three regular checkout workers. Customers arrive at checkout with groceries as a Poisson arrival process with a mean of 60 per hour. The time to serve each customer at checkout is an independent Exponential random variable with mean 3 minutes. The store's policy is that if there are more than 12 customers waiting to check out, including those who are currently in service, then the store manager goes to the fourth cash register and begins working to check people out, until the number of people waiting is reduced to 12 again. Simulate a 10-hour day at this store, starting with no customers waiting. In your simulation, what was the average number of waiting customers? What fraction of the time did the manager work at checkout? What was the maximum number of waiting customers at any point in the day? (*Note:* during 10 hours, there should be approximately 1200 events, including about 600 arrivals and about 600 service completions.)

**9.** For the inventory problem in Section 9.8, try some other inventory policies. Would you recommend changing the inventory level at which the shop orders new televisions, or the number of televisions that the shop would order at the reorder level? For the best policy that you can find, estimate the long-run average cost per day of shipping, of lost sales, and of holding inventories.

Myerson_C09.qxd 07-09-2004 09:46 Page 376

Myerson_Appendix.qxd 07-09-2004 09:49 Page 377

# Appendix: Excel Add-Ins for Use with This Book

## Accompanying Software: Simtools and Formlist

Simtools.xla and Formlist.xla are add-ins for Microsoft Excel (version 5 and later). Simtools adds statistical functions and procedures for doing Monte Carlo simulation and risk analysis in spreadsheets. Formlist is a simple auditing tool that adds procedures for displaying the formulas of any selected range. Also available here is Torndiag.xls, an Excel workbook for making "tornado" sensitivity-analysis diagrams. Simtools and Formlist were developed to facilitate applications of probability and statistics and to help managers do complex decision analysis. These programs are distributed as freeware for individual use and may be freely redistributed to students and faculty in an academic institution. All other rights are reserved.

## Installing Simtools and Formlist

To install Simtools and Formlist, copy the Simtools.xla and Formlist.xla files from the accompanying disk or from the Internet, and save these files in the macro library subfolder of your Microsoft Office folder. The path for this library folder is often

C:\Program Files\MSOffice\Office\Library

This folder should already contain other .xla files that were automatically installed by Excel (including as ANALYSIS.XLA), and so you can find this add-in library folder by searching your computer for other *.xla files.

Then in Excel, select the Simulation Tools and Formula List options from the Tools:AddIns dialog box. After installation, SimTools and Formula List should appear as options directly under the Tools menu in Excel.

Myerson_Appendix.qxd 07-09-2004 09:49 Page 378

## Managing Problems with XLA Files

### Download problems

When saving .xla files from the Internet, you should make sure that each file's name has the correct .xla extension that signifies an Excel add-in, because Internet Explorer sometimes tries to change .xla extensions to .xls. If your browser changes the filename extension of Simtools.xla and Formlist.xla to .xls, you should be able to use these files after renaming them to restore the correct .xla filename. If your browser launches Excel instead of saving the XLA files to disk, you should right-click on the above links and select the "Save as . . ." option from the pop-up menu.

### Problems with add-in functions in spreadsheets

Excel may sometimes fail to recognize functions from an open .xla add-in file when these functions appear in an .xls workbook file that was made on another computer. The problem is that Excel is looking for a copy of the .xla file in the location where it was kept in the computer that made the workbook. The solution is to reconfigure the workbook's add-in links by using Excel's Edit-Links procedure. The following is a description of how to apply this procedure when you have this problem with workbook files that use functions provided by Simtools.xla.

To begin this Edit-Links procedure, the Simtools.xla add-in should already be open in Excel. (To check, look for "SimTools" in Excel's Tools menu.) The active workbook should be an .xls workbook file that was made on another computer. The problem is that formulas in this workbook that contain Simtools functions are getting "#NAME?" error values on your computer.

In Excel's Edit menu, select the Links option, which launches a dialog box called Links. In the Links dialog box, make sure that the source file Simtools.xla is selected, and then click the Change Source button, which launches another dialog box called Change Links. In the Change Links dialog box (which resembles a standard dialog box for finding files), navigate to the folder that contains your open copy of the Simtools.xla file (typically, in a folder named Library), and select this Simtools.xla file. Then click OK to close the Change Links dialog box. Finally, click OK to close the Links dialog box.

The functions in the workbook should now work on your computer. You may want to save the reconfigured workbook file now.

### Using Simtools features in other VBA programs

If you want to use Simtools functions and macros in other VBA programs, it is necessary to attach Simtools.xla as a reference in your VBA module, by

Myerson_Appendix.qxd 07-09-2004 09:49 Page 379

applying the Tools: References menu command in the Visual Basic Editor and checking Simtools.xla as an available reference.

### Source code

Simtools and Formlist are open code software. Their VBA source code is available on the accompanying disc in an Adobe PDF file called tools.pdf.

## Features of Simtools

Simtools.xla adds to Excel the following 3 macro procedures and 32 statistical functions, which are listed here in six categories. Optional parameters are shown here in *italics*.

### Inverse cumulative probability functions

These functions, along with Excel's NORMINV function, can be used for generating random variables when the first parameter (called "probability" or "randprob") is generated by a RAND( ) or CORAND function.

- BETINV(probability, mean, stdevn, *lowerbound, upperbound*) returns inverse cumulative values for a beta random variable, parameterized by its mean and standard deviation. When the first parameter is a RAND( ), BETINV yields a bounded random variable. (Default lower and upper bounds are 0 and 1.)
- BINOMINV(probability, *n*, *p*) returns inverse cumulative values for a binomial random variable. When the first parameter is a RAND( ), BINOMINV yields a bounded integer random variable between 0 and *n*, with mean *n*∗*p*.
- DISCRINV(randprob, values, probabilities) returns inverse cumulative values for a discrete random variable. When the first parameter is a RAND( ), DISCRINV returns a discrete random variable with possible values and corresponding probabilities in the given ranges.
- EXPOINV(probability, mean) returns inverse cumulative values for an exponential random variable. When the first parameter is a RAND( ), EXPOINV yields a nonnegative random variable (often used for random waiting times).
- GAMINV(probability, mean, stdevn) returns inverse cumulative values for a Gamma random variable, parameterized by its mean and standard deviation. When the first parameter is a RAND( ), GAMINV yields a nonnegative random variable.

Myerson_Appendix.qxd 07-09-2004 09:49 Page 380

- GENLINV(probability, quart1, quart2, quart3*, lowest, highest*) returns inverse cumulative values for a Generalized Lognormal random variable that has 25% probability below the quart1 value (the top of the first quartile), 50% probability below quart2, and 75% probability below quart3. A Generalized Lognormal random variable is a constant plus or minus a Lognormal random variable. When the first parameter is a RAND( ), GENLINV yields a random variable that could be positive or negative, but is bounded on the side of the narrower quartile range. If optional lowest and highest values are specified (satisfying lowest < quart1 < quart2 < quart3 < highest), then values of the Generalized Lognormal random variable are adjusted as necessary to keep GENLINV within these bounds (increasing to the lowest value from below it, decreasing to the highest value from above it).
- LNORMINV(probability, mean, stdevn) returns inverse cumulative values for a Lognormal random variable, parameterized by its mean and standard deviation. When the first parameter is a RAND( ), LNORMINV yields a nonnegative random variable.
- POISINV(probability, mean) returns inverse cumulative values for a Poisson random variable. When the first parameter is a RAND( ), POISINV yields a nonnegative integer random variable.
- TRIANINV(probability, lowerbound, mostlikely, upperbound) returns inverse cumulative values for a random variable with a triangular probability density. When the first parameter is a RAND( ), TRIANINV yields a bounded random variable.
- XTREMINV(probability, mean, stdevn) returns inverse cumulative values for an extreme-value (or Gumbel) random variable, parameterized by its mean and standard deviation. When the first parameter is a RAND( ), XTREMINV yields a random variable that may be positive or negative. (If $W$ is a Weibull random variable, then $-\text{LN}(W)$ has this extreme-value distribution.)

### Functions for working with correlations among random variables

- CORAND(CorrelArray*, RandSource*), entered as an array formula in a range of cells in a row, returns RANDom values for making random variables that have correlations as in the given CorrelArray. (See also NORMIZE.) Each value in an array returned by CORAND is like a RAND( ) in that it is generated uniformly between 0 and 1, but values in a CORAND array are not independent. For any integer $n$ greater than 1, the CorrelArray parameter may be a square $n \times n$ array of correlations for $n$ random variables (as returned by MCORRELS), and then CORAND returns an array of $n$ correspondingly

Myerson_Appendix.qxd 07-09-2004 09:49 Page 381

correlated values. Such an $n \times n$ correlation array should be symmetric and must have 1s on the diagonal. The CorrelArray parameter can also be a single number, in which case CORAND functions as if the CorrelArray parameter were a $2 \times 2$ array, returning two random values with the given correlation. When CorrelArray is a number and the optional RandSource parameter is a reference to another cell that contains a RAND( ) or CORAND formula, the CORAND function in one cell returns a uniform random value that is correlated with the RandSource cell according to the CorrelArray number. (For more about CORAND, see the following notes.)

- MCORRELS(dataRange) returns the matrix of correlation coefficients among the columns of the data range. If the data range has $n$ columns, then MCORRELS should be entered as an array formula in a square $n \times n$ range.
- MIDRAND(correlation, givenCoValue) returns the conditional median of a CORAND given another CORAND's value and the correlation between them. Used for subjective assessment of correlations.
- MSQRT(squarearray) returns a lower-triangular matrix square root (or Cholesky factor) of a given square array. It is an array function to be used in a range that is the same size as the given square range.
- NORMIZE(datacolumn) returns an array of normalized rank values, sampled from the standard normal distribution (at fractile medians) and rank-ordered as the data column. When CORANDs are used to make continuous random variables that are not normal, the correlation parameters of CORAND should be normalized rank correlations, which can be estimated from data by applying NORMIZE to each data series and then computing the correlations among these normalized arrays. (*Note:* NORMIZE may return an error message when it is entered into a range that is not the same size as the data column. But a calculation bug in Excel 97 may also cause this error message when a NORMIZE array is recalculated. In this case, the spreadsheet should be recalculated using the Ctrl+Alt+F9 keystroke, which tells Excel to recalculate all cells.)
- PRODS(values) multiplies each pair of values in the given range and returns the products as a square array. The values must be given in one row or one column. To illustrate the use of this function, suppose that a range named "correls" contains a square symmetric array that lists the correlations of the random returns-per-share of various stocks, a range named "stdevns" lists the standard deviations of these stock returns, and a range named "shares" lists the numbers of shares of these stocks in some investment portfolio; then the standard deviation of the total returns of the portfolio is

  SUMPRODUCT(PRODS(shares),PRODS(stdevns),correls)^0.5

Myerson_Appendix.qxd 07-09-2004 09:49 Page 382

### Functions for decision analysis

- ARGMAX(labels, values, *testCells*, *criterion*) returns a label corresponding to a maximal value, where the corresponding test cells (if any) match the criterion.
- CE(incomes, RiskTolConst, *RiskTolSlope*) returns the certainty equivalent of a random draw from the incomes range for a decision maker with constant risk tolerance (or linear risk tolerance when the optional RiskTolSlope parameter is used). When the risk tolerance is a positive number, the *CE* value is between the minimum and the average of the income values and the *CE* becomes closer to the average as the risk tolerance becomes larger. If the risk tolerance is a negative constant (denoting risk-seeking behavior), then the *CE* value is between the average and the maximum of the income values. If the RiskTolConst parameter is exactly equal to 0 and the RiskTolSlope is omitted, then the CE function returns the average of the income values. (The *CE* is computed by converting the incomes to utility values by UTIL, averaging these utility values, and converting this average utility back into monetary units by UINV. Nonnumerical entries in the incomes range are ignored. The *CE*'s dependence on the risk tolerance is discontinuous when the risk tolerance is zero, because a slightly lower risk tolerance yields a *CE* near the maximal income and a slightly higher risk tolerance yields a *CE* near the minimal income. If the RiskTolSlope parameter is used, then a zero or negative risk tolerance generates an error.)
- RISKTOL(HighIncome, LowIncome, CertainEquiv) returns the constant risk tolerance such that a lottery paying either the high or low income, each with probability 1/2, has the assessed certainty equivalent value.
- UTIL(income, RiskTolConst, *RiskTolSlope*) returns a utility value of monetary income for a decision maker with constant or linear risk tolerance.
- UINV(utility, RiskTolConst, *RiskTolSlope*) returns the monetary certainty equivalent of an expected utility from the UTIL function with the same risk-tolerance parameters.

### Functions for analyzing discrete probability distributions

- CEPR(values, probabilities, RiskTolConst, *testCells*, *criterion*) returns the certainty equivalent, for a decision maker with constant risk tolerance, of a random income drawn from the specified values according to the corresponding probabilities, conditional on the event where the corresponding test cells (if any) match the criterion. When the RiskTolConst parameter is 0, the CEPR function returns the conditional expected value.

Myerson_Appendix.qxd 07-09-2004 09:49 Page 383

- CORRELPR(values1, values2, probabilities) returns the correlation coefficient for a discrete probability distribution with corresponding values of two random variables.
- COVARPR(values1, values2, probabilities) returns the covariance for a discrete probability distribution with corresponding values of two random variables.
- STDEVPR(values, probabilities) returns the standard deviation for a discrete probability distribution with corresponding values of a random variable.

### Functions for regression analysis

- REGRESSN(XDataRange, YDataRange) actively returns multiple regression output for predicting *Y* as a random variable that depends linearly on the explanatory *X* variables. REGRESSN should be entered as an array formula in a range with seven rows and as many columns as the *X* data range.
- YHATSTE(XDataRange, NewXRow, RegressnStdErr) returns the standard error of the estimated conditional mean of *Y* (often called *Y*-hat) at the given new *X* row in a multiple regression.

### Functions for randomly generating discrete distributions

- DIRICH(alphaArray, *RandSource*), entered as an array formula in a range that has the same size as the alpha array, returns Dirichlet random fractions that sum to 1. These random fractions have means that are proportional to the corresponding values in the alpha array. Higher alpha-array values yield less variance. (RANDom inputs to drive the random DIRICH values can be provided externally in an optional RandSource range that is the same size as the alpha array. Otherwise, DIRICH gets random inputs automatically by implicit calls to Excel–VBA's random number generator.)
- DIRALPHA(dataRange), entered as an array formula in a range of cells in one row, returns estimated alpha parameters for a Dirichlet distribution. Each row of the data range should contain nonnegative numbers that sum to 1.
- LGT(x) is the logistic distribution. It transforms random variables from a logit model to fractions between 0 and 1 or (as an array formula) to discrete probability distributions proportional to the EXP($x(i)$) values. In the array-formula usage, $x$ must be an array of numbers in a row, and LGT must be entered into a similar array.
- LGTINV(p) is the inverse of the logistic or logit function. Applied to a fraction or (as an array formula) to a discrete probability distribution, LGTINV returns log-odds ratios for a logit model. In the

Myerson_Appendix.qxd 07-09-2004 09:49 Page 384

array-formula usage, *p* must be an array of probability values in a row, and LGTINV must be entered into a similar array.

- SHUFFLE(n, *RandSource*), entered as an array formula in a range of *n* cells in one row, returns a random ordering of the numbers from 1 to *n*. When entered into a row range of fewer than *n* cells, this function generates random samples from {1, . . . , *n*} without replacement. The values in a given range of *n* cells in a row can be shuffled by entering the array formula =INDEX(givenrange,1,SHUFFLE(*n*)) into another *n* cells in a row. (An optional RandSource cell containing a RAND( ) value can be used to determine the outcome of SHUFFLE. When RandSource is omitted, SHUFFLE uses instead an implicit call to Excel–VBA's random number generator.)

## Macro procedures added by Simtools to the Excel Tools menu

- SIMULATION TABLE, in a selected range, tabulates outputs from repeated recalculations of a Monte Carlo simulation model. The outputs to be tabulated should be in the top row of the selected range, but the top-left cell of this selected range should be unused. Recalculated values of the simulation outputs will fill the lower rows of the selected range, with each row containing the output values from an independent recalculation of the simulation model. The left column of the selected range is used for a percentile index, which can be useful for making cumulative-distribution charts after the output data is sorted (but the Simulation Table procedure itself does not sort the output data).
- ITERATIVE PROCESS iteratively copies values to a state range from an update range, while tabulating output.
- COMBINE ROWS makes all combinations of rows from selected ranges. It can be used to make a table of possible event sequences that is equivalent to a decision tree.

## More technical notes about CORAND

With an $n \times n$ CorrelArray parameter, CORAND's optional RandSource parameter can be a reference either to a single cell or to *n* cells in a row. RandSource cells should contain independent uniform random values, as generated by RAND or CORAND. When RandSource is a single cell, this cell's value is returned in CORAND's output array as the last value on the right. When RandSource is a range of *n* cells in a row, CORAND looks in these cells for all the random inputs that it needs to generate its *n* correlated outputs, preserving the rightmost value but transforming the others to generate appropriately correlated random outputs. This usage (along with a similar RandSource usage

Myerson_Appendix.qxd 07-09-2004 09:49 Page 385

for the DIRICH and SHUFFLE function) can be applied when you want to work with another simulation add-in's random number generator. When RandSource is omitted, CORAND automatically generates its *n* random values by transforming the results of *n* implicit calls to Excel–VBA's random number generator. The seed for this VBA random number generator is randomly reset (with a reference to the computer's clock) whenever the SimulationTable or IterativeProcess macro procedure is run.

## Features of Formlist

Formlist.xla adds one macro procedure to the Excel Tools menu and adds two Lookup-Reference functions that can be used in auditing formulas.

- FORMULA LIST is a procedure on the Tools menu for auditing the formulas that are used in a selected range. Under this procedure, the formulas from all cells in the audit range can be listed as text in any selected column. If a cell is part of an array formula, then the range of this array is listed with the formula (new in v1.5). The list can also include the names that are defined in the workbook (new in v1.4). If the selected output range is the same as the audit range, then FORMULA LIST instead does a formula/text toggle, inserting a single quote ['] before the initial equal sign [=] in each formula, to display formulas in place as text. A second application of this formula/text toggle removes these quotes to restore the formulas. Array formulas are not affected by a formula/text toggle.
- FORMULAS(auditRange) is a function that actively displays the current formulas in the audit range as an array in a column. To display all formulas, the FORMULAS function should be entered as an array formula into a range in one column that includes more cells than there are formulas in the audit range. The words "FORMULAS FROM RANGE . . ." should appear below the last displayed formula.
- FORMRC(auditCell) is a function that returns the formula of the designated cell in R1C1-style notation (with braces on array formulas). If the cell does not contain a formula, then FORMRC returns the cell's value (or a zero-length string for a blank cell).

## Features of Torndiag.xls

Torndiag.xls is a normal Excel workbook file, not an add-in. It can be stored in any disk or folder where you keep Excel workbook (.xls) files. When you open this file in Excel, it adds a TornadoDiagram procedure to the Excel

Myerson_Appendix.qxd 07-09-2004 09:49 Page 386

Tools menu. This procedure can then be used to make a tornado-style sensitivity-analysis diagram in any open workbook. (Tornado diagrams show how an output value would change as various input parameters are changed, one at a time, from a given best estimate to a given low estimate and a given high estimate.) Instructions for using this TornadoDiagram procedure, along with an illustrative example, can be found in the Torndiag.xls workbook itself.

## Excel Spreadsheets for Chapters 1 through 9

The disk accompanying this book also includes spreadsheet files for making the figures that are shown in each chapter of the book: chapter1.xls, chapter2.xls, chapter3.xls, chapter4.xls, chapter5.xls, chapter6.xls, chapter7.xls, chapter8.xls, chapter9.xls.

Myerson_BibNote.qxd 07-09-2004 09:50 Page 387

# Bibliographic Note

For a more sophisticated development of the mathematical ideas of probability and utility that are used in this book, the reader may consult any advanced text on probability models and financial mathematics. I would particularly recommend De Groot's *Optimal Statistical Decisions* and Luenberger's *Investment Science*.

Within the field of decision analysis, I have focused here almost exclusively on simulation analysis, as the most versatile tool for handling realistic complexity. For a broader introduction to other methods of decision analysis, the reader may consult Robert Clemen's current textbook *Making Hard Decisions* and Howard Raiffa's classic *Decision Analysis*. Other authors who have led the way in exploring the use of spreadsheets for decision analysis include Cliff Ragsdale, Sam Savage, David Vose, Wayne Winston, and Christian Albright. David Vose's *Risk Analysis* may be particularly comparable to this book in its presentation of sophisticated probability models for applied decision analysis.

There are many good general references on how to use Microsoft Excel. In particular, I would highly recommend John Walkenbach's *Excel Bible*, in any of its versions.

## References

Albright, S. Christian, *VBA for Modelers*, Pacific Grove, CA: Duxbury, 2001.

Clemen, Robert T., *Making Hard Decisions* (2nd ed.), Belmont, CA: Duxbury, 1997.

De Groot, Morris H., *Optimal Statistical Decisions*, New York: McGraw-Hill, 1970.

Luenberger, David G., *Investment Science*, New York: Oxford University Press, 1998.

Myerson_BibNote.qxd 07-09-2004 09:50 Page 388

Ragsdale, Cliff T., *Spreadsheet Modeling and Decision Analysis* (3rd ed.), Cincinnati: South-Western, 2000.
Raiffa, Howard, *Decision Analysis* Reading, MA: Addison-Wesley, 1968.
Savage, Sam L., *Insight.xla*, Pacific Grove, CA: Brooks/Cole, 1998.
Vose, David, *Risk Analysis* (2nd ed.), Chichester, UK: Wiley, 2000.
Walkenbach, John, *Excel 2000 Bible*, New York: Wiley, 1999.
Winston, Wayne L., and S. Christian Albright, *Practical Management Science* (2nd ed.), Pacific Grove, CA: Duxbury, 2000.

Myerson_Index.qxd 09-09-2004 17:51 Page 389

# Index


## A

Absolute reference, 6
Add-in functions. *See* Excel add-ins
Adjusted average, 225
Adjusted average squared error, 225, 226
Akerlof, George, 255
Analysis ToolPak, 223
ANALYSIS.XLA, 377
AND, 38, 70
Annual growth ratio, 114
Arbitrage opportunity, 317
Arbitrage pricing theory, 75, 317–322
ARGMAX, 382
Array formula, 160
Asset pricing models, 305–317
@Risk, 2
Autocorrelation, 336
AVERAGE, 13, 63

## B

Backward analysis, 211–214
Bates Restoration, 251
Bayes' rule, 31
Beliefs, 203
Bell-shaped curve, 113
Beta distribution, 201
BETAINV, 139, 201
Beta random variable, 138–140
BETINV, 139, 201, 379
Bidding problems, 251–271. *See also* Decision analysis
Bidding to buy, 265
Bidding to sell, 252, 256, 273
BINOMDIST, 34–36
Binomial distribution, 34
Binomial random variable, 144–145, 201, 354–355
BINOMINV, 144, 201–202, 379
Branch, 27
Brown, Robert, 339
Brownian motion, 339–340
Brownian-motion growth models, 339–343

## C

Capital asset pricing model (CAPM), 317
Cases
  Bates Restoration, 251
  Global Sports, 148, 331
  Lepton Inc., 147
  Scotia Snowboards, 232
  Superior Semiconductor, 48–49, 82, 108
  Valdez Wilderness, 264
CE, 75, 93, 382
Central limit theorem, 68–71
CEPR, 75, 90, 382

Myerson_Index.qxd 09-09-2004 17:51 Page 390

Certainty equivalent, 83, 101, 105
Certainty equivalent of normal lottery, 135–137
Cholesky factorization, 193
Coin tossing, 3–7
Column-input data table, 15, 39
COMBINE ROWS, 39, 154, 384
Competitive behavior, 263–271
Conditional expectation, 199–230
  backward analysis, 211–214
  conditional expected value, 204
  correlation, 214–216
  dependence among random variables, 199–202
  estimating, 203–206
  expectation tree, 211–212
  expected-posterior law, 206–211
  linear regression models, 220–224
  regression analysis/least-squared error, 224
  uncertainty about a probability, 216–220
  what is it, 199
Conditional expected value, 204, 208
Conditional independence, 19
Conditional median, 184
Conditional probability, 14, 17, 29, 31
Confidence interval (95%), 93–94
Constant risk tolerance, 84–91, 101–104
Constraint
  moral hazard, 295–301
  Solver, 177
  time, 366–370
Continuous probability distribution, 109
Continuous random variable, 20, 107–150
  beta random variable, 138–140
  binomial random variable, 144–145
  certainty equivalent of normal lottery, 135–137
  discrete and continuous unknowns, 131–135
  exponential random variable, 140–141
  extreme-value random variable, 143
  gamma random variable, 141–142
  generalized-lognormal distributions, 122–124
  lognormal distributions, 116–122
  nonparametric distribution fitting sample data, 143–144
  normal distributions, 108–113
  Poisson random variable, 145
  probability density, 111–113
  subjective assessment, 124–131
  triangular random variable, 138
  uniform random variable, 137–138
  Weibull random variable, 143
Continuous sharing rules, 291
Copy, 4–5
CORAND, 167–171, 186, 380–381, 384–385
Corporate decision making, 305–317
CORREL, 162–166
Correlation, 157–158
  conditioned expectation, 214–216
  equation, 157
  estimating, 162–166
Correlation array, 160
Correlations, 380–381
CORRELPR, 157, 383
COUNT, 13
COUNTIF, 211
Covariance, 156–157
COVARPR, 157, 383

Myerson_Index.qxd 09-09-2004 17:51 Page 391

Criterion of expected value maximization, 71
Crystal Ball, 2
Cumulative probability, 51
Cumulative risk profile, 74

## D

Data:Group, 219
Data:Sort, 65
Data table, 11
Data-Table command, 15–18
Decision analysis, 231–279, 382
  bidding problems, 251–271
  competitive behavior, 263–271
  corporation, 305–317
  decision tree, 246–250
  game theory, 264–265, 269–271
  methods of simulation analyses, 233–241, 271
  partnership, 282–305
  strategic use of information, 241–245
  winner's curse, 255–263
Decision node, 248
Decision tree, 246–250
Decision variable, 231
Degrees of freedom, 226
Dependence among random variables, 199–202
DIRALPHA, 383
DIRICH, 383
Discrete and continuous unknowns, 131–135
Discrete probability distributions, 50–51, 382–383
Discrete random variables, 47–79
  accuracy of sample estimates, 66–71
  central limit theorem, 68–71
  charting a probability distribution, 50–53
  cumulative probability, 51
  decision criteria, 71–75
  discrete probability distribution, 50–51
  estimates from sample data, 62–66
  expected value, 58–59, 62
  expected value criterion, 71–75
  inverse cumulative probability curve, 52–53
  multiple random variable, 75–77
  normal probability distribution, 66–71
  simulating, 53–57
  standard deviation, 60–62
  unknown quantities, 49–50
  variance, 60–61
DISCRINV, 55–56, 379
Distributions. *See* Probability distribution; Random variable
Dollar sign ($), 6
Download problems, 378
Drift, 341
Duration of project, 366–370
Dynamic models, 325–375
  Brownian-motion growth models, 339–343
  exponential arrival models, 349–356
  forecasting model, 328–339
  Goeing Aircraft case, 332–339
  inventory model, 362–366
  iterative process, 356–366
  log-optimal investment strategies, 343–349
  net present value (NPV), 326–328
  no-memory property, 349–351
  project length, 366–370
  queuing models, 356–362

## E

Economic order quantity (EOQ), 364
Edit-Links procedure, 378

Myerson_Index.qxd 09-09-2004 17:51 Page 392

EOQ, 364
Equal sign (=), 3
Equivalent models, 41
Excel add-ins
  Formlist, 3, 377, 385
  Simtools. *See* Simtools
  source code, 379
  Torndiag.xls, 377, 385–386
  troubleshooting, 378–379
Exhaustive, 27
EXP, 113–115
Expectation tree, 211–212
Expected-posterior law, 206–211
Expected utility, 88
Expected value, 58–59, 62
Expected value criterion, 71–75
Expected value of information, 244
Explanatory variable, 220
EXPOINV, 141, 379
Exponential arrival models, 349–356
Exponential function, 113
Exponential random variable, 140–141, 349
Extreme-value random variable, 143

## F

Fallacy of averages, 60, 240
Fast situations. *See* Cases
First-price auction, 252
Forecasting model, 328–339
Formal quantitative model, 40–42
Formlist, 3, 377, 385
FORMRC, 385
Formulaic dependence, 200
FORMULA LIST, 385
FORMULAS, 385
Future value of stock, 306–317

## G

Game model, 264–265
Game theory, 264–265, 269–271
GAMINV, 141, 379
GAMMAINV, 141
Gamma random variable, 141–142, 352
Generalized-Lognormal distribution, 122–124, 128–129
Generalized-Lognormal random variable, 122
GENLINV, 122–126, 380
Global Sports, 148, 331
GoalSeek, 86, 185, 262
Goeing Aircraft case, 332–339
Gumbel distribution, 143

## I

IF, 3
Increasing risk tolerance, 94
Independent, 5
INDEX, 39
Infeasible, 179
Information, strategic use of, 241–245
Initial equal sign (=), 3
Insert:Chart:XY-Scatter, 65
Insert-Function dialog box, 139
Insert:Function dialog box, 42, 201
Insight.xla, 2
Insurance. *See* Risk sharing and finance
Intersection, 29, 153
Intersection probability, 30
Intuition, 40–41
Inventory holding cost, 362
Inventory model, 362–366
Inventory policy, 363
Inverse cumulative probability curve, 52–53
Inverse cumulative probability functions, 379–380
ITERATIVE PROCESS, 384
Iterative process, 356–366

## J

Joint probability distribution, 153

Myerson_Index.qxd 09-09-2004 17:51 Page 393

## L

Label, 27
Lagrange multiplier, 320
Law of expected posteriors, 206–211
Law of iterated expectation (expected-posterior law), 206
Law of large numbers, 63
Least-squared error, 224
Length of project, 366–370
Lepton Inc., 147
LGT, 383
LGTINV, 383–384
Linear functions of random variables, 158–161, 190–193
Linear regression models, 220–224
Linear risk tolerance, 94–96
Linear sharing rules, 291
LN, 113–115
LNORMINV, 121, 380
Local optimal solution, 181
Local optimum, 181
Local risk tolerance, 95
Logarithmic growth rate, 114–115
Log-Brownian motion, 341
Log-drift, 341
Log-mean, 116
Lognormal distributions, 116–122
Lognormal random variable, 116
Log-optimal investment strategies, 343–349
Log-standard-deviation, 116
Log-volatility, 341
Lottery, 135–137

## M

Macro procedures, 384
Marginal probability distribution, 154
Market-adjusted probability, 75, 314
"Market for Lemons, The" (Akerlof), 255
Market portfolio, 306
Markov process, 358
MATCH, 210
Mathematically equivalent, 41
MAX, 266, 299
MCORRELS, 166, 381
Median, 122
Method of point estimates, 49
MIDRAND, 184–186, 381
MIN, 238, 253
Model, 40–42
Modular simulation, 262
Moral hazard, 295–305
Moral-hazard incentive constraints, 295–301
Morgenstern, Oskar, 84, 96
Multiple random variable, 75–77
Multivariate Normal, 166–171
Multivariate normal asset returns, 171–175
Multivariate Normality, 175
Multivariate Normal random variable, 191, 223
Mutually exclusive, 27

## N

Nash equilibrium, 269, 271
Natural logarithm function, 113
Net present value (NPV), 326–328
Newsboy problems, 241–250. *See also* Decision analysis 95% confidence interval, 93–94
Node, 27
No-memory property, 349–351
Nonlinear sharing rules, 291, 292, 301
Non-Normal random variable, 186–188
Normal error term, 221
Normality assumption, 175
Normalized rank correlation, 187–189
Normal lottery, 135–137

Myerson_Index.qxd 09-09-2004 17:51 Page 394

Normal probability distribution, 67–68, 108–113
Normal random variable, 109, 113
NORMINV, 67–68, 111–112, 168–169
NORMIZE, 188–189, 381
NORMSDIST, 109–111
NPV, 326–328

## O

Oil exploration example, 25–37
Optimization problems, 176

## P

Partnership, 282–305
Paste, 5
Pearson correlation coefficient, 157, 189
PERCENTILE, 23, 261
PERCENTRANK, 22–23
Personal utility function, 83, 84
Piecewise-linear sharing rules, 291–292, 301–305
POISINV, 145
Poisson distribution, 145
Poisson random variable, 145, 352–353
Portfolio analysis (multivariate normal asset returns), 171–175
POSINV, 380
Posterior beliefs, 203
Posterior probability, 14
Predecessors, 367
Principal-agent problems, 281. *See also* Risk sharing and finance
Prior beliefs, 203
Prior probability, 14
Probability
  conditional, 14, 17
  cumulative, 51
  defined, 7
  market-adjusted, 75, 314
  posterior, 14
  prior, 14
Probability density, 111–113
Probability distribution, 125. *See also* Random variable
  charting, 50–53
  discrete, 50–51
  joint, 153
  marginal, 154
  normal, 67–68, 108–113
  shadow, 321
  subjective, 125
Probability tree, 25–37
Problems with XLA files, 378
PRODS, 160–161, 381
Projected forecast, 328–338
Project length, 366–370
Proportional risk tolerance, 102
Publicly held corporations, 305–317

## Q

Quantitative model, 40–42
Quartile boundary points, 122
Quartile median, 247
Quartile median selection, 247
Queuing models, 356–362

## R

RAND, 3, 4, 43
Random generation of discrete distributions, 383–384
Random variable, 9
  beta, 138–140
  binomial, 144–145, 354–355
  continuous, 20, 107–150. *See also* Continuous random variable
  defined, 50
  dependence, 199–202
  discrete, 47–79. *See also* Discrete random variables
  exponential, 140–141, 349
  extreme-value, 143
  gamma, 141–142, 352

Myerson_Index.qxd 09-09-2004 17:51 Page 395

generalized-lognormal, 122
linear functions, 190–193
lognormal, 116
multivariate normal, 166–171, 191
non-Normal, 186–189
normal, 109
Poisson, 145, 352–353
standard normal, 110
stationary, 328
triangular, 20–21, 138
uniform, 137–138
Weibull, 143
Random walk, 340
Regression analysis, 220–227, 383
REGRESSN, 222, 383
Relative reference, 6
Residual error, 338
@Risk, 2
Risk averse, 73, 83, 101
Risk-free interest rate, 308–309
Risk neutral, 82–84
Risk premium, 83, 101
Risk sharing and finance, 281–324
arbitrage pricing theory, 317–322
asset pricing models, 305–317
CAPM, 317
continuous sharing rules, 291
corporation, 305–317
linear sharing rules, 291
moral hazard, 295–305
moral-hazard incentive constraints, 295–301
nonlinear sharing rules, 291, 292, 301
partnership, 282–305
piecewise-linear sharing rules, 291–292, 301–305
senior manager (corporation), 305–306
sensitivity report, 320–321
shirking, 299–300
stockholders, 306
Risk tolerance, 91
constant, 84–91, 101–104
increasing, 94
linear, 94–96
local, 95
proportional, 102
total, 286
Risk tolerance constant, 85, 87
RISKTOL, 75, 86–88, 382
Root, 27
ROUND, 145

## S

Sales calls example, 7–25
Sample standard deviation, 65
Scotia Snowboards, 232
Sealed-bid auction, 252
Senior manager (corporation), 305–306
Sensitivity analysis, 20, 40, 235
Sensitivity Report, 320–321
Shadow price, 320
Shadow probability distribution, 321
Sharing risk. *See* Risk sharing and finance
Shirking, 299–300
SHUFFLE, 384
Simtools, 2, 377
CORAND, 384–385
correlations, 380–381
decision analysis, 382
discrete probability distributions, 382–383
getting started, 2–3
installing, 377
inverse cumulative probability functions, 379–380
macro procedures, 384
other VBA programs, 378–379
random generation of discrete distributions, 383–384
regression analysis, 383
troubleshooting, 378
Simulation model, discrete random variables, 53–57

Myerson_Index.qxd 09-09-2004 17:51 Page 396

SIMULATION TABLE, 11, 17, 62–63, 384
Solver, 176–182, 301
Solver Sensitivity Report, 320–321
Source code, 379
Squared error, 225
Standard deviation, 60–62, 74
Standard error, 220
Standard Normal, 110
Stationary random variable, 328
Stationary value, 338
Statistical dependence, 200
Statistical regression software, 228
STDEV, 61, 64, 65
STDEVP, 61
STDEVPR, 61, 65, 383
Stock market, 306–317
Strategic use of information, 241–245
Strategy, 231, 243
Subjective assessment, 124–131
Subjective assessment of correlations, 182–186
SUM, 13
SUMIF, 211
SUMPRODUCT, 58, 159, 161, 381
Superior Semiconductor, 48–49, 82, 108

## T

Table with multiple inputs, 37–40
Time constraint (project length), 366–370
Time-series forecasting models, 328–338
TornadoDiagram, 385–386
Torndiag.xls, 377, 385–386
Total risk of tolerance, 286
Tree diagram
  decision tree, 246–250
  expectation tree, 211–212
  probability tree, 25–37
Triangular random variable, 20–21, 138
TRIANINV, 21, 138, 380
Troubleshooting, 378–379

## U

UINV, 89–90, 94, 382
Unbounded optimal solution, 180
Uncertainty about a probability, 216–220
Uniform random variable, 137–138
UTIL, 85, 382
Utility function, 75
Utility theory, 75, 81–106
  advanced technical notes, 96–104
  certainty equivalent, 83, 101
  constant risk tolerance, 84–91, 101–104
  linear risk tolerance, 94–96
  personal utility function, 83, 84
  probabilities, 82–92
  risk tolerance. *See* Risk tolerance
  simulation data, 92–94
  Von Neumann-Morgenstern thesis, 96–101

## V

Valdez Wilderness, 264
Value at risk, 74
Value of information, 219, 241–245
Variable
  decision, 231
  explanatory, 220
  random. *See* Random variable
Variance, 60–61

Myerson_Index.qxd 09-09-2004 17:51 Page 397

Vignettes. *See* Cases
VLOOKUP, 55, 243
Volatility, 341
Von Neumann, John, 84, 96
Von Neumann-Morgenstern thesis, 96–101

W

Waiting in queue (queuing models), 356–362
Weibull random variable, 143
What-if table, 15
Winner's curse, 220, 255–263

X

XLA files, 378
XTREMINV, 143, 380

Y

*Y*-hat, 383
YHATSTE, 383

Myerson_Index.qxd 09-09-2004 17:51 Page 398

Myerson_Index.qxd 09-09-2004 17:51 Page 399

Myerson_Index.qxd 09-09-2004 17:51 Page 400

Myerson_Index.qxd 09-09-2004 17:51 Page 401

Myerson_Index.qxd 09-09-2004 17:51 Page 402

Myerson_Index.qxd 09-09-2004 17:51 Page 403

Myerson_Index.qxd 09-09-2004 17:51 Page 404